



The PCI Express* Architecture and Advanced Switching

Interconnects for Converged
Computing and Communications

Table of Contents

- Executive Summary** 1
- Fewer Interconnects Simplify Convergence and Modular Communications Platforms** 1
 - Two Industry Trends Benefit Communications 1
 - Fewer Interconnects Needed 3
- Advancing Chip-to-Chip Communications: From Buses to Serial Interconnects and PCI Express* Architecture** 3
- The Widely Supported PCI Express* Architecture** 4
- Advanced Switching—A System Fabric Interconnect** 5
- Conclusion** 6
 - Benefits of PCI Express* Architecture 6
 - Benefits of Advanced Switching Based on PCI Express* Architecture 6
- For More Information** 6

Executive Summary

Computing and communications technologies have advanced along with changing markets, bringing these technologies together at the silicon, board and system levels. The compute industry has evolved around a standard chip-to-chip interconnect—the Peripheral Component Interconnect (PCI) specification—while communications and embedded solutions have incorporated multiple interconnect technologies for chip-to-chip and system fabrics, often proprietary. Multiple, diverged interconnect technologies present complexities and added cost in developing next generation embedded and communications systems.

The enormous benefits of a common chip-to-chip interconnect for board-level designs, and a single system fabric interconnect for system-level designs, encompass the entire ecosystem of computing and communications silicon manufacturers, silicon Intellectual Property (IP) vendors and equipment developers. PCI Express* architecture as a new chip-to-chip interconnect and Advanced

Switching based on PCI Express architecture for system fabrics are positioned to offer overwhelming benefits to the communications and embedded industries over other niche technologies.

In the following pages, we look at the industry and market trends that are catalyzing the need for common interconnect topologies, the benefits of common interconnects and why PCI Express architecture and Advanced Switching based on PCI Express architecture offer compelling benefits to the industry.

Fewer Interconnects Simplify Convergence and Modular Communications Platforms

Communications is undergoing two significant industry trends: greater convergence with computing technologies and a changing value chain that is catalyzing development and adoption of modular platforms for deploying complex, converged solutions. Added to these trends, the silicon industry, at an interconnect level, is no longer segmented by compute and communications as it has been. The industry is represented by one body of silicon component manufacturers who look to leverage the collective industry investments in order to reduce costs and embrace market trends. Combined, these impact the choices of interconnect technologies within next-generation communication systems.

TWO INDUSTRY TRENDS BENEFIT COMMUNICATIONS

Computing and communications are converging. Catalyzed by the explosive growth of networking and the enormous adoption of the Internet by hundreds of millions of personal and business users, communications technologies are offered in nearly every computing device deployed. From desktop computers to industrial terminals to laptops and PDAs, wired and wireless networking devices are often integrated as part of the system (Figure 1, next page). In addition, computing is an integral part of nearly all communications functions for media, control plane and management processing in communications equipment.

