HyperTransport™ Technology: Simplifying System Design

Sean Cleveland
Scott Swanstrom
Chris Neuts

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

Table of Contents

Introduction

Designing an I/O bus requires more than a protocol definition and electrical interface specification. While these are two very important aspects of any bus architecture, other issues must also be addressed, such as system design considerations. In the past, a bus had to be made wider for higher performance through higher bandwidth. But the ever-shrinking desktop and server form factors required by IT and consumer markets has made this design methodology both unrealistic and expensive, both in terms of actual costs and power consumption.

The HyperTransport technology I/O link is a narrow, high speed, low power I/O bus that has been designed to meet the requirements of the embedded markets, the desktop, workstation, and server markets, and networking and communications markets. HyperTransport technology links are capable of extremely fast signaling with clock speeds of up to 800MHz, and Double Data Rate (DDR) memory signaling, to provide an effective throughput of 1.6 gigatransfers per pin-pair on a 32-bit link. The result is a maximum aggregate throughput of 12.8 gigabytes per second, per link.

HyperTransport technology provides high speeds while maintaining full software and operating system compatibility with the Peripheral Component Interconnect (PCI) interface that is used in most systems today. In older multi-drop bus architectures like PCI, the addition of hardware devices affects the overall electrical characteristics and bandwidth of the entire bus. Even with PCI-X1.0, the maximum supported clock speed of 133MHz must be reduced when more than one PCI-X device is attached. HyperTransport technology uses a point-to-point link that is connected between two devices, enabling the overall speed of the link to transfer data much faster. This also means that HyperTransport technology does not suffer from the overhead imposed by bus arbitration where resources need to be allocated inside of a shared bus.

The HyperTransport technology link is a "packetized" bus, which means addresses, data, and commands are sent along the same wires allowing designers to implement much narrower links. PCI, on the other hand, is a wider, slower bus that requires dedicated pins and traces for data, address, and sideband information, which is the upper and lower portion of the wave transmission. In a direct comparison, even though HyperTransport technology may take up to eight cycles of the bus to move as much data as PCI, the pure speed of HyperTransport technology more than makes up the difference, sending the data up to 48 times faster than PCI, providing a speed increase in a much narrower connection.

Routing narrower buses is far easier than routing wider buses. Narrower buses also reduce the need to add layers and additional costs to system board designs. In fact, HyperTransport technology has been designed specifically for lower-cost four-layer boards and Flame Retardant 4 (FR4) material, the fiberglass material used in most circuit boards. The true strength of HyperTransport technology is the ease at which system designers can integrate the technology into new and existing architectures. HyperTransport technology removes many of the constraints that impede high bandwidth, low latency bus architectures, while offering advanced capabilities at the same time.

This paper focuses on the benefits that HyperTransport technology gives system designers and the advanced capabilities that HyperTransport technology offers.

Foundation of HyperTransport™ Technology

Signaling Basics

The HyperTransport technology I/O link is two point-to-point unidirectional links employing differential signaling with speeds of 1,600 megatransfers or roughly 1.6 billion data transfers per second per pin pair. Differential signaling, a technique that uses two wires for each signal, with the result being the difference between the two signals sent, does not suffer from problems associated with the single-ended signaling of high-speed parallel buses (such as bouncing signals, interference, and cross-talk from adjacent signals.) With HyperTransport technology, these two wires or traces on a physical circuit

board are called the true and the complement, and they handle the routing of all commands, addresses, and data by way of the data path, a control signal, and one or more clock signals.

The type of differential signal used by HyperTransport technology is enhanced Low Voltage Differential Signaling (LVDS), which is actually lower in power than traditional LVDS because it uses a smaller voltage transition to transmit a signal. A lower voltage signaling was chosen over standard LVDS in order to guarantee support for future die size geometries and core voltages expected to be delivered over the next 7-10 years. Lower voltages also allow costs to be reduced and power requirements to be kept under control.

Maximum lengths of HyperTransport technology routes between chips can be upwards of 24 to 30 inches (0.6 to 0.75 meters) depending on the number of layers used in the printed circuit board (PCB). HyperTransport technology links can also be routed through card connectors and short cables for proprietary backplane purposes. This is because the combination of point-to-point links, differential signaling, and unidirectional data paths simplifies board-level electrical design, reduces latency and allows for greater distances between components.

Flexibility

The design benefits of HyperTransport technology make for a very compelling solution for a variety of different markets. HyperTransport technology has already achieved strong industry support for the full software compatibility with PCI and with all operating systems that use a PCI infrastructure. HyperTransport technology also offers the ability to serve as a mezzanine bus for PCI 66MHz/64-bit and PCI-X based systems. In fact, adding inline bridging to devices such as PCI, PCI 66MHz/64-bit, PCI-X, Gigabit Ethernet, InfiniBand architecture, and others is easy with device chaining supported by unique tunneling capabilities. A complete HyperTransport technology-based system consists of a processor with a HyperTransport port called a HyperTransport host, the HyperTransport bus (i.e., an input link and an output link) and any I/O channels connected to the HyperTransport bus.

Scalability

A significant advantage of HyperTransport technology is scalability. HyperTransport technology links can scale both in frequency (clock speed) and in the width of its data paths. This enables system designers to trade performance for power savings as required. HyperTransport technology data paths may be 2, 4, 8, 16, or 32-bits wide, and the data paths are not required to be symmetrical in width, meaning widths do not have to exactly match up. When devices supporting different clock speed and link widths are connected together, the HyperTransport technology protocol layer ensures the bus works correctly without the need for special I/O drivers. By offering variable asymmetric width, system designers can simplify PCB layout and reduce pin-count on application-specific integrated circuits (ASICs).

