

HyperTransport™ Consortium

Applications Overview White Paper

Meeting the I/O Bandwidth Challenge: How HyperTransport Technology Accelerates Performance in Key Applications

December 2002

The HyperTransport Consortium

www.hypertransport.org

The HyperTransport Technology Consortium disclaims all warranties and liability for the use of this document and the information contained herein and assumes no responsibility for any errors that may appear in this document, nor does the HyperTransport Technology Consortium make a commitment to update the information contained herein.

DISCLAIMER

This document is provided “AS IS” with no warranties whatsoever, including any warranty of merchantability, non-infringement, fitness for any particular purpose, or any warranty otherwise arising out of any proposal, specification or sample. The HyperTransport Technology Consortium disclaims all liability for infringement of property rights relating to the use of information in this document. No license, express, implied, by estoppels, or otherwise, to any intellectual property rights is granted herein.

TRADEMARKS

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

Modern system architectures require greater I/O performance

The need for speed

The demand for faster processors, memory and I/O is a familiar refrain in market applications ranging from personal computers and servers to networking systems and from video games to office automation equipment. Once information is digitized, the speed at which it is processed becomes the foremost determinate of product success. Faster system speed leads to faster processing. Faster processing leads to faster system performance. Faster system performance results in greater success in the marketplace. This obvious logic has led a generation of processor and memory designers to focus on one overriding objective – squeezing more speed from processors and memory devices.

Processor designers have responded with faster clock rates and superpipelined architectures that use level 1 and level 2 caches to feed faster execution units even faster. Memory designers have responded with dual data rate memories that allow data access on both the leading and trailing clock edges doubling data access. I/O developers have responded by designing faster and wider I/O channels and introducing new protocols to meet anticipated I/O needs. Today, processors hit the market with 2+ GHz clock rates, memory devices provide sub 5 ns access times and standard I/O buses are 32- and 64-bit wide, with new higher speed protocols on the horizon.

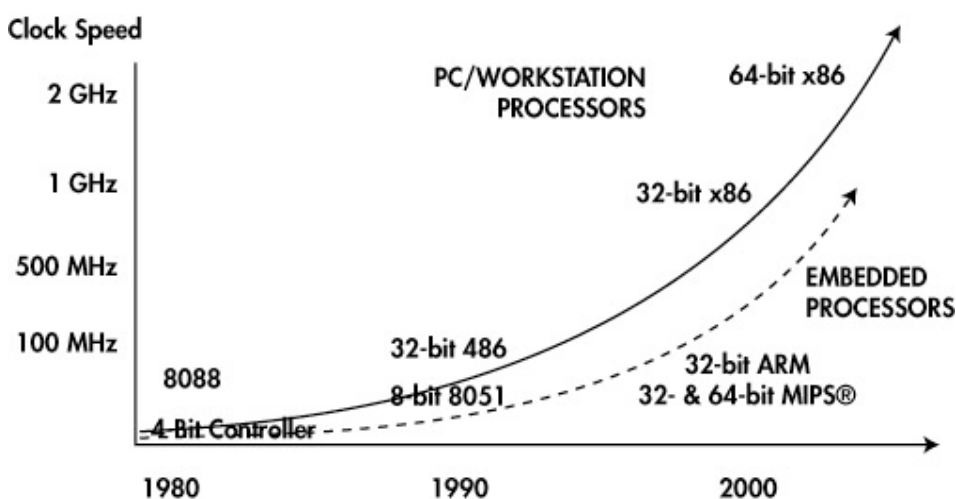


Figure 1: Processor clock speeds have driven the performance gains of the past two decades.

So where's the problem?

Speed Kills

Increased processor speeds, faster memories, and wider I/O channels are not always practical answers to the need for speed. The main problem is integration of more and faster system elements. Faster execution units, faster memories and wider, faster I/O buses lead to crowding of more high-speed signal lines onto the physical printed circuit board.

One aspect of the integration problem is the physical problems posed by speed. Faster signal speeds lead to manufacturing problems due to loss of signal integrity and greater susceptibility to noise. Very high-speed digital signals tend to become high frequency radio waves exhibiting the same problematic characteristics of high-frequency analog signals. This wreaks havoc on printed circuit boards manufactured using standard, low-cost materials and technologies. Signal integrity problems caused by signal crosstalk, signal and clock skew and signal reflections increase dramatically as clock speed increases.

The other aspect of the integration problem is the I/O bottleneck that develops when multiple high-speed execution units are combined for greater performance. While faster execution units relieve processor performance bottlenecks, the bottleneck moves to the I/O links. Now more data sits idling, waiting for the processor and I/O buses to clear and movement of large amounts of data from one subsystem to another slows down the overall system performance ratings.

New protocols abound, but compatibility and practicality are key issues

Chip-to-chip, inside-the-box, and outside-the-box I/O protocols are designed to meet different types of application requirements. For example, onboard communication links such as those between processor and local memory devices need high bandwidth and very low latency. WAN (Wide Area Network) I/O protocols such as TCP/IP and InfiniBand require high bandwidth for large amounts of data traffic but latency is not as important as reliable error detection and correction facilities. Inside-the-box protocols such as PCI define not only bus features such as clock speeds, signal characteristics and bus widths but also

important board level characteristics such as physical board form factors and connectors.

In attempting to address the need for greater I/O system speed, there have been several efforts to either extend existing I/O technologies or to introduce entirely new ones. The marketplace will eventually decide what standards emerge and which technologies survive the test of time. But, it is clear that two important issues must be resolved: a successful high-speed I/O technology must be compatible with both legacy I/O and emerging media and technologies and it must be practically applicable to real-world systems that employ extremely fast processors and memories and that interconnect large data streams to and from I/O links and processing subsystems.

Peripheral Component Interconnect (PCI) is the leading traditional board-to-board I/O technology and is an inside-the-box protocol that defines bus structures, board form factors and connector standards. PCI is multi-drop, parallel bus that defines a device configuration and addressing protocol that simplifies the interface to specialized I/O devices and is processor independent. PCI supports a robust system initialization, discovery and setup. During initialization, the system level operating system can discover all attached PCI compatible components, allocate resources and configure the I/O devices. The PCI load and store architecture and flat addressing space enables the easy configuration of I/O devices. Because of this robust and flexible architecture, PCI has been incorporated into a wide variety of products used in the personal computer and server markets. Its popularity has also extended to a number of embedded applications. Consequently, there are many chip and system developers with an extensive investment in PCI hardware and software driver technology.

