

AMD HyperTransport™ Technology-Based System Architecture

ADVANCED MICRO DEVICES, INC.
One AMD Place
Sunnyvale, CA 94088

Table of Contents

Introduction	3
AMD-8000™ Series of Chipset Components Product Overview	3
AMD-8111™ I/O Hub	3
AMD-8131™ I/O Bus Tunnel	4
AMD-8151™ Graphics Tunnel.....	4
Key Features of the AMD-8000™ Series of Chipset Components.....	5
HyperTransport™ Technology.....	5
AGP-8X.....	6
PCI-X	6
USB2.0	7
HyperTransport™ Technology I/O Architecture Overview	7
System Performance Enhancement	7
HyperTransport Technology Link.....	10
HyperTransport Technology Devices	11
HyperTransport Technology-Based System Architecture.....	13
Summary	14
References.....	15

Introduction

AMD introduces the AMD-8000™ series of chipset components designed to support AMD’s eighth-generation AMD Athlon™ and AMD Opteron™ microprocessors. Incorporating HyperTransport technology as the chip-to-chip interconnect, these core logic elements integrate the latest I/O technologies to form the system building blocks for a wide range of high-performance platforms harnessing the power of AMD’s eighth-generation AMD Athlon processor for mobile and desktop applications and AMD Opteron processor for workstation and server applications.

AMD-8000™ Series of Chipset Components

The AMD-8000 series of chipset components consists of the following HyperTransport technology-compliant tunnels and I/O hub.

- AMD-8111™ I/O Hub
- AMD-8131™ I/O Bus Tunnel
- AMD-8151™ Graphics Tunnel

AMD-8111™ I/O Hub

The feature-rich AMD-8111 I/O Hub integrates storage, connectivity, audio, I/O expansion, and system management functions into a single device.

- HyperTransport link
 - 8-bit HyperTransport interface
 - 800MB/s aggregate bandwidth
- Six USB2.0 ports
- Up to eight PCI masters
- Integrated 10/100 MAC
- AC-97 soft modem and 6-channel audio
- LPC bus
- Dual EIDE channels
 - ATA 33/66/100/133

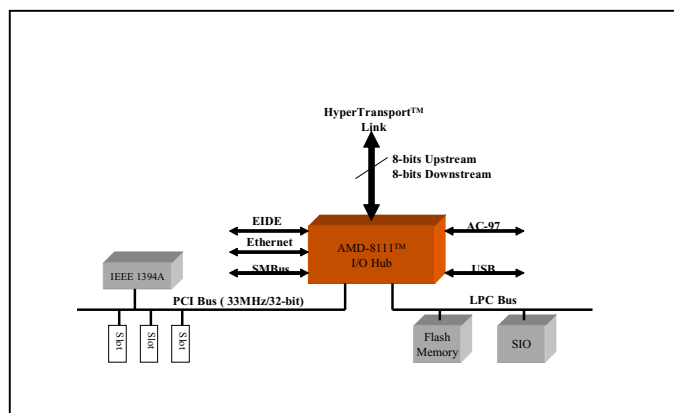


Figure 1: AMD-8111™ HyperTransport I/O Hub

AMD-8131™ I/O Bus Tunnel

The AMD-8131 I/O bus tunnel is a high-speed device that provides two independent, high-performance PCI-X bus bridges, with two high-speed HyperTransport technology links.

- HyperTransport link A
 - 16-bit HyperTransport interface
 - 6.4GB/s aggregate bandwidth
- HyperTransport link B
 - 8-bit HyperTransport interface
 - 3.2GB/s aggregate bandwidth
- Dual PCI-X channels: A and B
 - 133/100/66/33MHz operation
 - 66/33MHz PCI2.2 modes
 - Up to five PCI masters per bridge
 - SHPC-compliant hot plug support