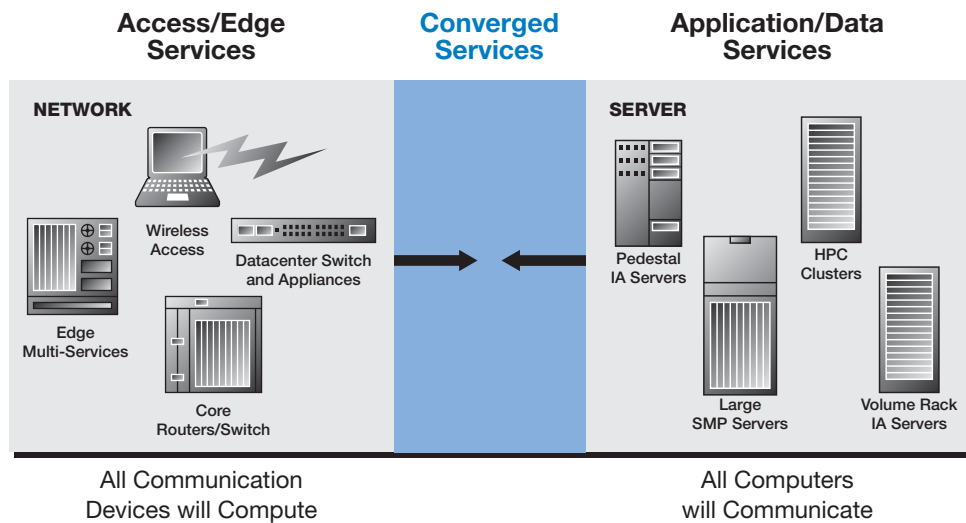


Figure 1. Convergence combines communications and computing technologies

Over their histories, computing has evolved around a single board-level interconnect, the PCI specification, while communications incorporated many board-level and system-level interconnects, some proprietary, others based on standards such as PCI. As the two disciplines converge, an abundance of interconnects creates complexity in interoperability, coding and physical design, all of which drive up cost. The use of fewer, common interconnects will simplify the convergence process and benefit infrastructure equipment developers.

In addition, today's telecommunications industry dilemma of growing network traffic, flat revenue and reduced capital and operational spending, has resulted in developing a modular approach to building converged communications solutions. Modularity allows complex systems to be integrated from system-level and board-level building blocks connected through common interconnects. The modularity model is attractive to many suppliers because of the cost and time-to-market benefits for building complex systems. The Advanced Telecom Compute Architecture (Figure 2) (AdvancedTCA/PICMG 3.x) specifies a modular platform for both computing and communications element to reside in a single chassis.

Industry-standard interconnects that can be reused among multiple platforms are key to both convergence and a modular system design approach. Common chip-to-chip interconnects enable greater design reuse across a board and improve interoperability between the computing and communications functions. A common system fabric enables board-level modularity by standardizing the switching interfaces between various line cards in a modular system. Fewer, common interconnects also reduce complexity in software and hardware and simplify system design. Simplification and reuse drive down costs and development time in modular components.




Figure 2. AdvancedTCA brings together communications and computing in a common, interoperable platform

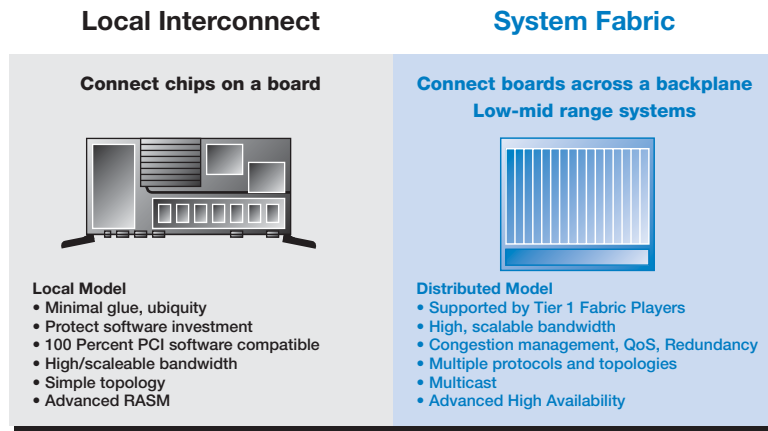


Figure 3. Communications platforms need two key interconnects

FEWER INTERCONNECTS NEEDED

As computing and communications converge and modular communications platforms are deployed, fewer interconnects at the board and system level that enable reuse of existing designs and meet the needs of communications bring enormous value (Figure 3).

In reality, there will always be multiple chip-to-chip interconnects on a board. The Optical Interface Forum and Network Processor Forum have standardized on specialized data-path interconnects that are widely adopted and address very specific needs. In addition to these, however, there are a number of I/O, control and management interconnects that would benefit from consolidation. PCI Express architecture meets the requirements for local chip-to-chip applications with its broad ecosystem support, 100% PCI software compatibility and high, scalable performance features (Table 1).

At the system level, Advanced Switching provides a system fabric interconnect that connects line cards in a modular system while reusing the physical and link layers of the PCI Express architecture.

