

PEX 8525

Features

- **PEX 8525 General Features**
 - 24-lane PCI Express switch
 - Integrated SerDes
 - Up to five configurable ports (x1, x2, x4, x8, x16)
 - 31mm x31mm, 644-ball PBGA package
 - Typical Power: 2.6 Watts
- **PEX 8525 Key Features**
 - **Standard Compliant**
 - PCI Express Base Specification, r1.0a
 - **High Performance**
 - Non-blocking switch fabric
 - Full line rate on all ports
 - Packet Cut-Thru with 115ns max packet latency (x8 to x8)
 - **Flexible Configuration**
 - Five highly flexible & configurable ports (x1, x2, x4, x8, or x16)
 - Configurable with strapping pins, EEPROM, I²C, or Host software
 - Lane and polarity reversal
 - **PCI Express Power Management**
 - Link power management states: L0, L0s, L1, L2/L3 Ready, and L3
 - Device states: D0 and D3hot
 - **Quality of Service (QoS)**
 - One Virtual Channel per port
 - Eight Traffic Classes per port
 - Weighted Round-Robin Ingress Port Arbitration
 - **Reliability, Availability, Serviceability**
 - 3 Standard Hot-Plug Controllers
 - Upstream port as hot-plug client
 - Transaction Layer end-to-end CRC
 - Poison bit
 - INTA# interrupt signal
 - Fatal Error (FATAL_ERR#) signal (legacy SERR equivalent)
 - PCIe baseline error reporting
 - Advanced Error Reporting
 - Port Status bits and GPO available
 - Per port error diagnostics
 - Bad DLLPs
 - Bad TLPs
 - CRC errors
 - JTAG boundary scan



High-Performance 24-lane, 5-port PCIe Switch

Multi-purpose, High Performance *ExpressLane*[™] Switch

The *ExpressLane* PEX 8525 device offers PCI Express switching capability enabling users to add scalable high bandwidth, non-blocking interconnection to a wide variety of applications including **servers, storage systems, communications platforms, blade servers, and embedded-control products**. The PEX 8525 is well suited for **fan-out, aggregation, dual-graphics, peer-to-peer, and intelligent I/O module** applications.

Highly Flexible Port Configurations

The PEX 8525 offers highly configurable ports. There are a maximum of 5 ports that can be configured to any legal width from x1 to x16, in any combination to support your specific bandwidth needs. The ports can be configured for **symmetric** (each port having the same lane width and traffic load) or **asymmetric** (ports having different lane widths) traffic. In the event of asymmetric traffic, the PEX 8525 features a **flexible central packet memory** that allocates a memory buffer for each port as required by the application or endpoint. This buffer allocation along with the device's **flexible packet flow control** minimizes bottlenecks when the upstream and aggregated downstream bandwidths do not match (are asymmetric). Any of the ports can be designated as the upstream port, which can be changed dynamically.

High Performance

The PEX 8525 architecture supports packet **cut-thru with a max latency of 115ns (x8 to x8)**. This, combined with large packet memory (**1024 byte maximum payload size**) and non-blocking internal switch architecture, provide full line rate on all ports for performance-hungry applications such as **storage servers or storage switch fabrics**.

End-to-end Packet Integrity

The PEX 8525 provides **end-to-end CRC** protection (ECRC) and **Poison bit** support to enable designs that require **end-to-end data integrity**. These features are optional in the PCI Express specification, but PLX provides them across its entire *ExpressLane* switch product line.

Configuration Flexibility

The PEX 8525 provides several ways to configure its operations. The device can be configured through strapping pins, **I²C interface**, CPU configuration cycles, or an optional serial EEPROM. This allows for easy debug during the development phase, performance monitoring during the operation phase, and driver or software upgrade.

Interoperability

The PEX 8525 is designed to be fully compliant with the PCI-SIG revision 1.1 specification. Additionally, it supports **auto-negotiation, lane reversal, and polarity reversal**. The PEX 8525 also undergoes thorough Interoperability testing in PLX's **Interoperability Lab**.

PLX Confidential

Low Power with Granular SerDes Control

The PEX 8525 provides low power capability that is fully compliant with the PCI Express power management specification. In addition, the SerDes physical links can be turned off when unused for even lower power.

Flexible Port Width Configuration

The width of each port can be individually configured through auto-negotiation, hardware strapping, host software configuration, I²C interface, or through an optional EEPROM.