Standard PCI is defined as a 32- or 64-bit bus with a 33 or 66 MHz clock rate that is no longer adequate for GHz and up processor clock rates. To extend PCI, the PCI Special Interest Group or PCI-SIG first defined a PCI-X 1.0 specification with clock rates boosted to 133 MHz, and in the second half of 2002, introduced the PCI-X 2.0 specification with clock speeds defined to 533 MHz. The problem with high-speed PCI is the legacy of wide 32- and 64-bit bus definitions. As noted previously, the speed issues associated with parallel buses apply to these fast,

wide buses. So the tradeoff is how much design and manufacturing cost can be absorbed in pursuit of near-compatibility with legacy I/O technology. In some cases, it may be feasible, but especially in emerging embedded applications, the cost and difficulties of supporting higher speeds in constrained I/O technologies may limit their application.

The industry has seen several new I/O technologies proposed to overcome these constraints including PCI Express, RapidIO, InfiniBand and HyperTransport.

Built on top of the PCI model, PCI Express, formerly known as 3GIO, has been proposed as a new layered architecture for supporting high-speed media. While the software model is compatible with PCI, the physical media and interface specification is completely different and features an embedded clocking scheme, QoS (Quality of Service) features, isochrony, RAS, and hot insertion/surprise removal support. A key attribute of the PCI Express model is the introduction of low-voltage, differentially driven pairs of signals, a packetized transaction protocol and a set of software and physical abstraction layers that insulate the physical media interface from the higher order software. While it has been proposed as a solution for applications as diverse as chip-to-chip interconnect, board-to-board buses, graphics I/O channels, and as an interconnect to 1394b, USB 2.0, Ethernet and InfiniBand channels, it is not yet clear how PCI Express will emerge or whether it will become as widely adopted as PCI. According to PCI-SIG documents, first devices are currently targeted to be available in the latter half of 2003.

One still evolving high speed I/O interconnect proposal is RapidIO. RapidIO, like PCI Express, has been proposed for both chip-to-chip and board-to-board applications. Although first proposed in 1997 as a fixed 8-bit or 16-bit parallel I/O scheme with little PCI compatibility, RapidIO has been extended to include a serial physical layer definition as well with some characteristics designed to support board-to-board signal requirements and a PCI-like configuration method. Like PCI Express, RapidIO uses packetized transactions, but unlike PCI and its extensions, RapidIO uses source-based addressing.

Unlike PCI, and as one of its most interesting attributes, RapidIO has seriously tried to meet the needs of many diverse applications by incorporating a large

number of different technologies and approaches. For example, the parallel specification is targeted to more tightly coupled processor-to-processor type transactions while the serial specification is targeted more to subsystem-to-subsystem and backplane applications. Within the serial specification, two types of transmitters have been defined, one is low power for short interconnects and one is high-power for longer connections. While it is obvious that the high power requirements of a long, backplane compatible transmission line is overkill for the short, chip-to-chip connections needed on a printed circuit board to connect processor-to-processor or processor-to-I/O channel, it is not clear that it serves the industry to have one “multi-faceted standard” for all possible applications. It remains to be seen whether the evolution of the RapidIO specifications will encourage wider adoption or limit its use due to fragmentation and loss of clear application focus.

One area where the RapidIO parallel approach with its ability to support many peer-to-peer transactions has proved useful is in DSP farm applications. DSP farms are one application where PCI compatibility is not an issue.

Among the “outside-the-box” interconnect technologies, InfiniBand, promoted by the InfiniBand Trade Association, has received the most attention. InfiniBand is designed to provide a reliable and easily scalable outside-the-box technology for large-scale storage and server networks. Designed to bring a single unified I/O fabric solution to the Internet data center, InfiniBand is an interconnect fabric that can connect disk arrays, storage area networks, local area networks, servers and server clusters. It provides a channel-based I/O scheme with host channel adapters and target channel adaptors. Channel adapters are intelligent I/O engines that employ InfiniBand links at data rates of 2.5 Gbps and up using both copper wires and fiber optic media. With intelligent I/O engines linked with high-speed data channels, the data center architecture can evolve from centralized servers with limited I/O capabilities to processing engines and I/O engines connected by the InfiniBand fabric. I/O resources can be accessed across the fabric through target channel adapters and shared between many servers. This approach takes the “nuts-and-bolts” connection work out of the equation and gives the data center manager a cleaner, and higher level means of growing and expanding high-speed server systems.

InfiniBand technology is a point-to-point interconnect based on a channel-based architecture that depends upon intelligent target and host nodes that use a message-passing protocol. Because it is targeted to data center applications, the InfiniBand fabric is designed to support multiple, redundant links between systems. InfiniBand links provide 2.5, 10 and 30 Gbps bandwidth in each direction and can link host channel adapters (processor-based platforms), target channel adapters (I/O devices), switches (connecting link to link) and routers (subnets to subnets).

InfiniBand has garnered a great deal of support from industry leaders in the storage network, server and data center markets and it is clear that it will be, at the very least, one of the high speed technologies that will impact the future direction of the industry. Like any other emerging technology, its ultimate market acceptance will depend upon its ability to meet the needs of its targeted application in a cost-effective and easy-to-adopt form.

Additional LAN technologies such as TCP/IP over Gigabit Ethernet, SPI-4.2 and Fiber Channel are widely utilized in data centers and the Internet backbone and provide 1 to 10 Gigabit/second data channels. Easy connection to high-speed data links such as these are a requirement for next generation systems of all types.

HyperTransport™ Technology – A universal chip-to-chip I/O solution

With the proliferation of pumped up traditional I/O technologies and new proposed standards, why has AMD, Apple, Broadcom Corporation, Cisco Systems, nVIDIA, PMC-Sierra, SGI, SiPackets, Sun Microsystems, and Transmeta led a consortium of industry leaders in developing HyperTransport technology?

The goal of HyperTransport is not to supplant other I/O technologies, but rather to provide a highly focused and standardized chip-to-chip interconnect that not only meets the data transfer requirements of the core onboard processing, memory and I/O elements, but that can also easily connect to both low speed traditional I/O devices and high speed new media I/O channels.