Adoption

HyperTransport technology has already made tremendous inroads into the networking and communications markets. In networking applications, HyperTransport technology allows manufacturers of routers, switches, and other network equipment to immediately extend the number of ports and to significantly increase the throughput of their equipment with little or no software or architecture changes. Further, HyperTransport technology provides a scalable network fabric with the ability to meet the needs of switching equipment for years to come.

Networking Extensions

The addition of the HyperTransport technology networking extensions adds a number of new features to the HyperTransport technology I/O specification. They include:

1. A message passing protocol that adds the ability to stream a sequence of packets to a given address.
2. The addition of 16 streaming point-to-point flow controlled virtual channels that support millions of end-to-end flow controlled individual streams and the ability to bridge Systems Packet Interface (SPI) 4.2 traffic that is typically used in communications data plane chips.
3. An enhanced error recovery protocol that automatically detects and recovers from data errors that are likely to occur when HyperTransport technology links become even faster in the future.

4. Support for direct peer-to-peer transfers where packets can be sent directly between peer devices without having to be reflected via the HyperTransport technology host device.

5. Support for 64-bit addresses that will be required for large memory models in excess of one terabyte.

6. Increased support for concurrent host transactions that allows up to 1K transactions to be outstanding, further improving flexibility and performance for communications applications.

7. A formalized method for defining HyperTransport technology hubs and switches within the standard and provides a roadmap for each of these device types.

HyperTransport™ Technology Supports Multiple Topologies

HyperTransport technology supports multiple connection topologies including daisy chain topologies, switch topologies and star topologies.

Daisy Chain Topology

The HyperTransport technology tunneling feature makes daisy chains of up to 31 independent devices possible. A HyperTransport technology tunnel is a device with two HyperTransport technology connections containing a functional device in-between. Essentially, the HyperTransport technology host initializes a daisy chain. A host and one single ended slave is the smallest possible chain, and a host with 31 tunnel devices is the largest possible daisy chain.
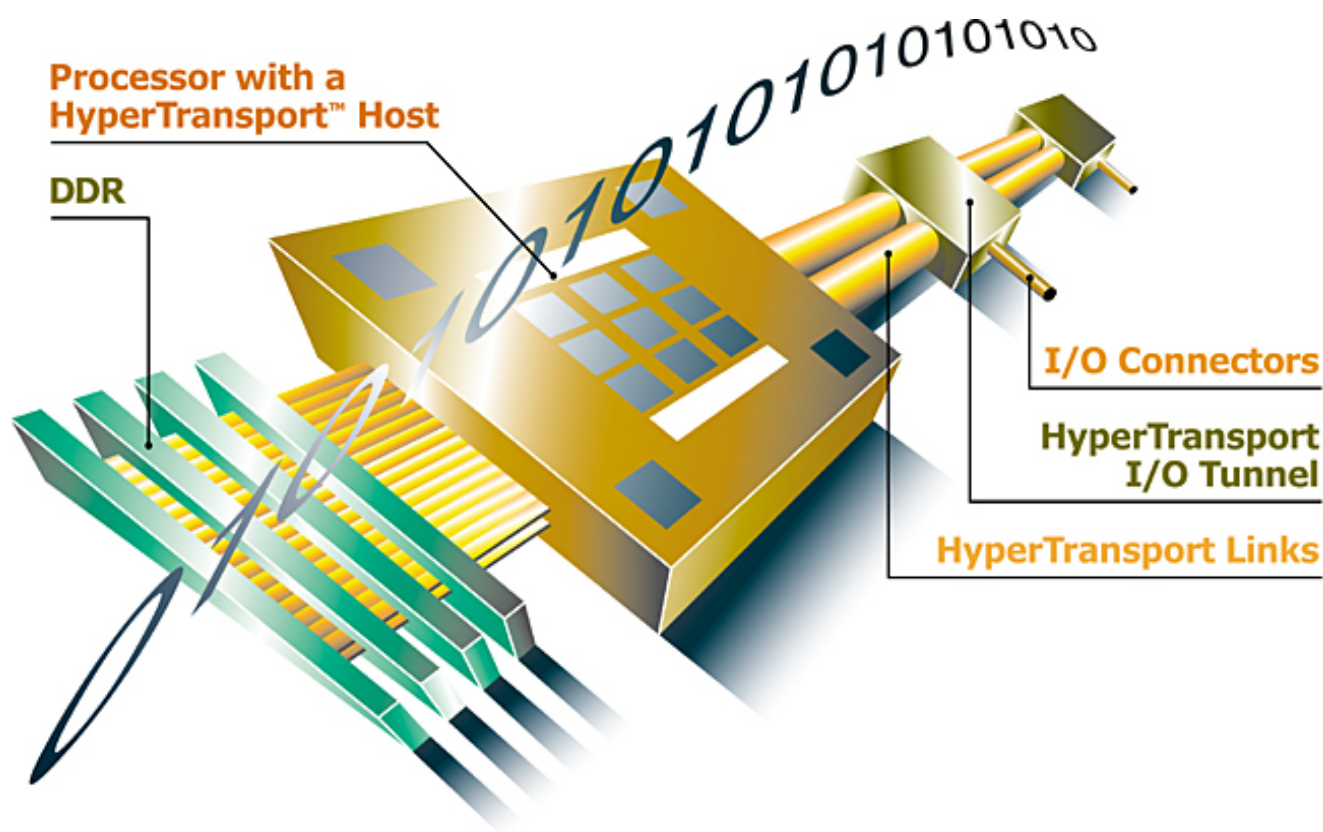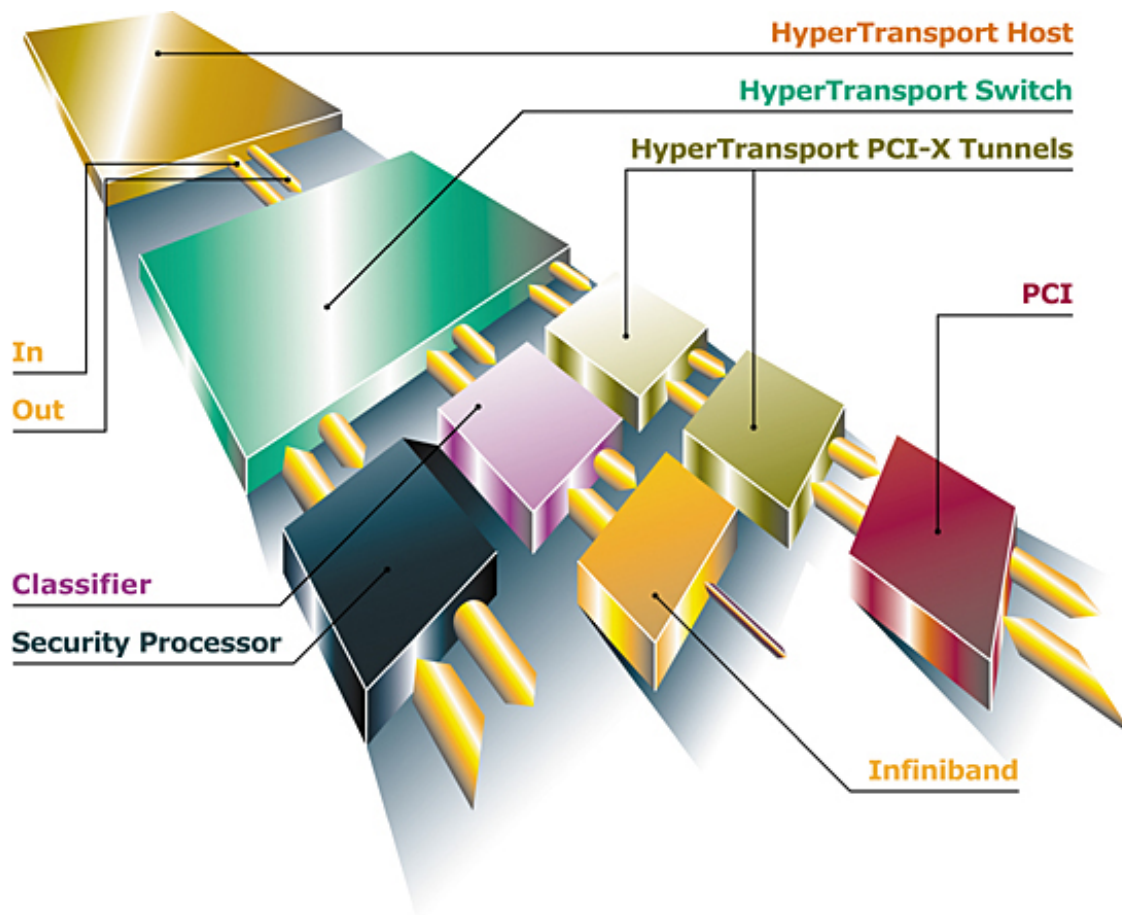
Figure 1: HyperTransport™ Technology Daisy Chain Configuration.