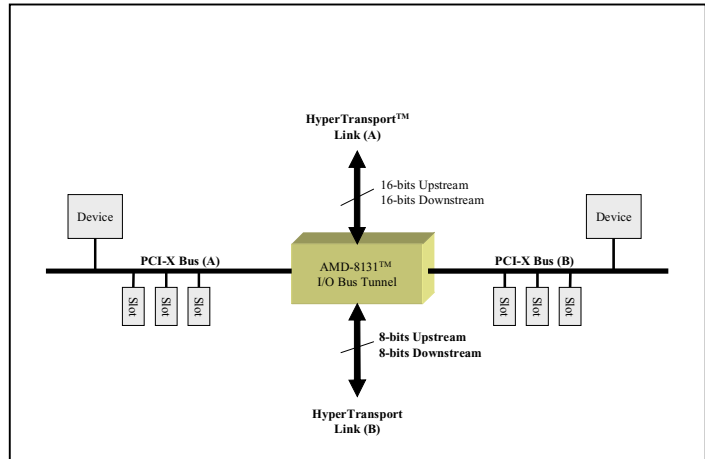


Figure 2: AMD-8131™ I/O Bus Tunnel

AMD-8151™ Graphics Tunnel

The AMD-8151 graphics tunnel is an AGP3.0-compliant graphics controller offering stunning graphics performance modes up to AGP-8X. The tunnel function provides connection capability to other downstream HyperTransport technology devices, allowing greater system flexibility.

- HyperTransport link A
 - 16-bit HyperTransport interface
 - 6.4GB/s aggregate bandwidth
- HyperTransport link B
 - 8-bit HyperTransport interface
 - 1.6GB/s aggregate bandwidth
- AGP3.0 specification compliant
 - 8X/4X transfer mode
- AGP2.0 specification compliant
 - 4X/2X/1X transfer mode

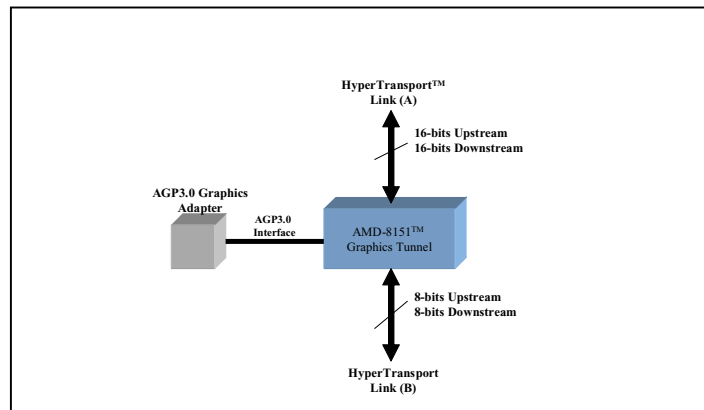


Figure 3: AMD-8151™ Graphics Tunnel

Key Features of the AMD-8000™ Series of Chipset Components

The AMD-8000™ series of chipset components has integrated the latest I/O technologies to enable the next generation of high-performance platforms. Key features include HyperTransport technology, AGP-8X, PCI-X, and USB2.0.

HyperTransport Technology

HyperTransport technology is a high-speed, high-performance, point-to-point link for integrated circuits, and is designed to meet the bandwidth needs of tomorrow's computing and communications platforms. At a peak throughput of up to 12.8GB/s per link, HyperTransport technology provides an I/O solution for the most demanding system applications.