Industry leaders in processors, networking, telecommunications, software, FPGAs, silicon IP, and development tools have recognized the advantages of HyperTransport technology as a potentially universal chip-to-chip interconnect and it has rapidly gained momentum and support. As of mid-2002, the HyperTransport Consortium, the association that steers the technology, passed the 50-member milestone and there were over 35 announced HyperTransport-enabled products. Systems using HyperTransport technology that were shipping in high volume included personal computers from leading personal computer makers Compaq, Hewlett-Packard, and NEC and the Xbox game console from Microsoft.

Key HyperTransport Consortium Promoters	
Company	Industry
Alliance Semiconductors	Networking devices
AMD	High performance processors
Apple	Personal computers
Broadcom Corporation	Networking and high performance processors
Cisco Systems	Networking
NVIDIA	Graphics
PMC-Sierra	Networking devices and high-performance processors
SGI	High-end workstations and graphics software
Sun Microsystems	Workstations
Transmeta	Low-power processors

Key Announced HyperTransport-enabled Products	
Product	Application Area
Ali M1685/M1563	Northbridge/Southbridge chipsets
SiPackets SP1011 HyperTransport to PCI Bridge and 4-way Switch	HyperTransport to PCI bridge and HyperTransport Switch device for personal computer, server and embedded HyperTransport applications
AMD 32-bit Athlon	Personal Computers and Servers

processor, 64-bit Opteron processor, AMD-8111 HyperTransport Hub, AMD-8131 HT PCI-X Tunnel, AMD-8151 HT AGP 3.0 Graphics Tunnel	
American Megatrends	AMIBIOS8 BIOS for Athlon and Opteron processors
Broadcom single, dual and quad 64-bit MIPS® processors	Networking equipment and high performance embedded systems
Cavium NITROX	Security macroprocessor for networking equipment
FuturePlus FS2240 HT Analysis Probe	Logic analyzer connection for link verification and performance analysis
GDA Technologies OperVera Verification IP and HyperTransport Tunnel Core	Prebuilt and reverified modules for generating HyperTransport transactions and tunnel IP
Hifn HIPP III Security Processor	Security protocol processor for networking equipment
NPTest ITS 9000ZX	High-performance HyperTransport test system
Nurlogic AmberBridge HyperTransport PHY core	Silicon IP for integration into ASIC and SOC devices
NVIDIA nForce IGP and MCP	Integrated Graphic Processor and Media and Communication Processor – graphics coprocessing platform for next generation PC graphics and media processing
Phoenix Technologies ServerBIOS	BIOS for 32- and 64-bit x86 processors for personal computers and servers
PLX Technology PowerDrive HT7520	HyperTransport to PCI-X bridge device with 6.4 GB/s bandwidth and 8- and 16-bit HyperTransport interface
PMC-Sierra RM9000x2 Dual 64-bit MIPS® Multiprocessor	Dual 1 GHz multiprocessor for networking equipment

Teradyne J973EP VLSI Test System	HyperTransport tester
TransEDA HyperTransport Verification Suite	HyperTransport speed verification process
VIA Technologies K8 Series Chipset	Chipset for high performance AMD processors
Xilinx HyperTransport core	Single-ended slave IP core for Vertex®-II FPGAs

HyperTransport technology - A solid technical foundation that quickly evolved into a universal onboard interconnect solution

The HyperTransport technology specification development illustrates the benefit of starting with a non-denominational core technology definition and encouraging participation by leading industry developers. Unlike many proprietary technologies that were presented as “open standards,” HyperTransport has achieved remarkable industry success because it was quickly transferred from AMD and its partners to a general industry consortium that defined a clear specification and opened the technical development to member companies. In addition, HyperTransport technology is provided royalty-free with just an administrative cost for the license itself (included in the low-cost membership in the consortium).

Originally conceived as a high bandwidth solution for chip-to-chip communications in high performance servers and personal computer workstations, HyperTransport technology was shaped by the early participation of leading networking and telecommunications providers. Thanks to evolutionary progress in the protocol specification (as opposed to revolutionary revisionism), it is now poised to be a universal chip-to-chip standard for embedded systems as well.

The initial HyperTransport technology developers were not driven by a specific processor architecture, but rather with the goal of developing a “processor-agnostic” I/O architecture that could become a universal onboard interconnect. Consequently, HyperTransport developers reviewed not only how increased

bandwidth could be obtained, but also at how this bandwidth could be implemented in practical, cost-effective and complementary ways to existing and emerging I/O technologies.

The first HyperTransport technology specification defined a scalable point-to-point 1.6 GigaTransfers/second signal rate technology that not only provided a high-bandwidth processor-to-memory, processor-to-processor, and processor/memory to I/O interconnect, but was software transparent to legacy PCI. While providing headroom to support emerging technologies such as InfiniBand, 10 Gigabit Ethernet and SPI-4, it could still be implemented using popular, low-cost four layer PCB manufacturing technologies.

Key HyperTransport Technology Characteristics	
Characteristic	Description
Bandwidth	Scalable to 12.8 Gigabytes/second aggregate bandwidth
Clock Rate	200, 300, 400, 500, 600 and 800 MHz
Data Rate	Dual Data Rate resulting in up to 1.6 GigaTransfers per second at an 800 MHz clock rate
Physical Link	Point-to-point, unidirectional data pairs using enhanced low-power 1.2-V LVDS (Low Voltage Differential Signaling) with low EMI and crosstalk and low power consumption
PCB Implementation	Designed to support common, low-cost four-layer Printed Circuit Board manufacturing flows with high signal integrity. Naturally compensating trace length matching simplifies PCB layout and reduces trace swiveling, on-chip differential termination saves board space and reduces discrete component costs.
Legacy I/O Compatibility	PCI software compatible and operating system transparent
Emerging I/O Compatibility	Bandwidth capable of direct connect to 10 Gigabit Ethernet, SPI-4.2, InfiniBand and other off-board links such as Fibre Channel, serial ATA, and iSCSI.
Data Payload Protocol	Packet-based in multiples of four bytes with shared link for data, addresses and commands

Scalability	Exceptional scalability, with asymmetrical links and variable clock rates and data path widths. 2, 4, 8, 16, and 32 bit wide links are supported with variable clock speeds. Bandwidth can be applied in the exact amount needed by the subsystem. There is bandwidth headroom to support future technologies.
Topologies Supported	Daisy-chained, star and switch topologies
Maximum Length	Up to 30 inches (0.75 meters) depending on PCB technologies employed
Error recovery	Automatic error detection and recovery handling for high reliability systems featuring large number of high speed HyperTransport links
Networking Extensions	Message passing protocol to support streaming of packets, 16 streaming point-to-point flow controlled virtual channels, enhanced error recovery, peer-to-peer transfers, 64-bit addressing for very large memories, concurrent host transactions support.