HyperTransport technology can route the data of up to 31 attached devices at an aggregate transfer rate of 3.2 gigabytes per second over an 8-bit HyperTransport technology I/O link, and up to 12.8 gigabytes per second over a 32-bit link. This gives the designer a significantly larger and faster fabric while still using existing PCI I/O drivers. In fact, the total end-to-end length for a HyperTransport technology chain can be several meters, providing for great flexibility in system configuration.

Switch Topology

A HyperTransport technology switch is a device designed for use in latency-sensitive environments supporting multiple processors or special-purpose processors. These processors are designed to increase available bandwidth, reduce latency, and support the effective use of multiple HyperTransport technology links running at different speeds. A HyperTransport technology switch passes data to and from one

HyperTransport technology chain to another. This allows system architects to use HyperTransport technology switches to build a switching fabric while isolating a HyperTransport technology chain from a host for offloading peer-to-peer traffic. By extending a fabric beyond a single chain of 31 tunnel devices, HyperTransport technology enables a multi-processor/host fabric, or creates a minimal latency tree topology.

Figure 2: HyperTransport™ Technology Switch Configuration.

In a switch topology, a device is inserted into the chain allowing it to be branched. This device is invisible to the original host controller, which believes the devices on the chain are daisy-chained. All devices stemming from the switch appear to be on the same bus while the intelligence of the switch determines bandwidth allocation. Branches

connected to HyperTransport technology switches may contain links of varying widths based on the discretion of the system designer.

The HyperTransport technology host communicates directly with the switch chip, which in turn manages multiple independent slaves including tunnels, bridges, and end-device chips. Each port on the switch benefits from the full bandwidth of the HyperTransport technology I/O link because the switch directs the flow of electrical signals between the slave devices connected to it.

A four-port HyperTransport technology switch could aggregate data from multiple downstream ports into a single high speed uplink, or it could route port-to port connections. For downstream chains that are connected to the switch, the HyperTransport technology switch port functions as a host port on that chain. So, for peer-to-peer traffic on that chain, the host reflection is done locally rather than having to be forwarded back to the actual host. This improves performance considerably.

HyperTransport technology switches can also support hot-pluggable devices. If slave devices are attached to switch ports via a connector, they can be hot-plugged while the rest of the HyperTransport technology fabric stays up. This also provides system designers the ability to reroute around failing does without having to shut the entire network down.

