Alliance Systems

The Move to PCI Express in Next Generation Systems

The next generation input/output (I/O) interconnect standard for the PC industry is PCI Express and PICMG 1.3 specifications. PCI Express extends the software interface and mechanical orientation of standard PCI but replaces its bandwidth throttling parallel bus with a serial architecture. PCI Express is a dual-transaction protocol that prioritizes packets to more specific targets, enabling a higher performance interface.

However, as platforms change to incorporate these new PCI standards, they are virtually eliminating space for high-end telephony boards. When building telephony system solutions, vendors need to be aware of the changes in these specifications so that they can better use current and future systems for communications applications.

PCI HITS THE CELING

Parallel Verses Serial Technology

The standard PCI implementations today, PCI 2.2 and PCI-X, are no longer able to support the increased bandwidth requirements of today's high throughput I/O devices and chipsets as well as future generations of I/O. At the core of all PCI specifications previous to PCI Express is a signaling technology that utilizes a multiple drop, shared bidirectional I/O bus that is at the end of its effective use in the modern PC platform. (See Figure A).

Figure A: PCI Shared Bidirectional I/O Bus Architecture

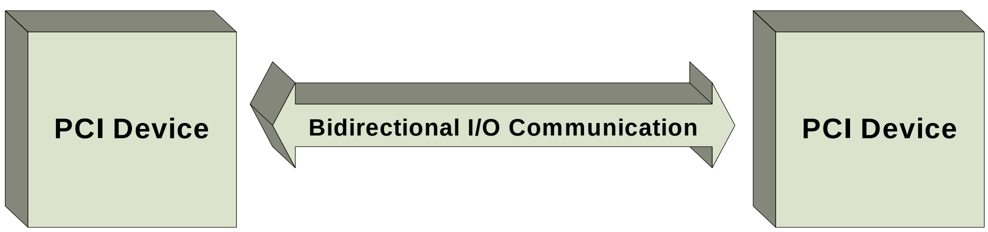

Greater performance requires higher frequencies and lower voltages. This presents a barrier for PCI because it must synchronize multiple signals with a bus clock for viable data transfer. Any deviations from the published specification caused by such things as signal load, voltage fluctuations, and impedance can cause signal skew. Skewed signals to devices and across multiple bus nodes can degrade the I/O performance of the PCI bus.

PCI Express defines a serial, self-clocking, packetized protocol that utilizes dual unidirectional I/O pathways for simultaneous communication between two devices instead of the bi-directional topology. (See Figure B).

Figure B: PCI Express Dual Unidirectional I/O Bus Architecture

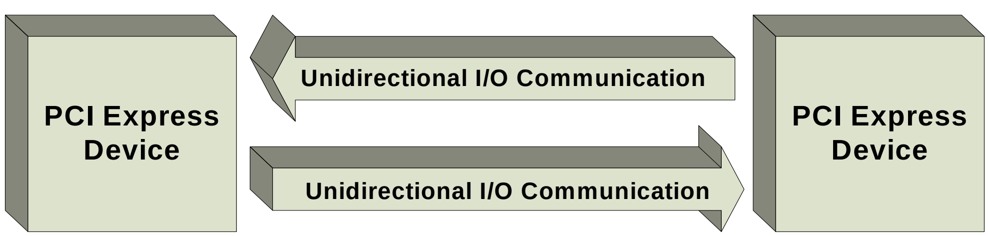

Lower Pin Count
PCI Express also implements lower pin counts and lower signaling voltages than PCI and has double the theoretical bandwidth. The PCI Express specification defines a number of different serial links per slot as x1, x2, x4, x8, x12, x16, and x32 (e.g. x8 = eight serial paths to and from the serial device). Because of this expandability, PCI Express is scalable to much higher throughputs than PCI.

Higher System Throughput
PCI has grown from 32-bit PCI at 33 MHz with a maximum throughput of 133 MB/sec to 64-bit PCI-X at 133 MHz with a maximum throughput of 1066 MB/sec and PCI-X 2.0 is on the way to supporting up to 4.3 GB/sec. PCI Express (PCI-E) will scale from x1 link at 500 MB/sec to a x32 PCI-E link at 16,384 MB/sec. (See Figure C). The theoretical limit of PCI-E generation 3 x32 links is calculated at over 60 GB/sec.