From the first specifications release, the combination of high bandwidth, scalability, PCI compatibility and low-cost, yet reliable physical implementation won converts immediately in the personal computer, server, communications and embedded market spaces. With the support of additional Consortium members, additional refinements were added to define the operation of HyperTransport switches and hubs, to support PCI-X bridges, and to add 64-bit addressing to support very large-scale systems. In 2002, a group of “networking extensions” were added to Revision 1.10 of the specification that attracted additional converts in the telecommunications and networking space. Networking extensions include enhanced error recovery, peer-to-peer linking, additional virtual channels with 16 priority levels, and new messaging protocols that enable packet streaming.

The combination of the robust definition of the original specification, the addition of communications-oriented features, and widespread support throughout the computing, networking and embedded industry positions HyperTransport technology as a leading contender to be the universal onboard interconnect technology.




Figure 2 – HyperTransport technology is a universal technology that can be used to provide high bandwidth chip-to-chip interconnect for processor, memory, I/O as well as to provide bandwidth compatible links to high-speed external I/O technologies.

Performance advantage at lower cost in real-world applications – a key reason for popular adoption

Regardless of widespread industry support, it is the real-world benefits obtained when used in popular applications that will make a given technology a success in the marketplace. The high-technology mantra of “smaller, cheaper, faster” is a continuing market directive that can only be met if developers can achieve extreme performance gains while simultaneously driving down costs.

HyperTransport technology plays an important role because it brings definitive advantages to a number of key applications, including personal computers, servers, network equipment and embedded systems.

HyperTransport technology provides bus integration, higher bandwidth and lower cost for next generation personal computer

A key industry driver is the personal computer. For years the personal computer motherboard has been the spawning ground of processor, system, I/O and specialty bus structures. Legacy buses such as ISA, VL-Bus, AGP, LPC, PCI-

32/33 and PCI-X have been utilized to support processors, memories, graphics engines, and a wide assortment of I/O devices and subsystems. Now that the processor designers and developers have taken advantage of new silicon design and manufacturing technologies to deliver GHz plus clock rate processors and combined them with dual-data rate (DDR) memories, these legacy buses have become the choke point of the personal computer and workstation motherboard.

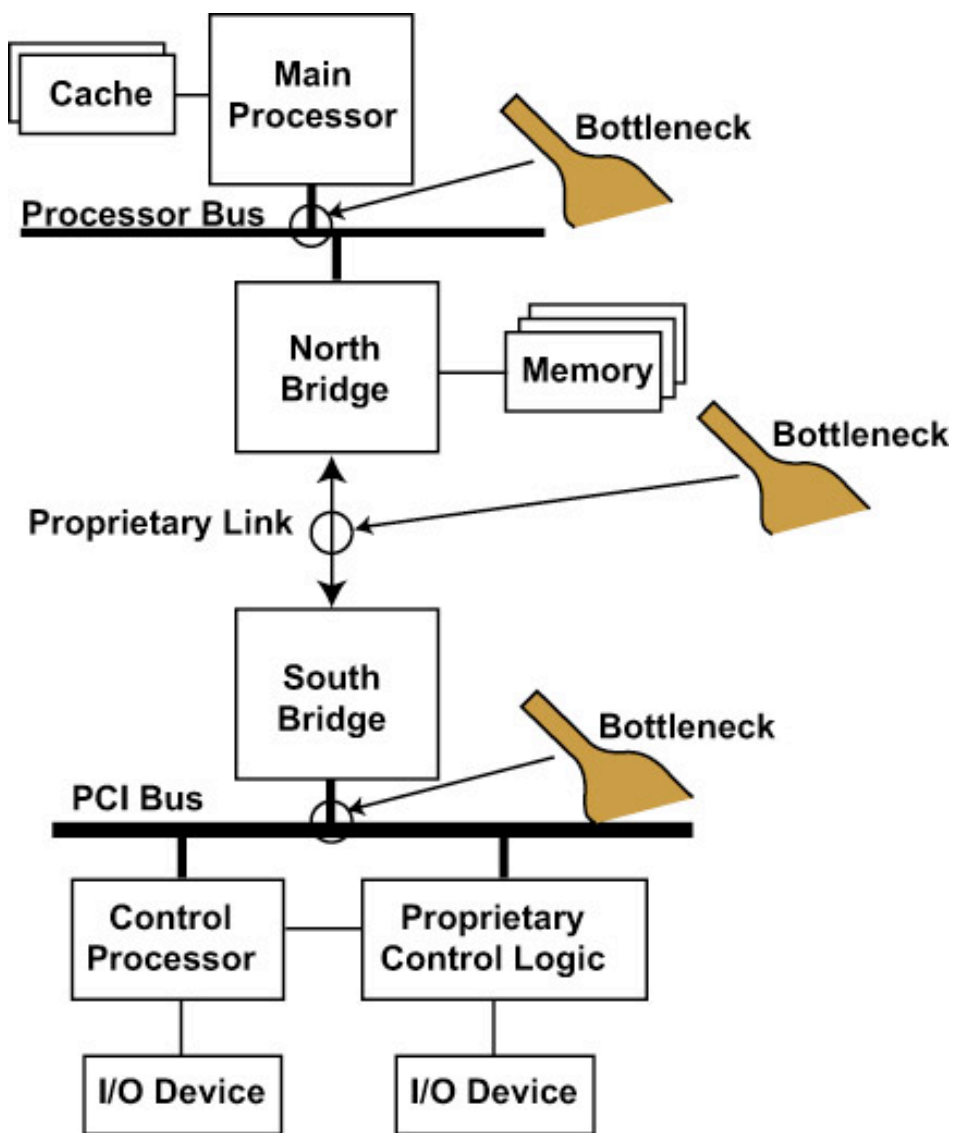


Figure 3 – The traditional personal computer system architecture has multiple layers of bus structure. Increasing processor clock speeds will not solve the I/O and memory bottlenecks.