The PEX 8525 supports a large number of port configurations (see Figure 1). For example, if you are using the PEX 8525 in a fan-out application, you may configure the upstream port as x8 and the downstream ports as four x4 ports; two x8 ports; or other combinations as long as you don't run out of lanes (24) or ports (5). In a peer-to-peer application, you can configure all five ports as x4. In a port aggregation application you can configure four x2 or x4 ports for aggregation into one x8 port.

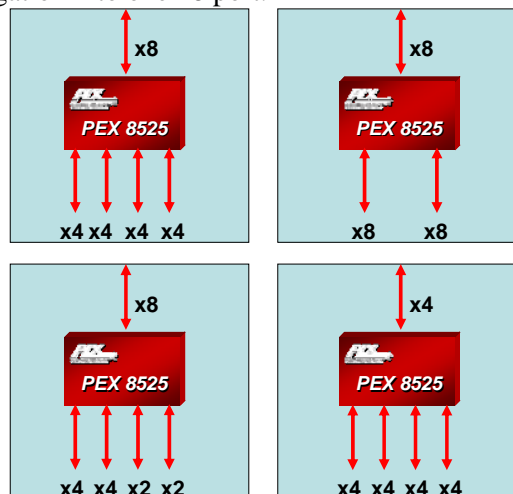


Figure 1. Common Port Configurations

Low Packet Latency

The PEX 8525 supports packet **cut-thru** with a max packet latency of 115ns between symmetric x8 ingress and egress ports. The low latency enables applications to achieve high throughput and performance. In addition to low latency, the device supports a packet payload size of up to 1024 bytes, enabling the user to achieve even higher throughput.

Hot Plug for High Availability

Hot plug capability allows users to replace hardware modules and perform maintenance without powering down the system. The PEX 8525 hot plug capability and

Advanced Error Reporting features makes it suitable for **High Availability (HA) applications**. Three downstream ports include a Standard Hot Plug Controller. If the PEX 8525 is used in an application where one or more of its downstream ports connect to PCI Express slots, each port's Hot Plug Controller can be used to manage the hot-plug event of its associated slot. Furthermore, its upstream port is a **hot-plug client**, allowing it to be **used on hot-pluggable adapter cards, backplanes, and fabric modules**.

Fully Compliant Power Management

For applications that require power management, the PEX 8525 device support both link (L0, L0s, L1, L2/L3 Ready, and L3) and device (D0 and D3hot) power management states, in compliance with the PCI Express power management specification.

SerDes Power and Signal Management

The PEX 8525 supports **software control** of the **SerDes outputs** to allow optimization of power and signal strength in a system. The PLX SerDes implementation supports four levels of power – off, low, typical, and high. The SerDes block also supports **loop-back modes** and **advanced reporting of error conditions**, which enables efficient debug and management of the entire system.

Applications

Suitable for **host-centric** as well as **peer-to-peer traffic patterns**, the PEX 8525 can be configured for a wide variety of form factors and applications.

Host Centric Fan-out

The PEX 8525, with its symmetric or asymmetric lane configuration capability, allows user-specific tuning to a variety of host-centric applications.

Figure 2 shows a typical **server-based** design where the root complex provides a PCI Express link that needs to be expanded to a larger number of smaller ports for a variety of I/O functions. In this example, the PEX 8525 would typically have an 8-lane upstream port, and as many as




Figure 2. Fan-in/out Usage

four x4 downstream ports. The downstream ports can be of differing widths if required. Almost all (non x86 based) high-end microprocessor manufacturers are offering PCI Express interfaces. The PEX 8525 can be directly connected to a processor to fan-out its PCIe port to a larger number of ports for enhanced connectivity as illustrated in Figure 3.




Figure 3. Fan out for PowerPC/MIPS CPUs

Embedded Systems

The PEX 8525 can also be utilized in embedded applications. Figure 4 shows several independent modules connecting through the PEX 8525. The port widths for each module can be configured as required. The peer-to-peer communication feature of the PEX 8525 allows these modules to communicate with each other without any centralized control.