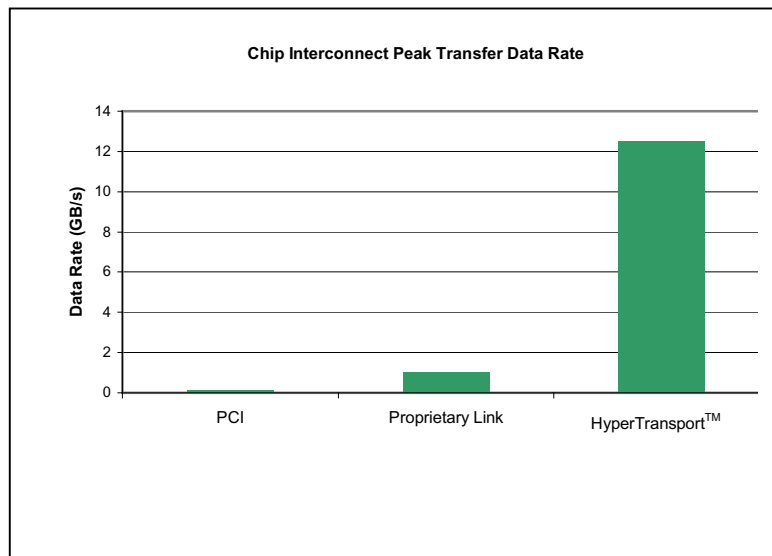


Figure 4: System Interconnect Bandwidth

AGP-8X

AGP-8X is the latest performance enhancement to the AGP graphics interface. It is compatible to the existing standard while increasing the peak data transfer rate to 2.1GB/s to provide the performance headroom for the next generation of graphic processors.




Figure 5: Graphics Interconnect Bandwidth

PCI-X

PCI-X is a performance extension of the existing PCI bus. PCI-X doubles the data bus to 64-bit and quadruples the clock speed to 133MHz to provide up to 1GB/s of data transfer rate. PCI-X provides the performance headroom to meet the demanding high-speed requirements of applications such as SCSI adapter, Gigabit Ethernet, Fibre Channel, and many others.

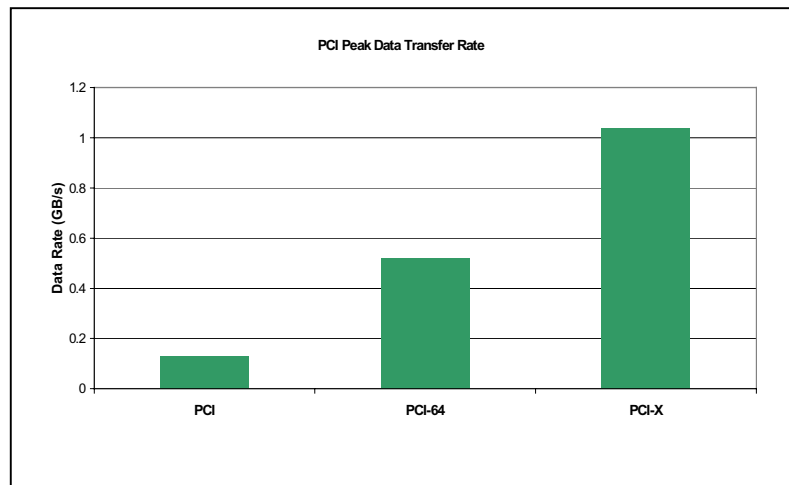


Figure 6: PCI Bus Bandwidth

USB2.0

USB2.0 is the latest performance evolution of the USB standard. At 480MB/s transfer rate, USB2.0 will enable a new generation of high-performance peripherals.




Figure 7: USB Bandwidth

HyperTransport Technology I/O Architecture Overview

HyperTransport technology is a high-speed, high-performance, point-to-point link for integrated circuits, and is designed to meet the bandwidth needs of tomorrow's computing and communications platforms. At a peak throughput of up to 12.8GB/s per link, HyperTransport technology provides an I/O solution for the most demanding system applications.

System Performance Enhancement

HyperTransport technology is designed to increase overall performance by helping to remove I/O bottlenecks, which improves bandwidth and reduces latency. Figure 8 shows some of the performance bottleneck areas that HyperTransport technology alleviates.

1. The processor front-side bus
2. Memory interface
3. Chip-to-chip interconnect
4. I/O expansion capability to high-speed industry buses