PCI Technology Summary	Maximum Theoretical Throughput (MB/sec)
PCI 32-Bit/33 MHz	133
PCI-X 64-bit/133 MHz	1066
PCI-X 2.0	4,300
PCI-E x1	500
PCI-E x32	16,384
PCI-E g3 x32	60,000

For the last few years, PCI-X has provided sufficient throughput for server applications such as Gigabit Ethernet network cards, Ultra320 (U320) SCSI and Fibre Channel host bus adapters and other high I/O bandwidth devices. For desktop applications, the highest throughput device has been video acceleration through AGP 8X. With network adapters moving into the 10 Gigabit Ethernet range and host bus adapter (HBA) technology continuing to scale higher with implementations such as Serial Attached SCSI (SAS), the bandwidth provided by PCI-X has become the bottleneck. The requirement for a separate, high-pin, high-throughput bus, like AGP just for video, increases the cost and complexity of the PC motherboard.

With PCI Express, applications and technologies such as video streaming, USB 2.0, IEEE 1394, Infiniband, and future generations of I/O will be supported within the PC architecture for years to come. The high-throughput support for current and future applications, reduced I/O architecture complexity and homogenization of bus designs makes PCI Express the natural successor to PCI.

Figure C: Throughput Comparison

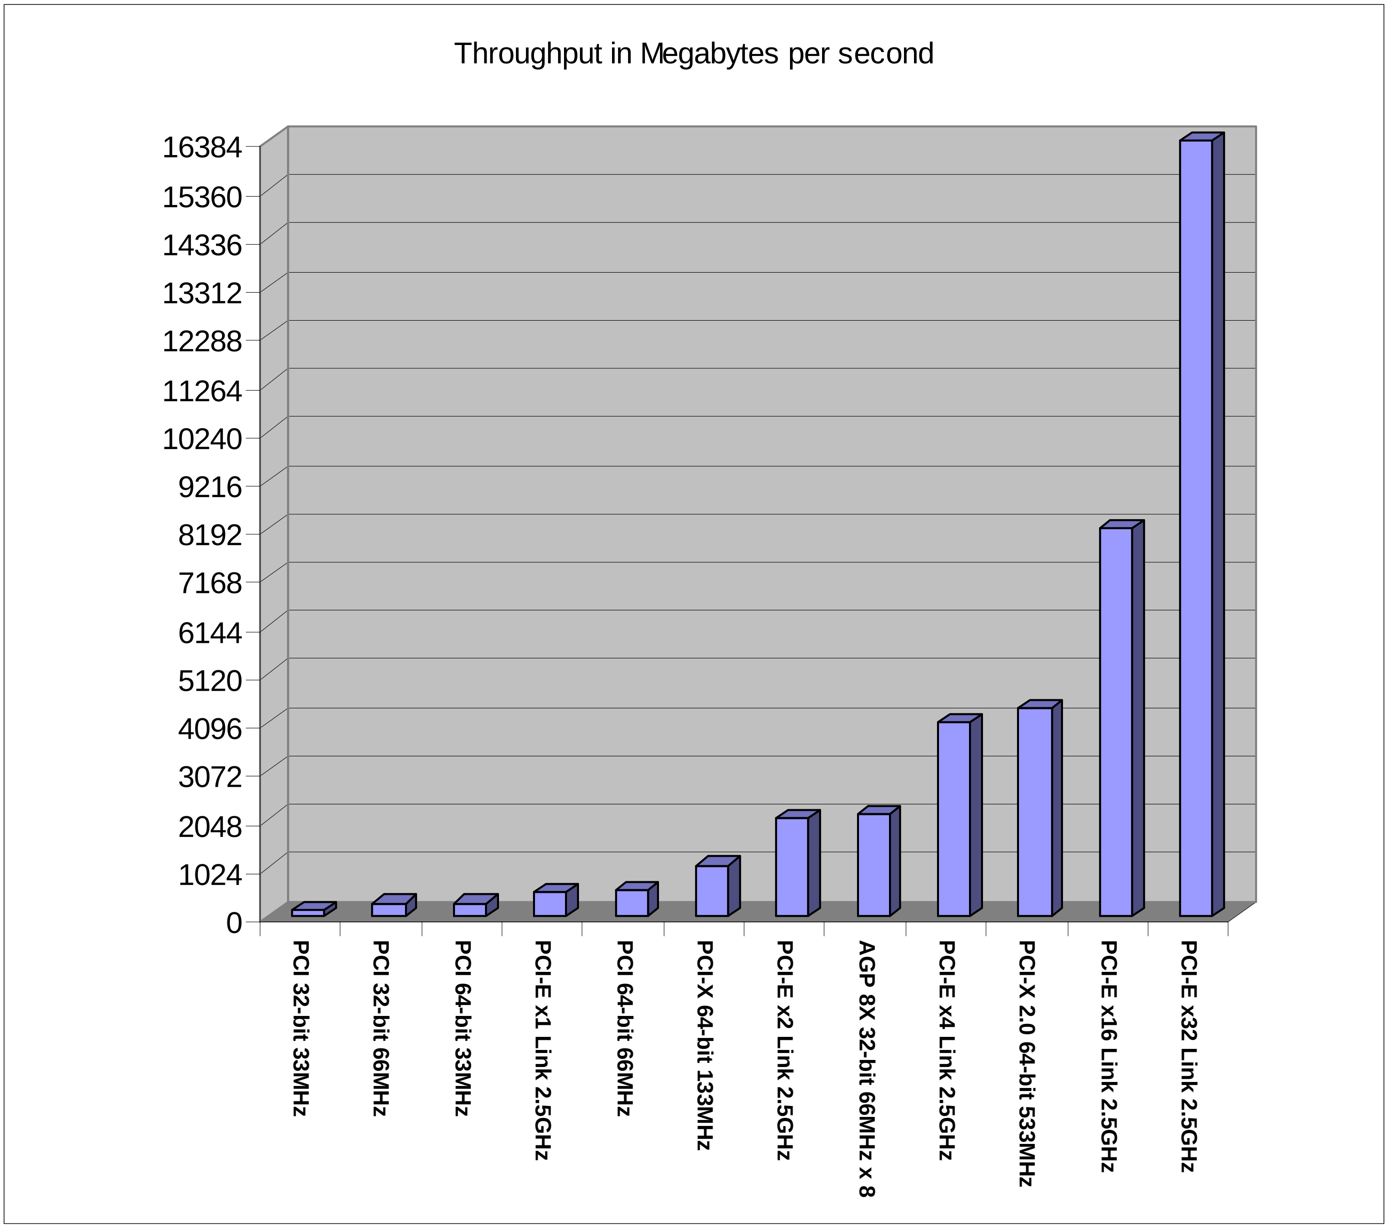

The one area where this increase in throughput will have the quickest impact on server performance is in the RAID HBA. Consider six U320 SCSI hard drives running in RAID 5. The theoretical maximum I/O throughput of this array is 1920 MB/sec. Standard PCI-X RAID HBAs are capable of pulling a maximum of 1024 MB/sec through the PCI-X slot. PCI Express RAID adapters would quickly remove this bottleneck from consideration. Figure D (below) shows a comparison of I/O performance between a U320 SCSI PCI-E RAID HBA and that of a U320 PCI-X 100 MHz RAID HBA measured in MB/sec for reads and writes of 2 MB files. (The test environment is six U320 SCA SCSI drives configured in a 292 GB RAID 5 array and running IOMeter.)