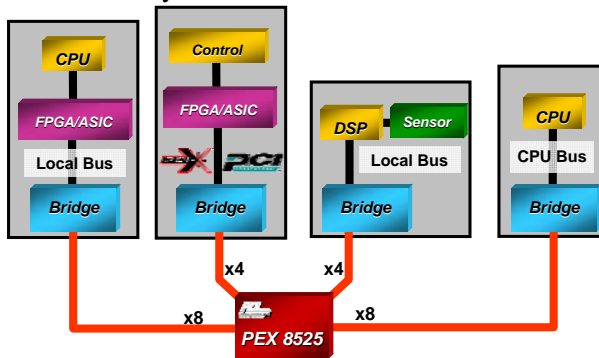


Figure 4. Embedded Systems

Peer to Peer Communication

Figure 5 represents a backplane where the PEX 8525 provides peer-to-peer data exchange for up to four line cards where the CPU/Host plays the management role.




Figure 5. Peer-to-Peer Usage

Graphics Fan-out Switch

As PCIe based graphics cards become more mainstream, it will be necessary to take a x8 port on the root complex device and fan it out to two x8 ports for dual graphics applications. Root Complex (Northbridge) devices are available with multiple PCIe ports. These ports can be further expanded to connect to a larger number of I/Os or to support dual-graphics using the PEX 8525 as shown in Figure 6. In this dual graphics application example, a x8 port on the root complex device is fanning out to two PCIe graphics cards via x8 links. The figure also shows how some of the ports can be bridged to provide PCI or PCI-X slots through the use of the ExpressLane PEX 8114 and PEX 8111 PCIe bridging devices.

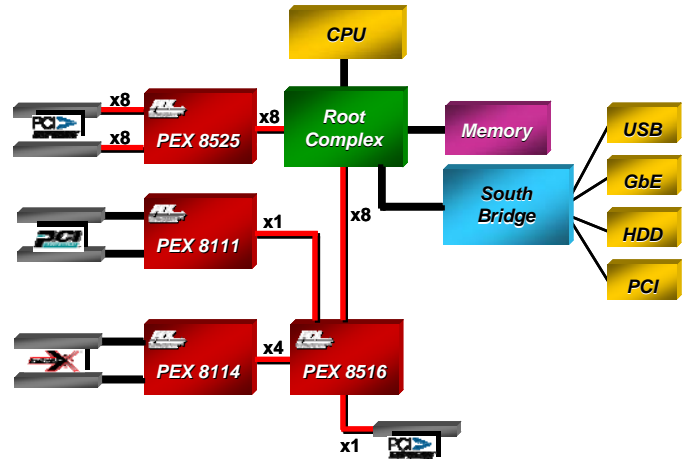


Figure 6. Dual Graphics Fan-out

PCI Express Port Expansion

The PEX 8525 enables designers to take, for example, two x8 PCIe ports and expand them into eight downstream ports. Some of these PCIe ports can be bridged to PCI or PCI-X using bridging products from PLX. Figure 7 illustrates one of the many configurations the PEX 8525 can support.

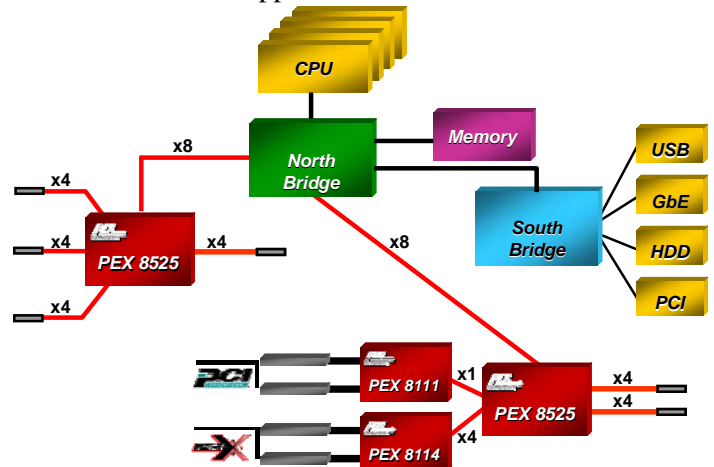


Figure 7. PCIe Port Expansion

Software Usage Model

From a system model viewpoint, each PCI Express port is a virtual PCI to PCI bridge device and has its own set of PCI

Express configuration registers. It is through the upstream port that the BIOS or host can configure the other ports using standard PCI enumeration. The virtual PCI to PCI bridges within the PEX 8525 are compliant to the PCI and PCI Express system models. The Configuration Space Registers (CSRs) in a virtual primary/secondary PCI to PCI bridge are accessible by

type 0 configuration cycles through the virtual primary bus interface (matching bus number, device number, and function number).