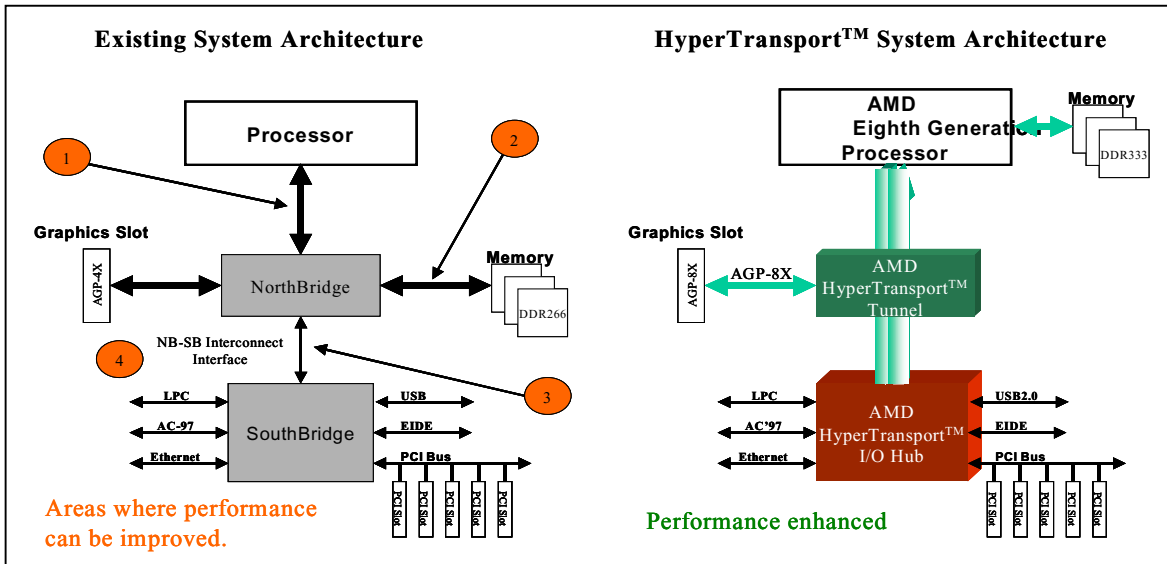


Figure 8: System Architectures

Performance Enhancement 1: Processor Front-Side Bus

For optimal performance, the front-side bus bandwidth must scale with increasing processor speeds. Current front-side bus bandwidth on AMD’s seventh generation platforms is on the order of 2.1GB/s. Replacing what has traditionally been the system “front-side bus” with a HyperTransport technology-based I/O connection dramatically extends processor to system communication bandwidth from 2.1GB/s up to 6.4GB/s (and potentially 12.8GB/s with future devices).

Performance Enhancement 2: Memory Interface

When cache misses occur, the processor must fetch information from main memory. In the Northbridge/Southbridge architecture shown in Figure 8, memory transactions must traverse through the Northbridge element, creating additional latencies that reduce performance potential. To help resolve this performance bottleneck, AMD incorporates the memory controller into its eighth-generation processor. The direct interface to the memory can significantly reduce the memory latency seen by the processor. This latency will continue to drop as the processor frequency scales. Additionally, hardware and software memory prefetching mechanisms can further reduce the effective memory latency seen by the processor.

This reduction in memory latency coupled with the additional increase in memory bandwidth available directly to the processor resulting from this platform architecture design optimization cannot be overstated as it tremendously benefits system performance across all application segments.

Performance Enhancement 3: Chip-to-Chip Interconnect

Current interface schemes offer throughput performances on the order of 266MB/s to 1GB/s. These rates may be sufficient for desktop platforms; however, a more robust interface is required for workstation, server, and other future platforms. The simultaneous integration of high-speed technologies such as (AGP-8X, Gigabit Ethernet, PCI-X, and Infiniband architecture, etc.), onto the high-end platforms will quickly dwarf the bandwidth capabilities of the existing interfaces. HyperTransport technology provides a high-speed, chip-to-chip interconnect that virtually eliminates the I/O performance bottleneck with ample performance headroom for future growth.

Performance Enhancement 4: I/O Expansion Capability to High-Speed Industry Buses

The Northbridge/Southbridge architecture shown in Figure 8 is not intended to support more than two “core-logic” elements. Adding additional high-speed functionality (such as Gigabit Ethernet, PCI-X, Infiniband architecture, and combinations, etc.), would have to occur in one of three ways:

1. The functionality would have to be attached to an existing bus interface such as the PCI bus. However, an existing bus may not have sufficient bandwidth to support high-speed technologies, especially in instances where multiple buses or combinations of buses must be supported simultaneously.
2. The functionality would have to be directly attached to the higher-speed proprietary chip-to-chip interconnect bus via a bridging device. However, the proprietary nature of this solution may limit the number of components available from vendors, thus impacting cost and availability.
3. The functionality would have to be integrated into one of the core logic components. This solution is the least flexible, as a wide range of components would have to be created for each desired combination of feature-set buses.