Star Topology

Whereas daisy chain configurations offer linear bus topologies much like a network "backbone," and switch topologies expand these into parallel chains, a star topology approach that distributes HyperTransport technology links in a spoke fashion around a central host or switch offers a great deal of flexibility.
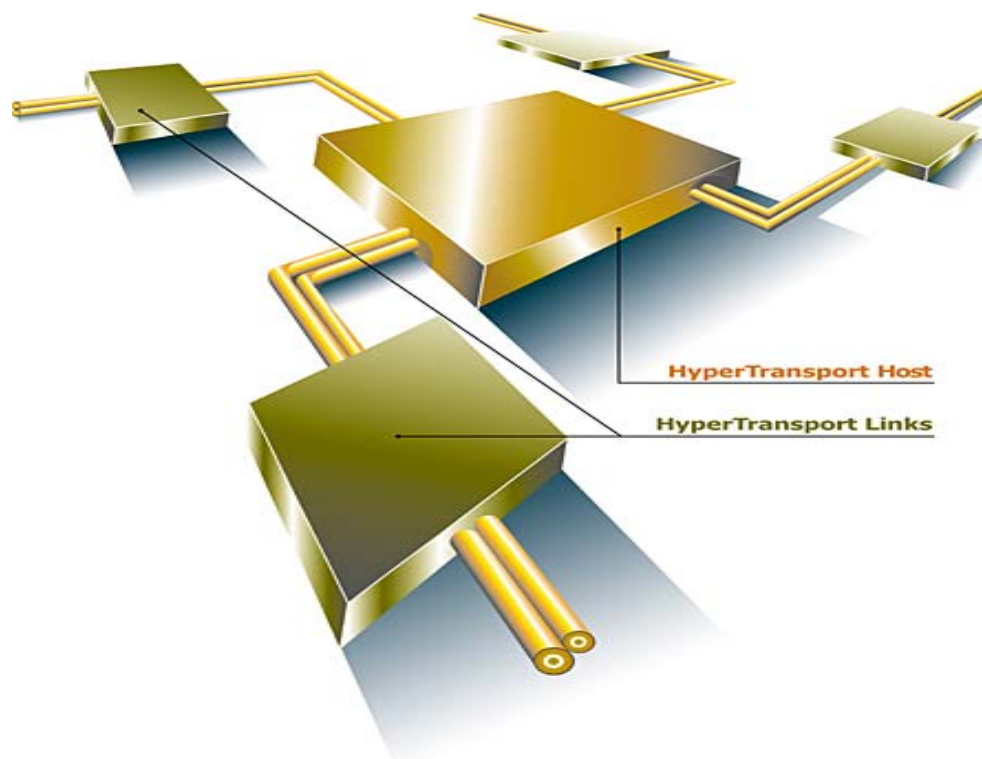
Figure 3: HyperTransport™ Technology Star Configuration.

With HyperTransport technology tunnels and switches, HyperTransport technology can be used to support any type of topology, including star topologies and redundant configurations, where dual star configurations are utilized to create redundant links.

## HyperTransport™ Technology Makes System Design Easier

Unlike bus designs that must support existing low cost connectors, long cables, and device sharing, HyperTransport technology has been designed as a chip-to-chip interconnect for moving data over short distances between chips in a system. It provides maximum bandwidth with minimum latency at low cost, and it achieves this by eliminating many of the constraints in system layout and design.

By architecting a new high-speed bus from the ground up, the designers of HyperTransport technology were able to take many aspects of board layout and implementation into consideration with the goal of lowering design costs and time to

market. HyperTransport technology takes a 180 degree turn away from today's wide, unwieldy bus architectures, to one that is scalable in size and speed, much narrower at comparable data throughputs, is lower in power due to native supply voltages that match deep sub-micron technologies, and far easier to implement.

Support for Multiple Packages

The HyperTransport technology interface built onto chips can be packaged in either ceramic or organic substrates with plastic pin grid array (PGA) or ball grid array (BGA) connectors.
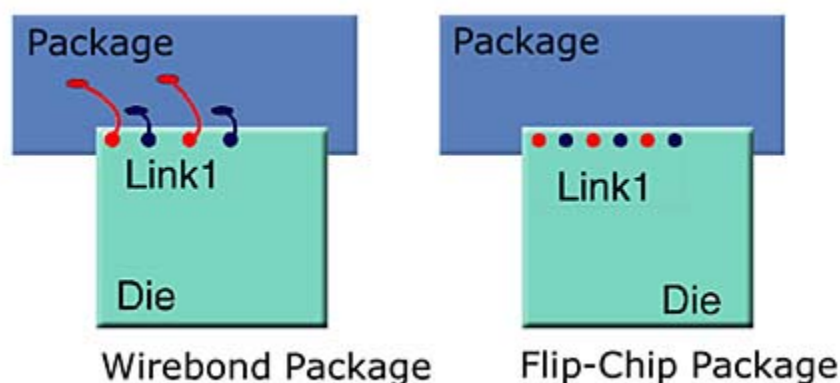

Figure 4: Differences Between Wirebond and Flip-Chip Packages.

Typically, chipset packages are plastic BGAs with the die wire bonded to them. A wire bond technique uses tiny wires to connect the bare die to the metal leads of the chip package. Alternatively, flip chip-pin grid array packaging can be utilized where the actual chip is turned over to reduce the inductance in both signal and power distribution paths.

For HyperTransport technology, flip-chip die attachment is recommended over wire bonded die attachment for all but the lowest link frequencies. The benefits include both lower inductance power connections made directly under the die, and reduced inductance and parasitic capacitance on the signals. This is required to achieve higher performance power delivery solutions, especially for chipsets with very fast HyperTransport technology interfaces where data rates of 800 megatransfers per second and above are required.