There are several problems to be solved in the personal computer and workstation application space.

On the one hand, to cope with the proliferation of high-speed I/O data generated by MP3 audio (data capture, compression and recording to CD-Rs), 3D graphics, USB ports, Firewire ports, and 100 Mbps and Gigabit Ethernet links, more bandwidth must be available in the core system. With more high-speed data paths comes the requirement for more signal pins – with sufficient electrical headroom to avoid signal corruption. This usually means more signal pins, more power and ground connections and more external components like resistors and capacitors – all clogging the motherboard and complicating PCB layouts.

In addition, the existing PCI bus must be retained because there is too much industry expertise and product legacy issues to abandon it. However, its role as the central personal computer I/O bus technology must change to a supporting role as the bridge to legacy products.

Finally, as it becomes easier to produce high-speed processors, the next logical step is the integration of multiple processors into the personal computer. This raises the problem of integrating each CPU with each other and creating high speed data paths between local and shared memory and I/O.

In the next generation personal computer, the primary internal bus pathway will be built around the main processor complex. The processor complex includes one or more processors, coprocessors, local cache, local memory, shared memory and I/O pathway buses to slower I/O and links to the next generation higher speed I/O technologies. This is the heart of the next generation personal computer and workstation platform.

These bus requirements can quickly overwhelm the traditional bus approaches. They require many parallel signal pins, along with many associated clock, power and ground signals and significant numbers of passive components to ensure signal integrity at high speeds. HyperTransport greatly reduces the number of signals due to its use of enhanced LVDS signals, point-to-point implementation, and its in-band commands (packetized data and commands) that eliminate sideband signals.

As a result, HyperTransport provides an excellent chip-to-chip solution for the next generation personal computer. HyperTransport provides faster bandwidth, narrower signal paths, lower-cost implementations and integration with both legacy I/O and compatibility with high-speed communications technology.

With even low-cost 8-bit HyperTransport I/O links, aggregate bandwidth can easily accommodate multiple HyperTransport devices and HyperTransport to PCI bridges. If additional bandwidth is required, the links may be boosted to 16-bit wide. Since HyperTransport technology supports asymmetrical links, HyperTransport-to-PCI bridges may have 16-bit links for one channel and 8-bit links for slower speed ports.

Since the HyperTransport protocol encompasses the PCI enumeration and configuration protocols, existing operating systems need no modifications to take advantage of the greater bandwidth and integration made possible by HyperTransport technology. This is of supreme importance in the personal computer marketplace as deploying new technologies usually force system manufacturers to delay market introduction until software drivers can be rewritten, tested and made available to the consumer. HyperTransport requires no such effort making its deployment a seamless transition to a far greater performance internal system technology.

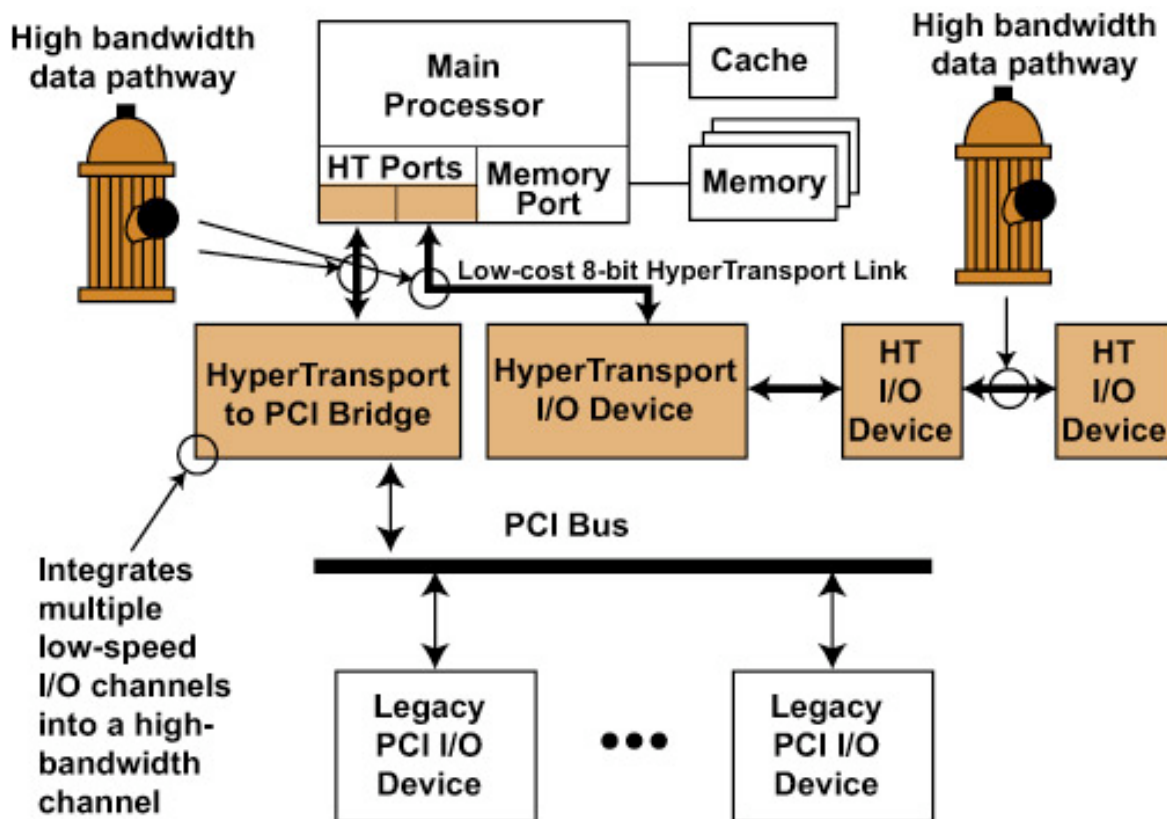


Figure 4 – The next generation HyperTransport-based personal computer system architecture provides bandwidth and integration for increased processor clock speeds, faster I/O technologies and legacy I/O.