HyperTransport technology, an industry standard, provides system designers a high-speed, daisy-chained interconnect between system components. Specific components can be connected in a building-block fashion to achieve a platform with specific feature-set and performance objectives.

HyperTransport Technology Link

HyperTransport technology devices (processor and core-logic components) are connected in a peer-to-peer fashion using high-speed HyperTransport technology I/O links. Each **link** is composed of a **Transmit (Tx)** sub link and a **Receive (Rx)** sub link, which operate independently and concurrently. Each HyperTransport sub link can potentially deliver up to 6.4GB/s of throughput¹, creating an aggregate bandwidth of up to 12.8GB/s per link. Shown in Figure 9, compositions of the HyperTransport link: Transmit, Receive.

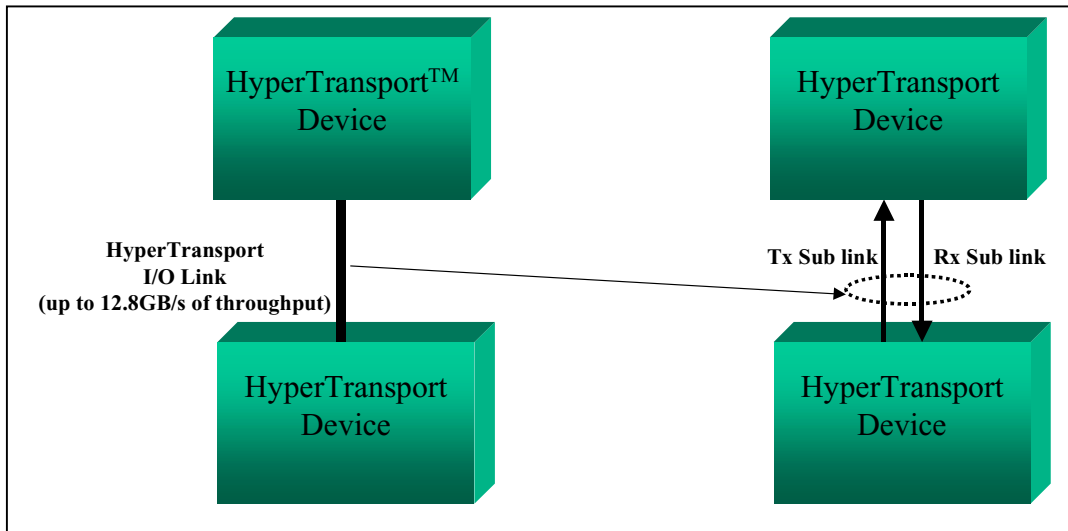


Figure 9: HyperTransport Links Consist of Two Independent Sub Links – Tx Sub Link and Rx Sub Link

¹ This throughput is a theoretical maximum determined by the selected clock rate (200MHz to 800MHz) and data path width (2-bit, 4-bit, 8-bit, 16-bit, or 32-bit). In this case, the assumption is an 800MHz clock frequency with 32-bit data paths on all sub links.

HyperTransport technology links and devices are **scalable** in both frequency and data-path width. With a default clock frequency of 200MHz, component designers can implement clock frequencies from 200MHz to 800MHz. **Double Data Rate (DDR)** memory technology techniques are implemented such that data is transferred on every clock edge (i.e., a 200MHz clock produces 400 mega transfers per second), enhancing system performance.

To provide further flexibility, HyperTransport technology allows the system designer to control cost and performance by providing scalable data path widths for each sub link. Designers can choose sub link data path widths of 2, 4, 8, 16, or 32 bits. Devices of different widths can be connected and will operate at the smaller bus width (auto-negotiation techniques are used to detect and configure the system data paths). Since data path widths impact the component's pin-count, designers have the flexibility to target specific package, price, and performance objectives. This flexibility also allows a single device to be used in multiple applications, each with a different performance objective. A device with 16-bit HyperTransport link can be connected to devices with slower 2, 4, and 8-bit link. The same 16-bit device can also be designed connected to other devices with 16-bit or to a higher 32-bit link.