Easier to Layout

One of the key considerations to designing boards that can be reliably manufactured in large quantities is the elimination of "skew" as shown in Figure 5. Skew is the timing and phases anomalies in a transmission signal that occur when trace lengths are not matched. This is a big concern, especially in designs that incorporate high-speed I/O buses. By using an independent clock for every grouping of up to eight LVDS signals, skew compensation is significantly easier. In addition, clock recovery is a simple task that does not require large blocks of power-consuming clock recovery logic in the silicon.
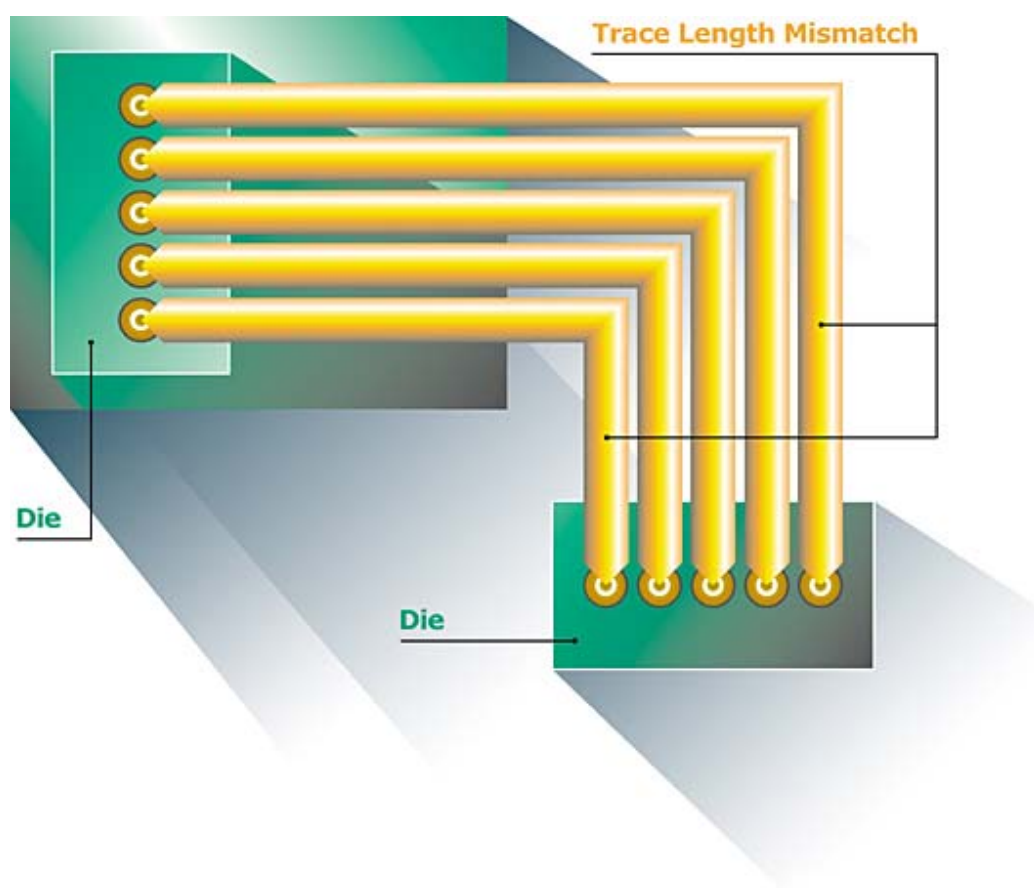
Figure 5: *Skew Caused by Trace Length Mismatch.*

Electrical Considerations

The issue with all I/O bus architectures is to keep the length bus traces on a motherboard equal so that the signals arrive at their destination with equal delays. The difference in the time it takes the signal to arrive between the two links is the amount of

skew built into the design. The higher the skew, the more likely data integrity will be compromised.

With regards to skew, typical board layout involves multiple designers maintaining a delicate balance between multiple factors. For example, the chip layout designer determines how best to keep skew controlled on the chip while the package designer worries about skew on the substrate with regards to the pad location and where the chip will be placed on the circuit board. Finally, the board designer needs to build a board with the least amount of mismatch while keeping propagation delays equal, reducing electrical issues, and minimizing skew in the routes by keeping them the same length if at all possible, because if the traces vary in length, the layout designer is required to embed "snake-like" patterns, often called "swizzles," along the trace paths, which wastes even more board real estate. This is not the most efficient process.

Packaging

Physical implementations of HyperTransport technology has been made easier than other high-speed I/O buses by simplifying bus layout. This is achieved by building constraints into the packages and designing a HyperTransport technology route layout from die pad to die pad instead of letting the designers do all of the work. This is called naturally compensating trace length matching.
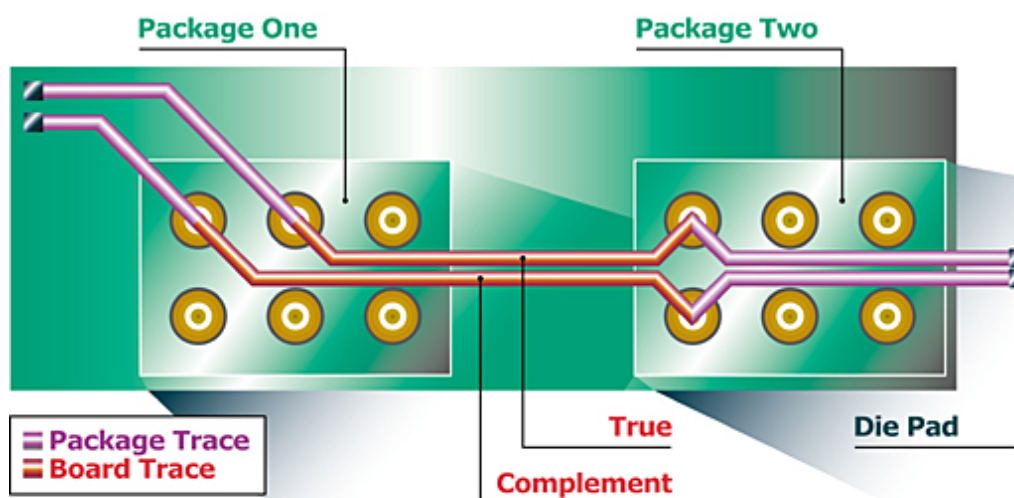
Figure 6: *Naturally Compensating Trace Length Matching.*

HyperTransport technology was designed to use naturally compensating trace length matching. As shown in Figure 6, the transmitter and receiver are physically designed such that the package traces are routed to lengths either matched or mismatched based upon the device ball out or pin out. Inside the package, a true is connected to a ball one row further away from the die than the complement in such a way that there is actually a mismatch in the package. When the package is then placed on a circuit board and the trace is run to another package, that mismatch is offset exactly by the mismatch on the board. This creates a zero mismatch and reduces skew inherent in traditional designs.