HyperTransport is ideal for multiple processor systems

The HyperTransport advantage becomes even clearer when additional processors and their local memories are added to the system. In a traditional bus system, when a second (or more) processor is added, the additional processing bandwidth is negated by the bottleneck caused by multiple processors trying to access the same shared memory through slow traditional buses. With HyperTransport-enabled processors, this bottleneck is eliminated while providing a high-speed, low-cost pathway between multiple processors and shared memories. HyperTransport also has provisions for cache coherency for multiprocessor systems. This provides a great deal of performance while avoiding traditional multi-access issues.

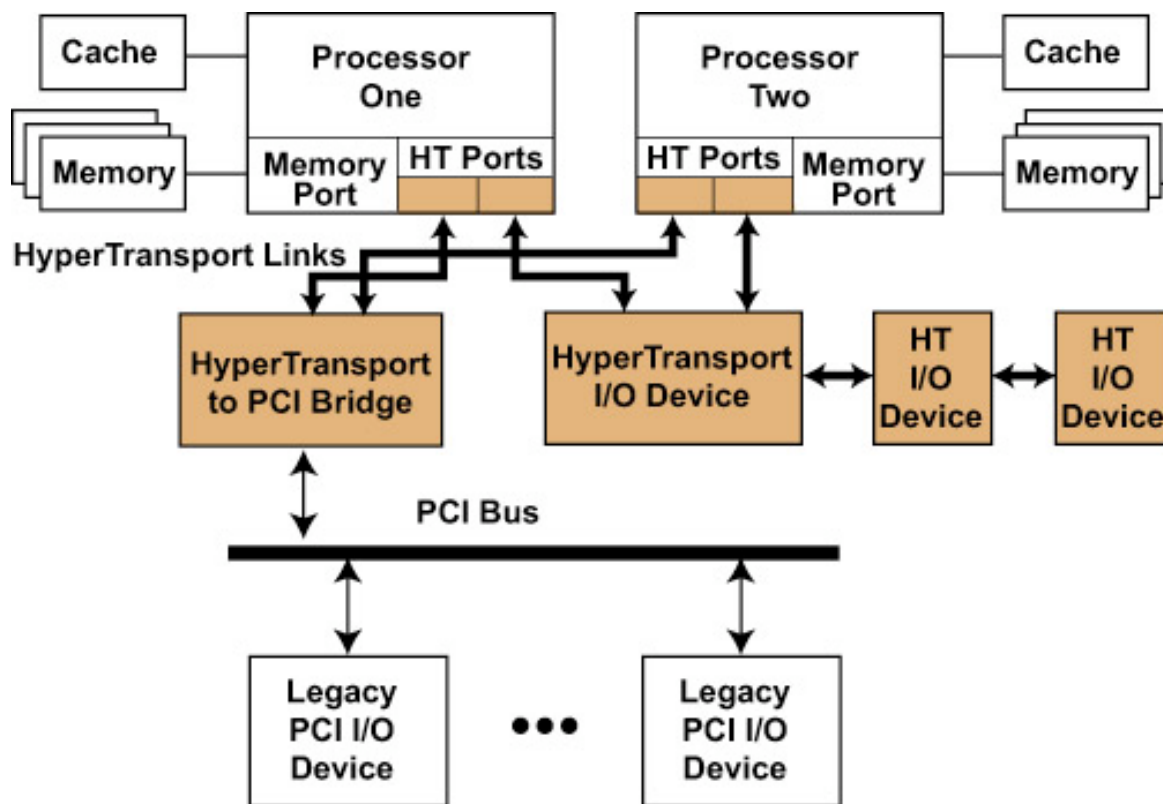


Figure 5 – With HyperTransport technology, multiprocessor based systems gain high bandwidth links between processors and memory and a low-cost multifunction link for processor-to-processor communications as well. HyperTransport enables better integration, faster data movement, lower cost and easier upgrade paths to higher performance systems.

With HyperTransport, as additional processor complexes are added to the system, the aggregate bandwidth of the system grows as well. This enables near linear increases in system performance as more processing, memory and I/O resources are added.

Servers Benefit from HyperTransport High Performance and Scalability

Servers also gain significant benefit from HyperTransport technology. While not exempt from the “smaller, faster, cheaper” mantra, servers gain even more from higher bandwidth capabilities and scalability.

At the server level system, multiple processors are becoming more and more the norm. Server farms can be consolidated into single systems with multiple

processor clusters. As overall system performance needs increase, it is simpler and most cost-effective to add additional processors rather than adding completely new systems. But multiprocessing systems have their own added complexities that must be overcome.

As shown in the multiprocessor personal computer example, servers can also benefit from the power of the HyperTransport technology model.