Figure D – I/O Comparison of PCI-E versus PCI-X in a U320 SCSI RAID 5

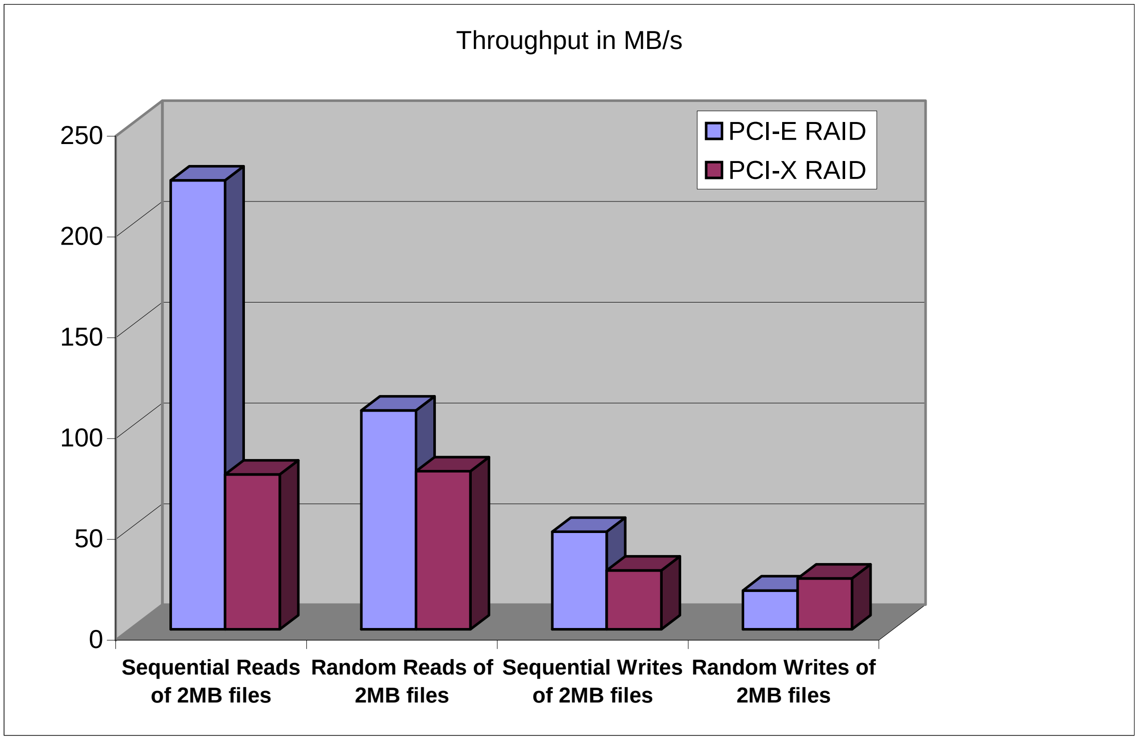

Depending on the type application as well as the way the HBA RAID controller is configured, you will typically see performance enhancements across multiple I/O profiles.

PCI EXPRESS UNIQUE CONNECTORS

PCI Express supports the same software interface as PCI as well as offering bridged PCI and PCI-X for legacy parallel technology. However, the standard connector for PCI Express is not mechanically compatible with prior versions of PCI. (See Figure E).

Figure E: Mechanical Comparison of PCI-E to PCI

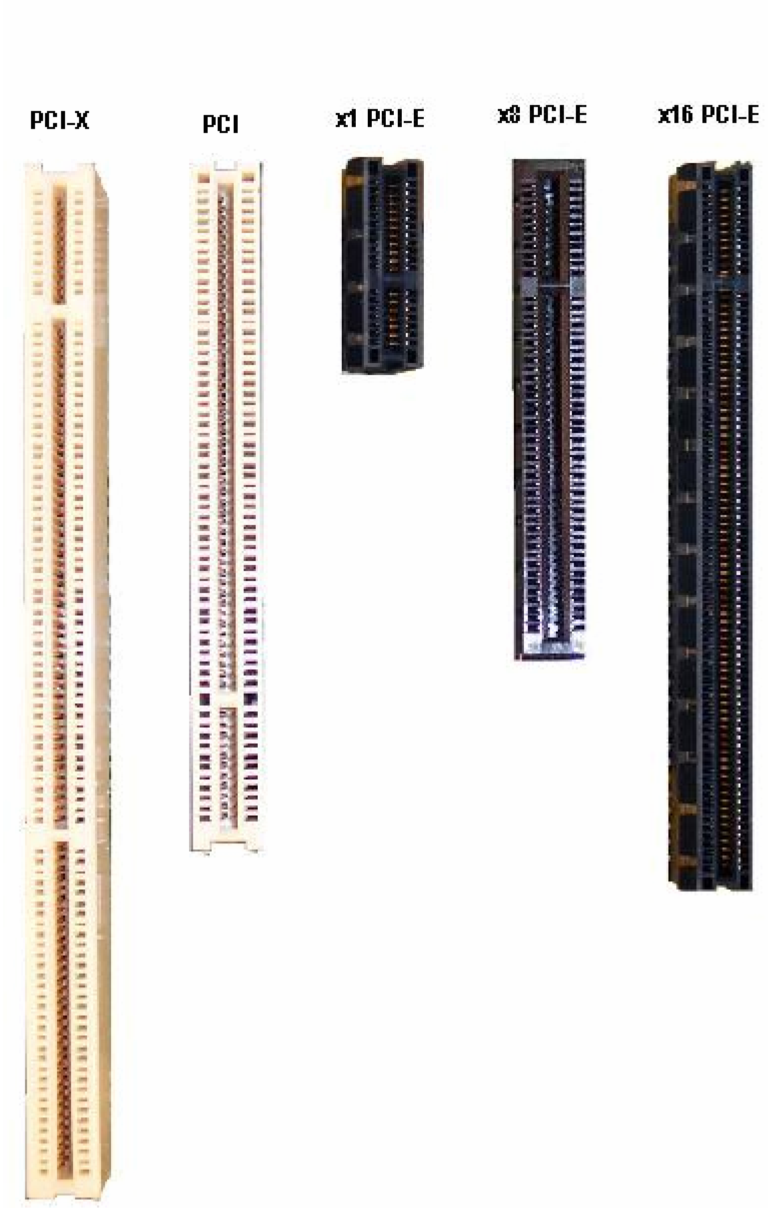

Specific Telephony Board Considerations
The adoption of PCI Express has a major impact on telephony applications that require +5V as well as PCI interfaces in general. Most new computing platforms use a combination of PCI Express x4, x8, and x16 slots in addition to PCI-X slots. There will be very few 32-bit +5V PCI slots available in the computing market for specific telephony boards.