The end result is the designer only needs to route directly from one package to another with minimal or no trace swiveling to guarantee trace length matching requirements. This makes HyperTransport technology a very straightforward, easy and economical bus to route. With other fast I/O bus architectures the actual physical space taken by the link would need to be expanded to meet the link matching rules implemented, which consume precious board space and complicate the routing of other signals. With HyperTransport technology, traces can be literally laid down in straight lines between packages while still meeting the very stringent requirements for HyperTransport technology. If links make the same number of left turns as right turns between packages on the board, the length requirements have again been met.

Of course, not all packages are mismatched. If a true and a complement trace are placed on pins parallel to the package edge, then the package trace lengths have been matched. In this case, the resulting system board trace lengths would naturally have minimum natural mismatch due to the package breakout routing. While the goal of HyperTransport technology layout design is to make sure the links from the die pad to the die pad on one signal are the same as the links on the other signal, there can be some flexibility. The HyperTransport Technology Design Guide does provide some flexibility with regards to package and board matching if the design is better served and physically works better.

Signal Routing

When system designers lay out boards, the extra length of board traces connecting pins in multiple packages must be dealt with in a way that minimizes board space while meeting matching requirements.

Figure 7 depicts an actual board design and illustrates how much easier HyperTransport technology is to route compared with AGP, a connector technology exclusively designed for video cards and the leading video connection technology used in PCs today.
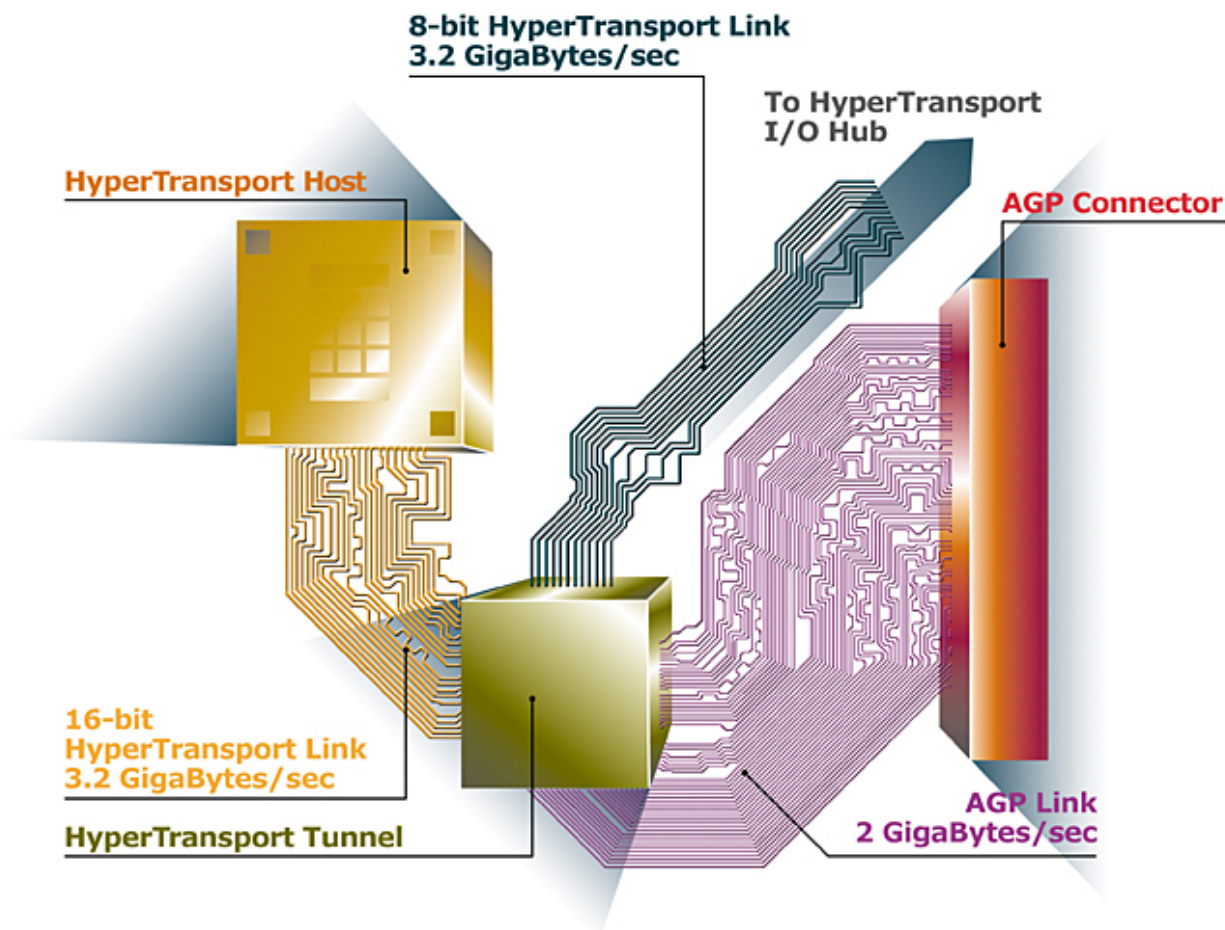
Figure 7: *HyperTransport™ Technology Traces Compared to AGP Traces.*

AGP was designed with two purposes:
1. To relieve the PCI bus of work with graphics data
2. To facilitate better bandwidth within the video system

HyperTransport technology:
- Is faster than AGP
- Has smaller matching requirements
- Is far easier to implement.