Advancing Chip-to-Chip Communications: From Buses to Serial Interconnects and PCI Express* Architecture

For over the last decade, the PCI specification has advanced with the evolving computer industry to accommodate faster peripheral device communications in servers, desktops and notebooks. With bandwidth now up to 8 Gbps, PCI continues to support many of the I/O needs of computing applications. The communications and embedded industries also adopted the standard in many applications, with PCI interfaces on communications

Equipment Industry Requirements	PCI Express* Architecture Features
Minimal glue, ubiquity and diversity of interoperable components	Widespread ecosystem adoption, resulting in a wealth of component choices
Maintain Control Processor legacy code	Breadth of PCI Express availability enables easy adoption for control processors.
Lowest cost	Compute industry adoption—huge volumes
100 percent PCI software compatible	No software changes needed
High, scalable bandwidth	2-64 Gbps (each direction) and future PHYs
Advanced Reliability, Availability, Serviceability and Management (RASM)	Hot add/remove, error management, reliable link layer etc.

Table 1. Mapping PCI Express architecture features to equipment industry requirements*

devices, FPGAs, and ASICs. At the platform level, CompactPCI*, based on the PCI architecture, is an industry specification used in many communications solutions.

Around the PCI specification, a vast ecosystem has evolved that enables silicon and board developers to focus on adding value in other parts of the design rather than on the interconnect. Companies—large and small—in the ecosystem offer intellectual property (IP) cores, IP verification code and tools and test equipment focused on PCI implementations. The world’s leading foundries support PCI, as well as ASIC developers and FPGA manufacturers. In addition, manufacturers offer bridges, switches and a host of other components for PCI interoperability. The PCI ecosystem has helped computing and communications developers deliver silicon and boards to market faster and more cost-effectively than proprietary bus implementations. The momentum of this ecosystem is a driving force in interconnect evolution for computing and communications.

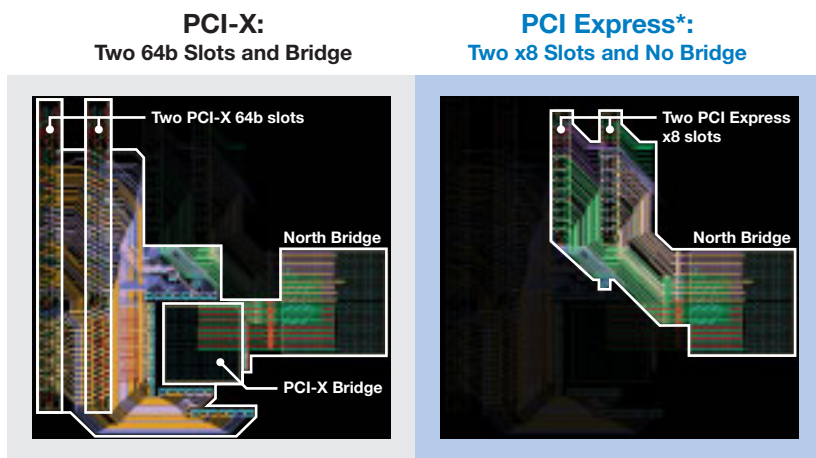
The PCI specification has an extensive following in multiple disciplines. However, as system performance has increased, the inherent limitations of a bus architecture create an overall system performance bottleneck. Increasing bus speed is not practical, because of the effects of increased signal skew. Widening the bus complicates board and silicon layout and connector design with high pin counts and increases required board

space and the number of PCB layers. As communications evolves, there is need for greater bandwidth, lower power requirements and higher port densities. These needs cannot be met by parallel bus architectures; a high-bandwidth, serial interconnect is needed.

Serial interconnect technologies have been under development for several years. Serial interconnects reduce pin count, simplify board layouts (Figure 4) and offer speed, scalability, reliability and flexibility not possible with parallel buses. All versions of these interconnect technologies rely on high-speed serial (HSS) technologies that have advanced as silicon speeds have increased. These new interconnects range from proprietary interconnects for core network routers and switches to standardized serial technologies, applicable to computing, embedded applications and communications, including PCI Express architecture.