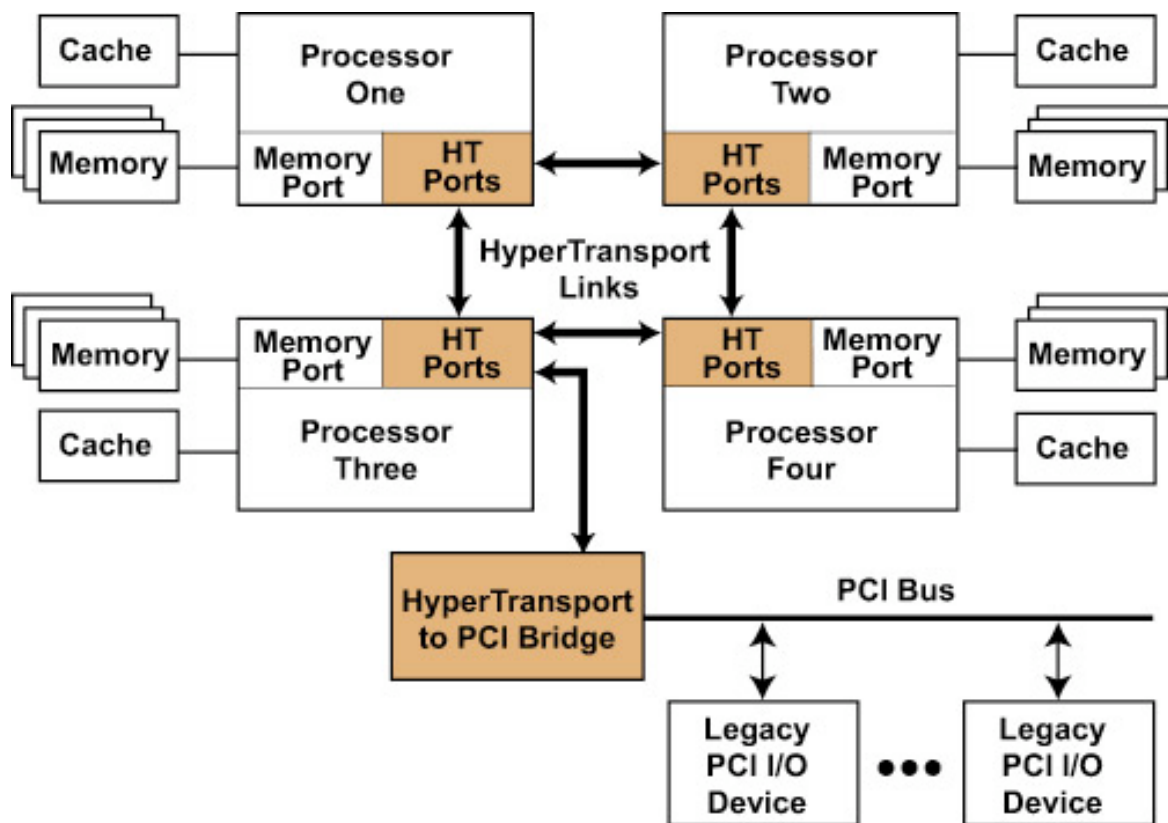


Figure 6 – HyperTransport technology enables large-scale multiprocessor-based servers with easily scalable performance based on adding multiprocessor subsystems. Scalable HyperTransport links can be asymmetric, with wider links between subsystems and narrow links used within a processor complex.

By using multiple processors with HyperTransport links, server providers can design very high performance systems that are easily scaled upwards in performance. A processor cluster of four processors, each with its local caches and memory can utilize a HyperTransport switch to connect to each other and shared memory. Each processor has the full bandwidth available for accessing

its own memory and can use the high-speed HyperTransport link to access the memory of the other processors.

As shown below, such systems are easily scaled to support many processor clusters. Each cluster can use HyperTransport links to connect to other processor clusters or to high-speed I/O channels or legacy PCI-type I/O devices.

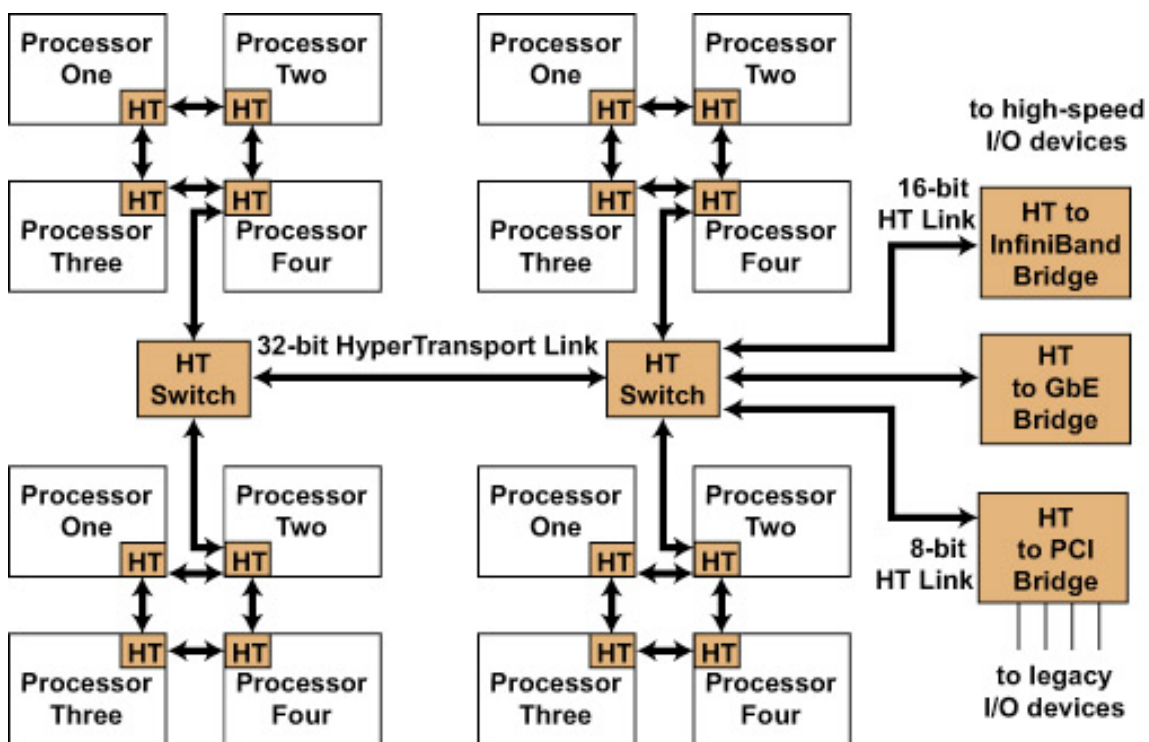


Figure 7 – Massive HyperTransport-based systems can be easily configured with multiple multiprocessor clusters linked together. Where additional bandwidth is required, 32-bit HyperTransport links can be used with 8- and 16-bit links used where the bandwidth needs are less.

For advanced server systems, access to high-speed communications technologies such as 10 Gigabit Ethernet, InfiniBand, and SPI-4 can be accommodated using HyperTransport bridges to those technologies. HyperTransport technology’s 12.8 Gigabit/second aggregate bandwidth enables it to easily integrate with these high-speed communications protocols.

HyperTransport Technology Provides Bandwidth and Integration for Network Equipment

Network equipment such as the routers that are the backbone of the Internet need high computational resources for such tasks as packet processing and forwarding, but they also need high-speed links to connect control processing elements with specialized data plane processing components. HyperTransport provides a universal interconnect that solves both problems. It is on the one hand a powerful interconnect technology for multiprocessing control plane processors (in similar fashion to the server application) and a fast, integrated communications link for connecting diverse specialized ASICs in the data plane.