HyperTransport technology embeds compliant **power-management** features for mobile, small form-factor, and desktop applications where power conservation is critical. Signal integrity, error detection, and data management/pacing features are also embedded to ensure **reliability** and **stability** in all applications.

HyperTransport Technology Devices

As shown in Figure 10, **tunnel** devices are introduced to provide connectivity solutions where HyperTransport technology must be extended to other downstream core-logic elements. In effect, a chain of HyperTransport technology devices can be created to provide greater system flexibility and functionality.

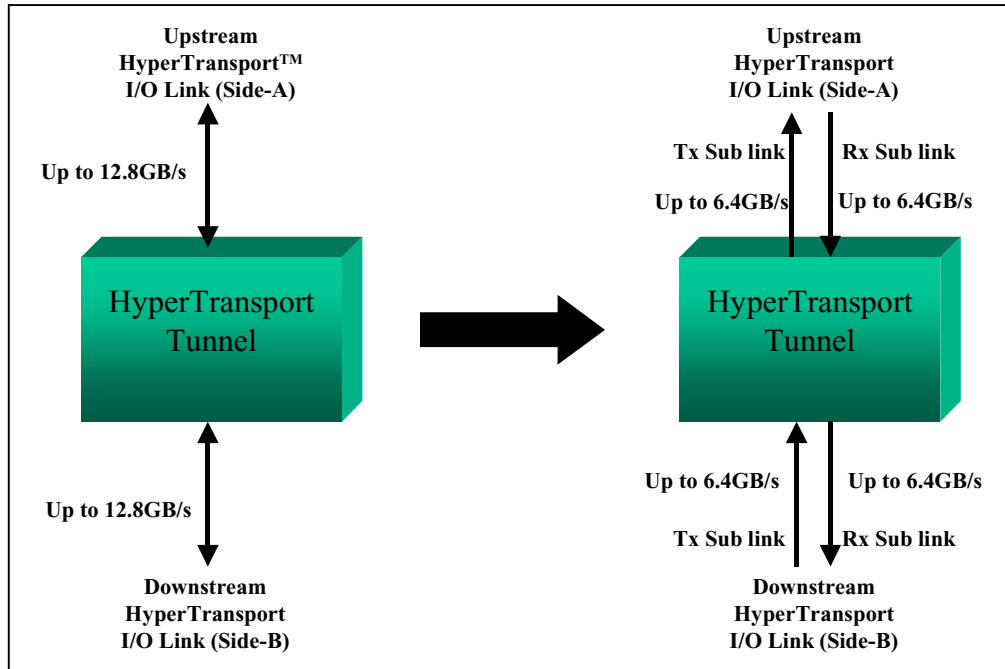


Figure 10: HyperTransport Tunnels Allow Transactions and Data to Flow Through the Component

This tunnel capability simply allows HyperTransport technology transactions to be forwarded upstream or downstream to the neighboring peer device. Note that the maximum aggregate throughput² of a HyperTransport technology tunnel device is 12.8GB/s and not the sum (12.8GB/s + 12.8GB/s) of its upstream and downstream links. The last device to terminate the chain is typically an **I/O hub** device, which does not implement tunnel capability.

HyperTransport technology tunnels and I/O hubs can incorporate bridging capability to other bus technologies such as AGP graphics, PCI bus, PCI-X, Infiniband architecture, EIDE controllers, etc. These devices can also integrate multiple functions onto a single component, such as that seen in traditional Northbridge and Southbridge devices, similar to the AMD-761™ system controller and the AMD-768™ peripheral bus controller, respectively.

² This maximum throughput assumes a device operating at 800MHz clock frequency and implementing a 32-bit data path on all sub links.

HyperTransport Technology-Based System Architecture

HyperTransport technology-based system architecture represents a significant enhancement to traditional system architectures by providing daisy-chained interconnects between system components. Specific components can be connected in a building-block fashion to achieve a platform with specific feature-set and performance objectives. The concepts of Northbridges, Southbridges, and PCI-bus chip-to-chip Interconnects are now replaced with tunnels, I/O hubs, HyperTransport technology links, and HyperTransport technology I/O chains (See Figure 11).