The Widely Supported PCI Express* Architecture

PCI Express architecture is the next-generation chip-to-chip interconnect for computing. Intel helped develop the PCI Express architecture along with many other companies that recognized the value of a high-speed serial chip-to-chip interconnect. Today, the PCI SIG owns and manages the PCI Express architecture. Within the compute industry, the PCI Express architecture has



Board area reduced by 53 percent
Board layer count reduction opportunity
Component count decrease

Figure 4. PCI Express* architecture reduces complexity and cost

earned a very large and committed following, developing an ecosystem of PCI Express architecture supporters:

- The major IP core developers, including FPGA vendors, offer a breadth of PCI Express architecture IP cores that can be incorporated into silicon designs.
- The leading IP verification suppliers offer PCI Express architecture verification code and tools to designers.
- Leaders in the tools and test equipment industry support PCI Express architecture in their products.
- Leading foundries support PCI Express architecture across their processes.
- ASIC and FPGA device manufacturers are incorporating PCI Express architecture into their designs.
- Ethernet, storage, I/O processors, switches and bridges and other components are being developed with PCI Express architecture.

A similar commitment is growing within the communications industry, as silicon vendors are able to reuse these foundation products and services from the compute ecosystem. The Intel® Developer Network (DevNet) for PCI Express Architecture consists of over 280 companies, with over 100 members providing products applicable to the communications industry.

The reason for this growing support lies in the architecture itself. PCI Express architecture is 100 percent PCI software compatible and requires no changes to drivers or BIOS. The enormous investment in PCI software over the last decade will not be lost. A foundation for a PCI Express architecture ecosystem that contains significant momentum already exists. More importantly, communications developers can reuse designs from computing in their new products, offering them significant cost savings and faster time-to-market opportunities.

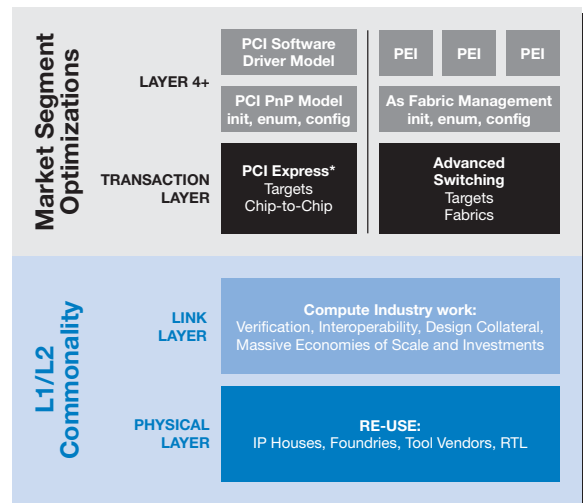
PCI Express architecture simplifies board layout. Only one routing scheme for board layout designers and one set of PHY IP for silicon designers is needed, making it easier for equipment and silicon developers to reuse the same design across multiple products, whether they are compute, storage or communications components.

PCI Express architecture also ensures wide interoperability across chosen devices through the disciplined, thorough validation efforts of the compute industry.

The breadth and depth of the PCI Express architecture ecosystem means silicon manufacturers and equipment developers have choice, availability, competitive costs, interoperability and lower barriers that enable them to innovate. The most important innovations in communications are going to come from developers with new ideas, equipment and services for their customers, not from solving interconnect problems. Due to its scalability and ubiquity, PCI Express architecture enables these innovators to focus on giving the world these creative solutions instead of spending time creating different interconnects to differentiate their products.

Advanced Switching— A System Fabric Interconnect

Advanced Switching based on PCI Express architecture is an emerging technology for data-plane and unified fabrics, and it has increasing support by fabric silicon vendors. Advanced Switching cleanly overlays the core of PCI Express technology (Figure 5), adding transaction layer capabilities to enable flexible protocol encapsulation,



Layered architecture enables modularity, reuse and scalability across Chip-to-Chip, fabrics and market segments

Figure 5. Advanced switching and PCI Express* architecture

peer-to-peer transfers, dynamic reconfigurations, multicast and more. It is specifically intended to complement the communications-centric portions of PCI Express.

Common physical and interconnect layers offer important cost savings to board developers and solution providers through reusability and backplane interconnect. In addition, a strong ecosystem foundation is already in place for device manufacturers and equipment vendors to more easily adopt the new system fabric interconnect.