Interrupt Sources/Events

The PEX 8525 switch supports the INTx interrupt message type (compatible with PCI 2.3 Interrupt signals) or Message Signaled Interrupts (MSI) when enabled. Interrupts/messages are generated by PEX 8525 for hot plug events, doorbell interrupts, baseline error reporting, and advanced error reporting.

Development Tools

PLX offers hardware and software tools to enable rapid customer design activity. These tools consist of a hardware module (PEX 8525 RDK), hardware documentation, and a Software Development Kit (SDK).

ExpressLane PEX 8525 RDK

The RDK hardware module includes the PEX 8525 with one x8 (card-edge slot) port, two x4 ports and one x8 port (see Figure 8). The RDK is available with a x8 card edge connector/adaptor. The adapters for x4 and x1 edge connectors are available to plug the RDK in smaller slots. The PEX 8525 RDK board can be installed on a motherboard or used as a riser card. The PEX 8525 RDK can be used to test and validate customer software. Additionally, it can be used as an evaluation vehicle for the PEX 8525 features and benefits. The PEX 8525 RDK provides everything that a user needs to get their hardware and software development started.

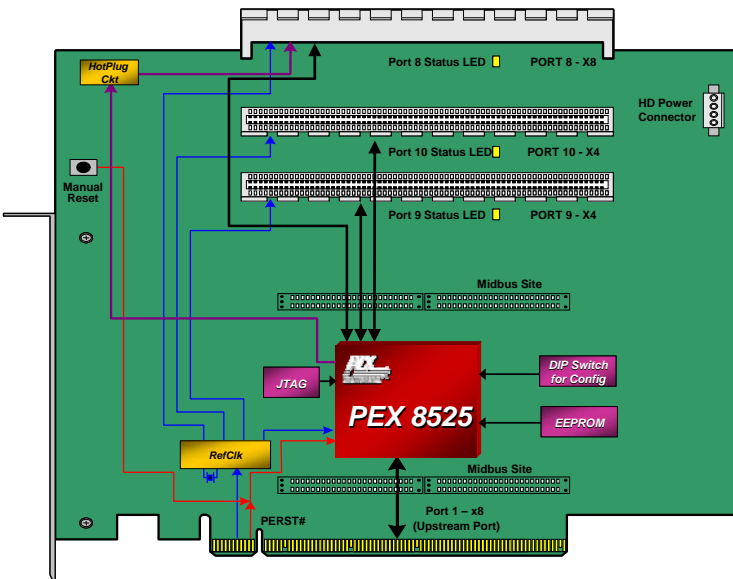



Figure 8. PEX 8525 RDK

SDK

The SDK tool set includes:

- Linux & Windows drivers
- C/C++ Source code, Objects, libraries
- User's Guides & Application examples



PLX Technology, Inc.
870 Maude Ave.
Sunnyvale, CA 94085 USA
Tel: 1-800-759-3735
Tel: 1-408-774-9060
Fax: 1-408-774-2169

Email: info@plxtech.com
Web Site: www.plxtech.com

Product Ordering Information

Part Number	Description
PEX8525-AA25BI	24-Lane, 5-Port PCIe Switch, I-Temp, 31x31mm ²
PEX8525-AA25BI G	24-Lane, 5-Port PCIe Switch, Pb-Free, I-Temp, 31x31mm ²
PEX 8525-AA RDK	PEX 8525 Rapid Development Kit w/ x8 Edge Connector

Please visit the PLX Web site at <http://www.plxtech.com> or contact PLX sales at 408-774-9060 for sampling.

© 2008 PLX Technology, Inc. All rights reserved. PLX and the PLX logo are registered trademarks of PLX Technology, Inc. ExpressLane is a trademark of PLX Technology, Inc., which may be registered in some jurisdiction. All other product names that appear in this material are for identification purposes only and are acknowledged to be trademarks or registered trademarks of their respective companies. Information supplied by PLX is believed to be accurate and reliable, but PLX Technology, Inc. assumes no responsibility for any errors that may appear in this material. PLX Technology, Inc. reserves the right, without notice, to make changes in product design or specification.