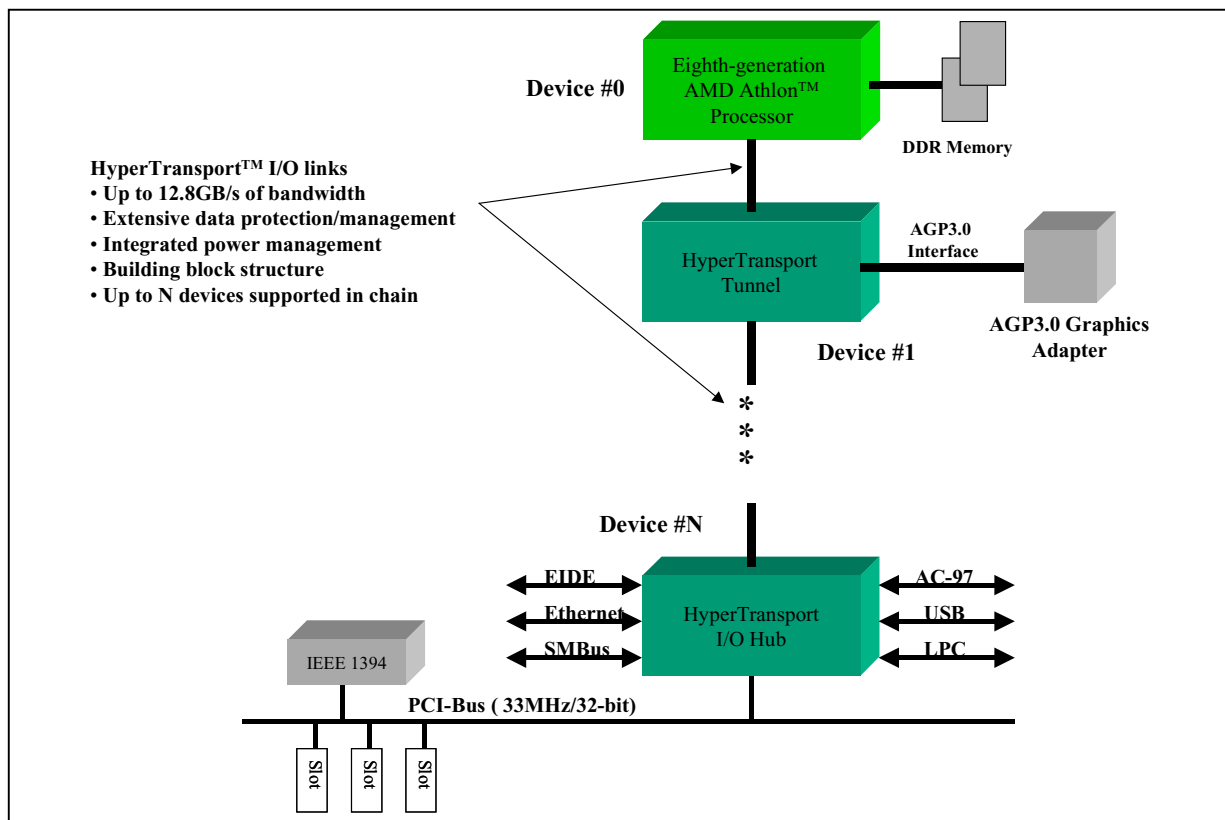


Figure 11: Connected HyperTransport Technology Elements

Systems based upon HyperTransport technology are no longer limited to a two-chip core-logic solution. Multiple components can be chained together in a building-block fashion to produce a multitude of platform configurations. HyperTransport technology has ample bandwidth to support the various high-speed bus technologies, even in combination implementations. The result is an extremely high-performance

system architecture that is flexible, adaptive, and reliable. From mobile and desktop application to workstation and server applications, all can leverage the benefits of this architecture without compromise.