Moving forward, most card technology companies will begin to produce native PCI Express products and so the number of standard PCI-X slots available on the latest computing platforms will become reduced in favor of true PCI Express slot interfaces.

EMBEDDED COMPUTING AND PASSIVE BACKPLANES

PICMG 1.3 is a specification that extends PCI Express to the passive backplane for embedded and industrial PC architectures. PICMG 1.3 replaces the standard 1.0 (ISA/PCI) specification by adding a PCI Express host slot and by replacing parallel bus interfaces to the backplane with x8 and x16 high-speed serial links. These links can be used for standard PCI-E x1, x4, x8, x16, and x32 slots or they can be bridged to legacy PCI and PCI-X interfaces. This specification also supports an additional 32-bit PCI/PCI-X to support bridgeless passive backplanes.

Specific Telephony Board Considerations
The passive backplane architecture is used primarily for high-density telephony applications with a high board count. With PICMG 1.3, the bandwidth is increased for the number of supported devices as well as extending the type of interfaces directly supported by the backplane technology. The PICMG 1.3 passive backplane provides for optional interfaces from the PCI Express host slot to devices such as SATA, USB 2.0, IPMB, SMBUS, and Power Management, which greatly increases the flexibility and extensibility of the passive computing platform.

Improved Power Distribution
PICMG 1.3 also improves power distribution. Since the PCI Express serial interface links use less traces as compared to the parallel bus of PICMG 1.0 systems, the PICMG 1.3 host boards have more edge contacts available to deliver system power. The total power delivery possible to a PICMG 1.3 host slot is double that of an equivalent full-size PICMG 1.0 Single Board Computer (SBC). (See Figure F.)

Figure F: Wattage Supported through the Host Slot

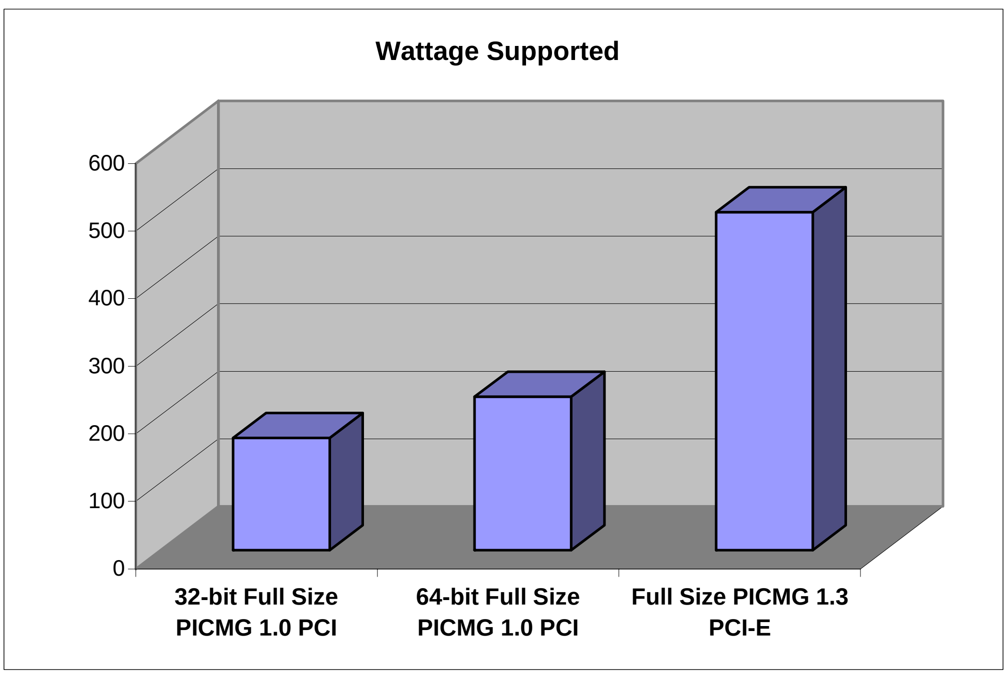

System Host Board Location and Connector Changes
The PCI Express SBC (now referred to as the System Host Board, SHB) has some mechanical changes that are noteworthy. The ISA/32-bit or 64-bit PCI edge connectors are replaced by a combination of x8 and x16 PCI Express connectors. (See Figure G.) Also, the component side of

a PCI Express SHB is the same as that of a standard PCI card but opposite from a PICMG 1.0 PCI/ISA SBC.