The same comparison can be made with PCI and most other high-speed I/O buses.

Reduction of Electrical Noise

HyperTransport technology uses a low voltage differential signaling (LVDS) transmission method for sending information. LVDS is a low noise, low power, low amplitude method for high-speed data transmission over copper wire. The differential nature of HyperTransport technology increases noise immunity and noise margins and allows protection to traditional single-ended noise sources including AC return paths and cross-talk aggressors.
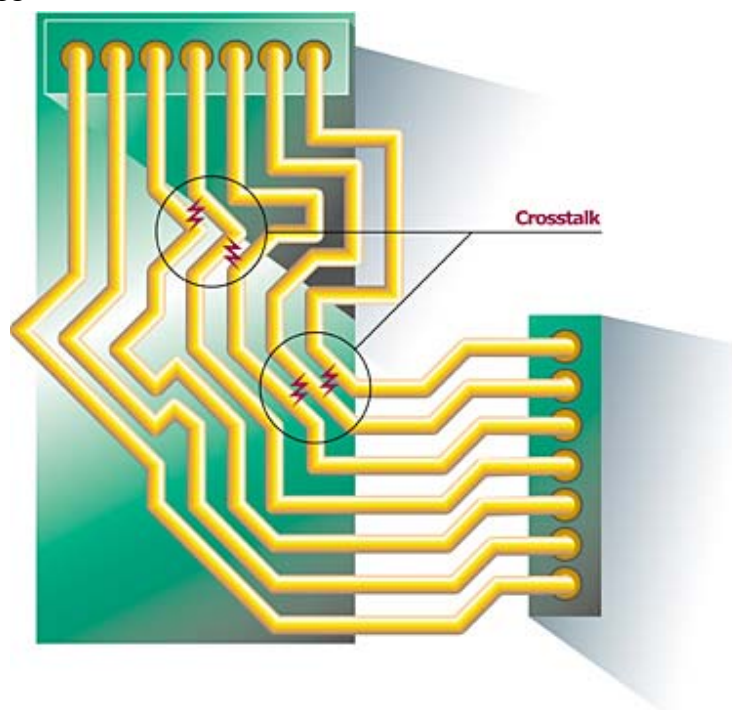

Figure 8: An Example of Crosstalk.

As shown in Figure 8, crosstalk is the disturbance caused by the electric or magnetic fields of one signal affecting the signal in an adjacent link. Whereas crosstalk is an inherent concern with other high-speed buses, it is a reduced noise source for HyperTransport technology. This impacts physical implementations in the way trace and space rules need to be implemented, but reduces susceptibility to plane split crossings where two signals cross a break in the ground plane and cause a higher AC impedance in

the return path or in layer to layer interactions. These are just some of the benefits HyperTransport technology gives designers in the area of electrical design.

Better Grounding

The HyperTransport technology I/O link receives return current both through the complement link and through a return ground link that typically resides in one of the middle layers of the circuit board. While most high-speed buses also have a return ground path that the electrical current flows through, a mechanism is not present through which current can flow if the ground is split or broken.
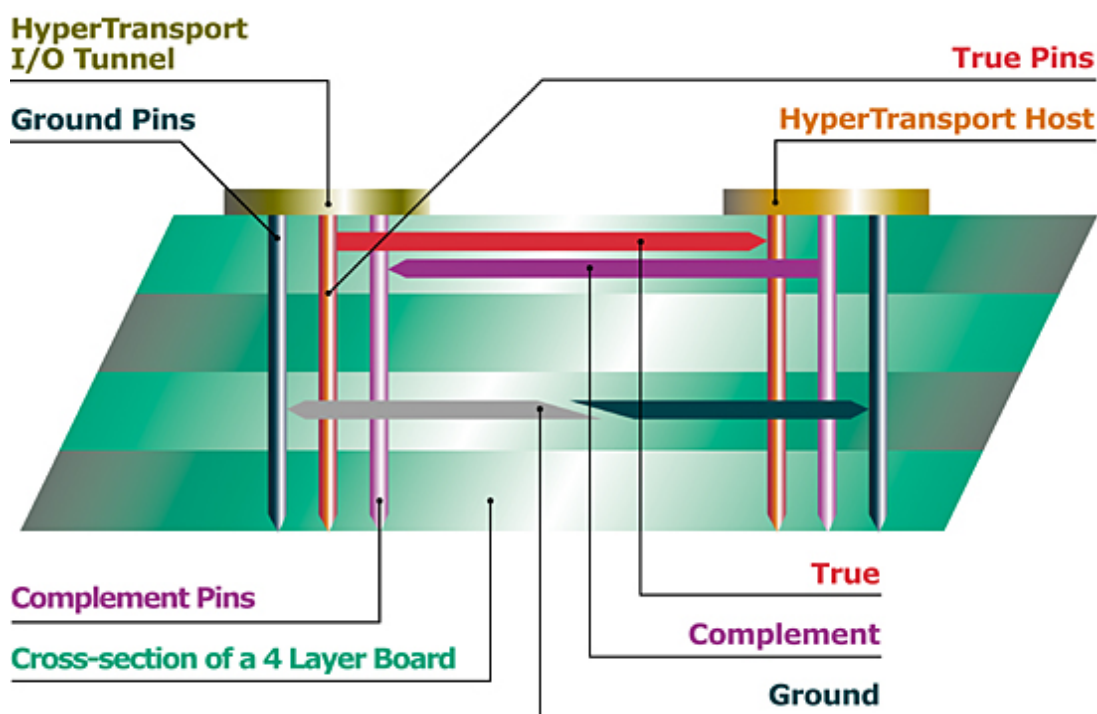
Figure 9: *Return path can use compliment if ground is broken.*

A split or broken ground reference path creates high impedance that can have a devastating effect on the stability of the bus. By using the complement as a secondary return path, the layout designer does not need to be as strict about maintaining continuous ground reference paths.