Summary

The AMD-8000 series of chipset components represent the latest innovation in platform solutions by enabling system designers to create highly scalable and high-performing platforms without compromise.

Figure 12 shows a high-performance platform based on the AMD-8151 Graphics Tunnel and the AMD-8111 I/O Hub. This configuration is similar to the traditional Northbridge and Southbridge solutions, however, sideband signals are virtually eliminated, and system performance is increased because of the high-speed HyperTransport technology connections.

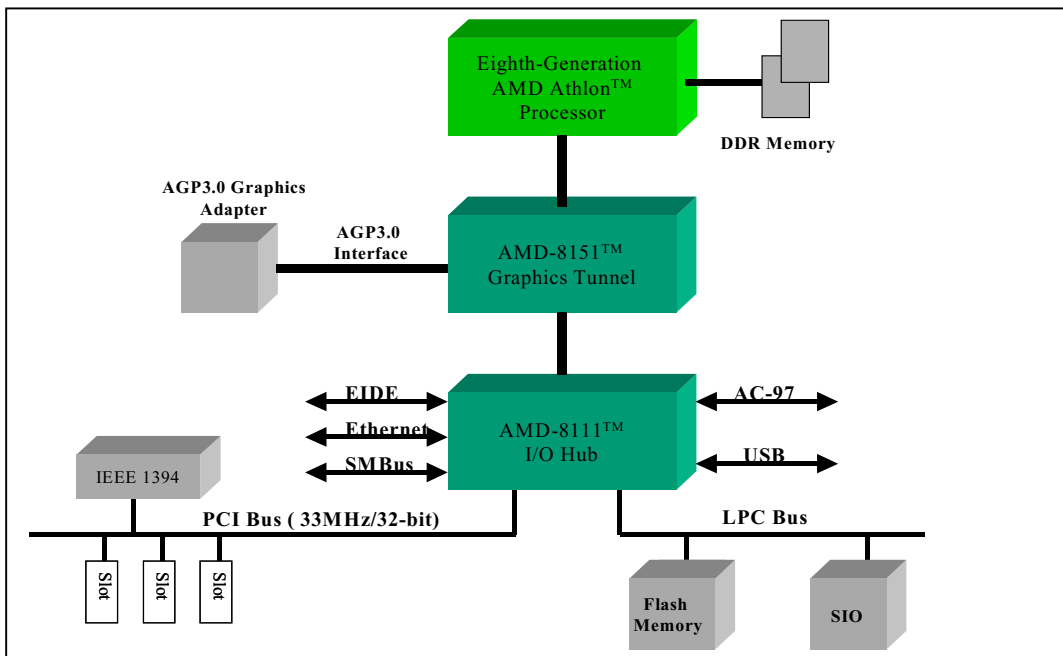


Figure 12: High Performance AMD HyperTransport Technology System Architecture

References

AMD-8111™ I/O Hub Product Brief (PID# 26036A)

AMD-8131™ PCI-X Tunnel Product Brief (PID# 26037A)

AMD-8151™ AGP3.0 Graphics Tunnel Product Brief (PID# 26035A)

For information on HyperTransport I/O Architecture, please refer to the HyperTransport I/O Link Specification, located on the www.HyperTransport.org website.

For Information on AMD's eighth-generation processor architecture, please refer to the "AMD Eighth-Generation Processor Architecture" Whitepaper located on the www.amd.com website.

AMD Overview

AMD is a global supplier of integrated circuits for the personal and networked computer and communications markets with manufacturing facilities in the United States, Europe, and Asia. AMD produces microprocessors, flash memory devices, and support circuitry for communications and networking applications. Founded in 1969 and based in Sunnyvale, California, AMD had revenues of \$3.9 billion in 2001. (NYSE: AMD).

© 2002 Advanced Micro Devices, Inc. All rights reserved.

AMD, the AMD Arrow logo, AMD Opteron, AMD Athlon and combinations thereof, and AMD-8000, AMD-8111, AMD-8131, AMD-8151, AMD-761, and AMD-768 are trademarks of Advanced Micro Devices, Inc. HyperTransport is a trademark of the HyperTransport Technology Consortium in the U.S. and other jurisdictions. Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.