In general, most system manufacturers should place the SHB slot in location similar to that of 1.0 SBC slots to maintain some mechanical backward compatibility for passive backplane system enclosures. The PICMG 1.3 specification also details requirements for a half-size card with two PCI Express connectors on the SHB.

Figure G: Example PICMG 1.3 Backplane

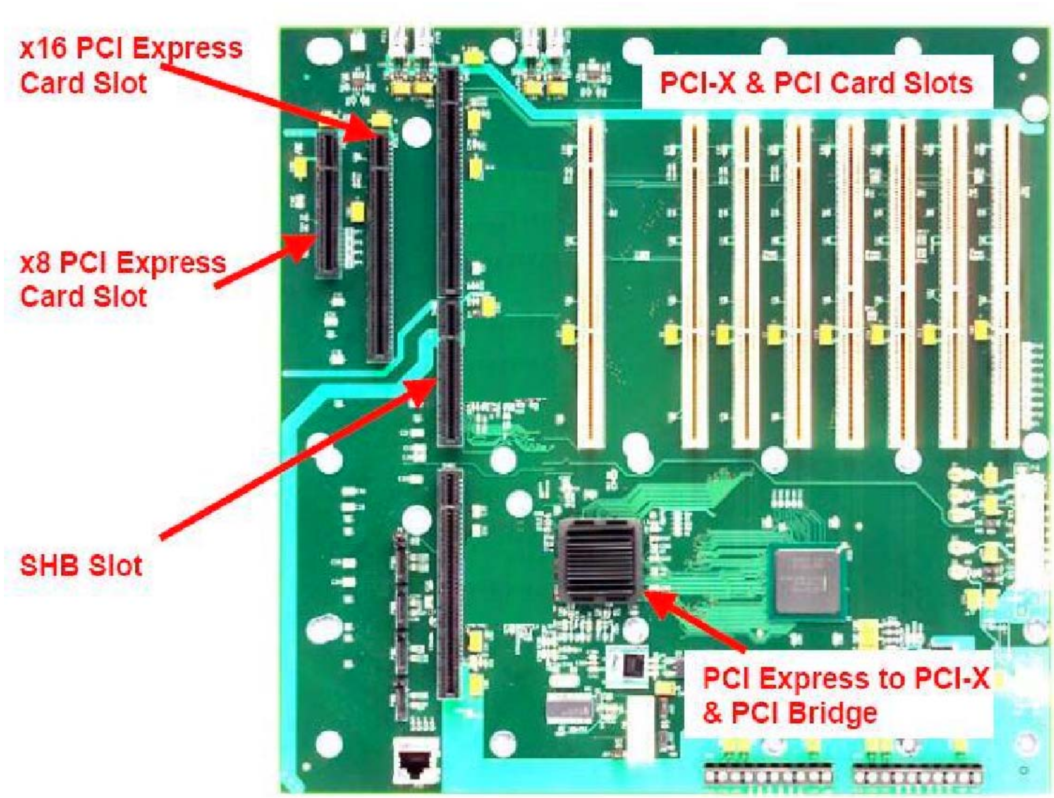

CONCLUSION

PCI Express and PCIMG 1.3 represent the next major leap in PC interconnect I/O and should support the throughput requirements of the PC industry for many years to come. The majority of new card technologies including storage, networking, and other I/O fabric devices will adopt PCI Express as the replacement for PCI and PCI-X. The PCIMG 1.3 specification will also become the new passive backplane interconnect for embedded and industrial PC applications.

BIBLIOGRAPHY

Introduction to PCI Express
A Hardware and Software Developer's Guide
Adam Wilen, Justin Schade, and Ron Thornburg
Intel Corporation

The new PICMG 1.3 specification
Bringing PCI Express to an SBC near you
Michael Bowling
Trenton Technology Inc.
Embedded Industry Product Review

Creating a Third Generation I/O Interconnect
Ajay V. Bhatt
Desktop Architecture Labs
Intel Corporation

www.picmg.org

www.trentonprocessors.com