The AdvancedTCA specification (PICMG 3.4) accommodates several standard system fabric interconnects, including Ethernet, Infiniband* Architecture, PCI Express architecture and Advanced Switching based on PCI Express architecture.

Conclusion

PCI Express architecture is the only chip-to-chip interconnect that embraces the communication and embedded equipment industry requirements for an industry standard, common interconnect that can be used to build standard building blocks for the communication/ embedded market as well as the compute market. PCI Express architecture allows a smooth transition from PCI to a faster, more scalable interconnect technology widely accepted throughout the compute industry. Adoption of PCI Express will allow the communications and embedded industry to take advantage of the lowest cost, most widely available interconnect in the world.

BENEFITS OF PCI EXPRESS* ARCHITECTURE

- Fast, efficient intercommunication among devices through a scalable and flexible interconnect architecture.
- High reliability using a robust protocol with precise error detection and recovery within the link layer.
- High system availability and hot swappability.
- Low pin count, leading to improved routeability, simple board layout, reduced board space and fewer PCB layers than parallel technologies.

- Because it is 100 percent PCI software compatible with no changes to drivers or BIOS, PCI Express architecture allows seamless migration of software with existing products based on PCI. PCI Express architecture utilizes the enormous installed base of PCI code throughout the compute, embedded and communications industries.

BENEFITS OF ADVANCED SWITCHING BASED ON PCI EXPRESS* ARCHITECTURE

- Industry-standard system fabric interconnect for data plane and unified fabrics.
- Supports direct peer-to-peer communications between any two devices.
- Encapsulation to route any protocol across the fabric. Also provides standardized encapsulations for basic load/store, message queuing, RDMA, PCI Express tunneling and others.
- Path based routing and advanced congestion management for traffic engineering and QoS with low latency.
- High availability support, including multiple mastering of the fabric and redundant paths.
- Supports a variety of switch topologies.

For More Information

Intel, as a leading silicon provider for both the communications and computing industries, is committed to aligning all of its compute and communications silicon, platforms, solutions and ecosystem-enabling with PCI Express. Intel supports PCI Express architecture and Advanced Switching based on PCI Express architecture with many programs for device manufacturers and equipment vendors that help them realize the benefits of these two technologies.

FOR MORE INFORMATION, VISIT THE FOLLOWING SITES:

- Intel® Developer Network (DevNet) for PCI Express Architecture: www.express-lane.org
- PCI SIG: www.pcisig.com
- PICMG: www.picmg.org

Intel around the world

United States and Canada

Intel Corporation
Robert Noyce Building
2200 Mission College Boulevard
P.O. Box 58119
Santa Clara, CA 95052-8119
USA
Phone
General information: (408) 765-8080
Customer support: (800) 628-8686

Europe

Intel Corporation (UK) Ltd.
Pipers Way
Swindon
Wiltshire SN3 1RJ
UK
Phone
England: (44) 1793 403 000
France: (33) 1 4694 71701
Germany: (49) 89 99143 0
Ireland: (353) 1 606 7000
Israel: (972) 2 589 7111
Italy: (39) 2 575 441
Netherlands: (31) 20 659 1800

Asia-Pacific

Intel Semiconductor Ltd.
32/F Two Pacific Place
88 Queensway, Central
Hong Kong, SAR
Phone: (852) 2844 4555

Japan

Intel Kabushiki Kaisha
P.O. Box 300-8603 Tsukuba-gakuen
5-6 Tokodai, Tsukuba-shi
Ibaraki-ken 300-2635
Japan
Phone: (81) 298 47 8511

South America

Intel Semicondutores do Brasil
Avenida Dr. Chucri Zaidan, 940, 10t
São Paulo
Brazil
Phone: (55) 11 3365 5500

For more information

To learn more about Intel Corporation
visit our site on the World Wide Web
at www.intel.com



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, life saving, life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

*Other names and brands may be claimed as the property of others.

Copyright © 2003 Intel Corporation. All rights reserved.
Order Number: 253053-001 Printed in the USA/0603/MS/KF