Elimination EMI Issues

LVDS is easy to implement on system boards and has no inherent Electromagnetic Interference (EMI) issues due to its low power requirements. Designers should continue to use good high-frequency design practices in order to minimize EMI. Addressing potential EMI issues early in the design process is also important as fixing EMI problems late in the design process is usually more difficult and more expensive.

The return path immunity of HyperTransport technology acts as a shield against EMI as well. The biggest culprit for EMI is when the trace that the signal sends to a chip and the path that the return wave sends back forms a circle of current. This circle of current induces a magnetic field and creates EMI. To battle this, both paths should be kept in parallel so there is no circle configuration. Since the HyperTransport technology I/O link has the true and complement right next to each other, if there is any circular return path through the ground plane, the current will be seen as high impedance and all of the return path will be redirected back through the complement instead, effectively eliminating EMI.

Easier to Power

HyperTransport technology is built such that the power supplied to a link can be attached to either side of the link rather than to both sides of a link as is required of other buses. This makes for a very economical board and reduces the need to route power across traces.

Power is supplied to both sides of the link by hooking up a regulator on one side and decoupling it on the other or vice versa. Because the current power requirements of HyperTransport technology are low enough, system flexibility is increased because the power regulator can be placed where it fits without having to fight the issue of placing it directly next to the chip or on a side that is hard to route to.

On-Chip Termination

Another benefit of HyperTransport technology is the use of on-chip termination as opposed to discrete termination resistors. Discrete termination resistors, commonly mounted on the top surface of a system board, require so much space that they must be staggered with surface traces routed around the bodies of the resistors. This type of configuration tends to result in a high number of unequal trace lengths. To compensate

for this problem, the layout designer is required to embed swizzles along the trace path, reducing board real estate. In addition, this procedure must be accomplished while maintaining adequate trace spacing to minimize signal crosstalk.

By contrast, HyperTransport technology uses on-chip, differential termination with both receiver and transmitter matching interconnect impedance. The on-chip termination saves significant board area and eliminates difficult to place components, greatly simplifying the board designer's task. This also results in trace lengths that can be more equal due to the inherent nature of using a standard semiconductor package with consistent leads and spacing. Also, the use of LVDS lowers the voltage across the termination resistors and lowers the overall power dissipation, reducing overall power consumption.

Simpler Power Distribution

Power distribution for HyperTransport technology links is very adjustable. Normally, signals are transmitted higher compared to the ground. Since HyperTransport technology uses LVDS, the differential nature of the signal allows the high and low states to change. When these states change, there is a point where both states become the same. The result can be a spike is that sent across the ground.

HyperTransport technology has been designed to overcome this issue by using adequate power decoupling. Through a power supply distribution network, power is supplied from the pin boundary to the die through the package. This practical power distribution network for HyperTransport technology include the on-die decoupling capacitance, package and motherboard interconnects, optional on-package and on-motherboard discrete decoupling capacitors, as well as the voltage regulator with the associated bulk decoupling capacitors. And again, because the HyperTransport technology power supply is not as demanding as the core supply or the power supplies for other I/O interfaces, several layout alternatives may be used in the design. However, since HyperTransport technology is a very high data-rate I/O interface, careful design practices must be exercised to ensure a low DC and AC impedance path from the regulator to the die.

Unlike other I/O interfaces, the HyperTransport technology interface may consist of several independent HyperTransport technology links on a single die. The power to

each link on the die must be delivered through the package by an independent interconnect in order to provide maximum isolation between the individual links. This will minimize the high frequency noise on the individual HyperTransport technology power supplies due to the switching of high data-rate signals on other HyperTransport technology links.

Summary

By establishing a point-to-point link between two devices, HyperTransport technology enables integrated circuits to exchange information at much higher speeds than currently available interconnect technologies.

HyperTransport technology provides a number of important features:

- The tunneling capability enables up to 32 devices to be connected in a chain which can serve as a backplane or system bus.
- Since HyperTransport technology-enabled devices exchange packets of data, these devices require significantly fewer traces or wires to connect signals, greatly reducing system design complexity.
- The use of enhance low-voltage differential signaling helps to eliminate many of the problems associated with single-ended signaling, including crosstalk and EMI.
- HyperTransport technology extends the lifespan of PCI by providing full backwards compatibility for PCI software, drivers, and operating systems while eliminating bottlenecks and providing the bandwidth necessary for future high-speed chips and interconnect standards.
- HyperTransport technology's low-voltage differential signaling gives system designers efficient power usage and power manageability to design smaller low-cost boards.
- Support for asymmetrically links and scalability in speed, width, frequency, and direction allows system designers to have the flexibility to trade performance for cost and power savings where it is appropriate.
- HyperTransport technology extends the lifespan of devices by providing enough bandwidth headroom to allow system designers to add new high performance parts.

- HyperTransport technology makes for reusable board designs as the addition of new chips can be done without the need to redesign board for placement or the need to make traces smaller, further cutting design cycles. And given initial measured results on different types of systems that use HyperTransport technology, very good timing margins have been observed which may equate to higher board yields.

By offering aggregate data rates of 12.8 gigabytes per second, per link over multiple high-speed scalable fabrics including daisy chain topologies, switch topologies, and star topologies, the HyperTransport technology I/O link provides the bandwidth needed for scalable, single and multi-processor servers and next generation network equipment. New HyperTransport technology networking extensions, including peer-to-peer communications and 64-bit addressing, are being developed by some of the largest and most influential networking companies in the silicon world. These enhancements to the HyperTransport technology protocol are expected to stimulate further adoption of HyperTransport technology in networking applications.

HyperTransport technology is licensed royalty-free to all members of the HyperTransport Technology Consortium. More information about HyperTransport technology can be found by visiting the consortium's web site at www.hypertransport.org.

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).