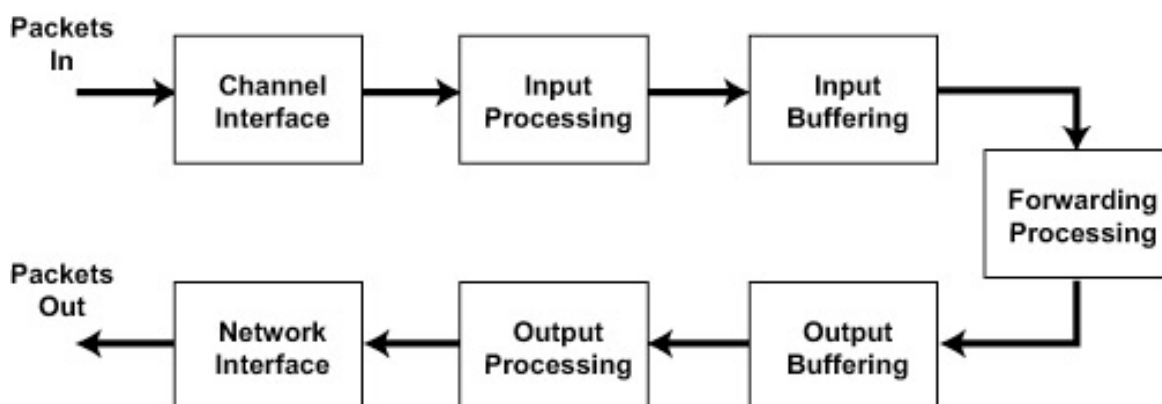


Figure 8 – Network routers essentially have two tasks. The first task is accepting input data streams from the network, stripping out the payload information and performing any processing that is required on the protocol envelop and data payload (data plane). The second task is determining the routing path for the data stream and controlling the packet processing units (control plane).

Traditional routers utilize general-purpose control processors, usually 32- or 64-bit RISC processors, to manage system resources and the overall data flow through the system, and specialized ASICs to perform data packet processing tasks. Large memory buffers are required to hold the packet data along each step of the processing and routing tasks.

Routers need large-scale computational resources to handle the control functions of packet processing and forwarding. They also need “fat pipes” to move massive amounts of data through the data plane. As the personal computer market has a

need for maintaining legacy software and PCI-type I/O investments, network equipment has a similar need for maintaining the millions of lines of code written for the general purpose 32- and 64-bit RISC processors that drive most of the network equipment installed today.

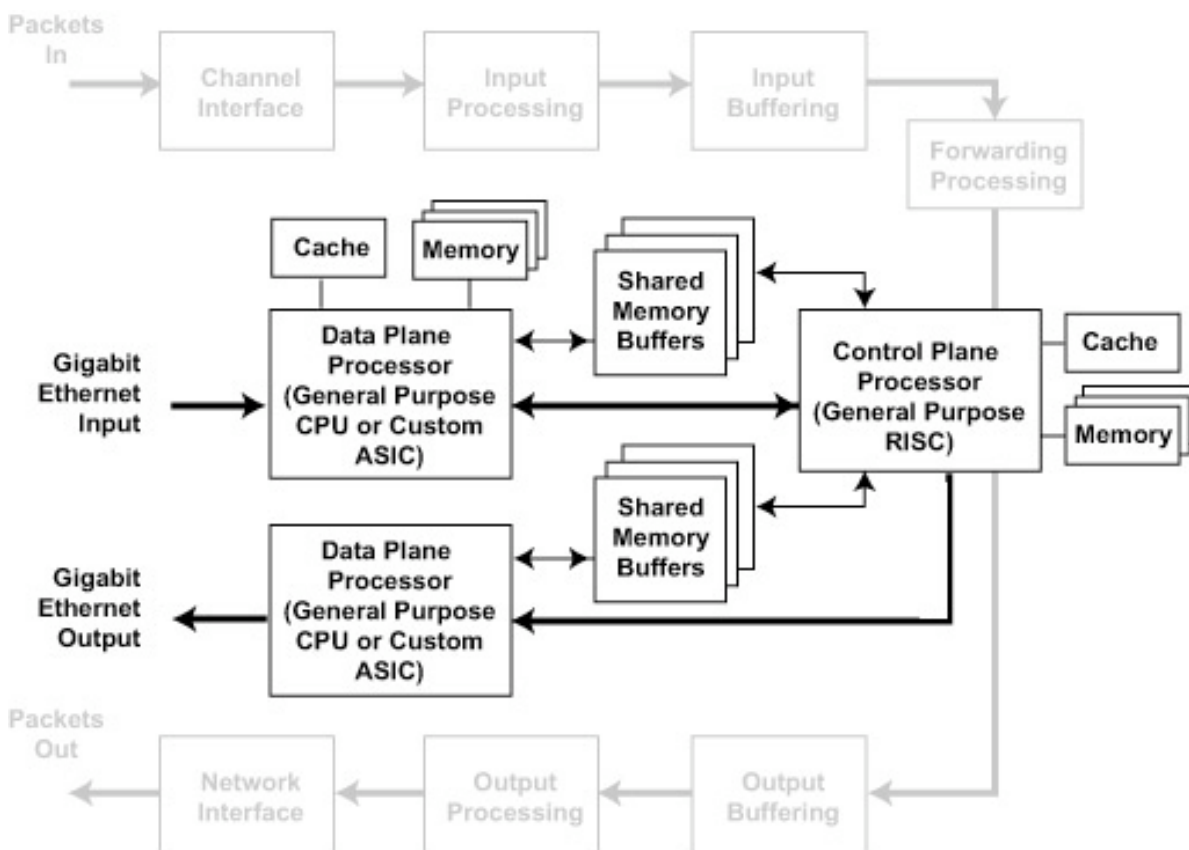


Figure 9 – Bottlenecks in control and data plane applications occur as large amounts of packet data must be shifted from data processing units through the system. Memory buffering and processor to local memory buses quickly run into bandwidth limitations.

While general-purpose processors are good in that they can handle many different tasks and have a well-known architecture, they aren't usually designed to be easily scalable. Gluing multiple RISC processors together runs into the same problems found in scaling up servers and personal computers. Additional processing power is useless without a corresponding increase in the bandwidth between system elements and the elimination of I/O and memory bottlenecks.

Using HyperTransport solves the problem by providing an integrated pathway between processor complexes. Even multiple 64-bit RISC processors can be linked using HyperTransport links with bandwidth to spare.

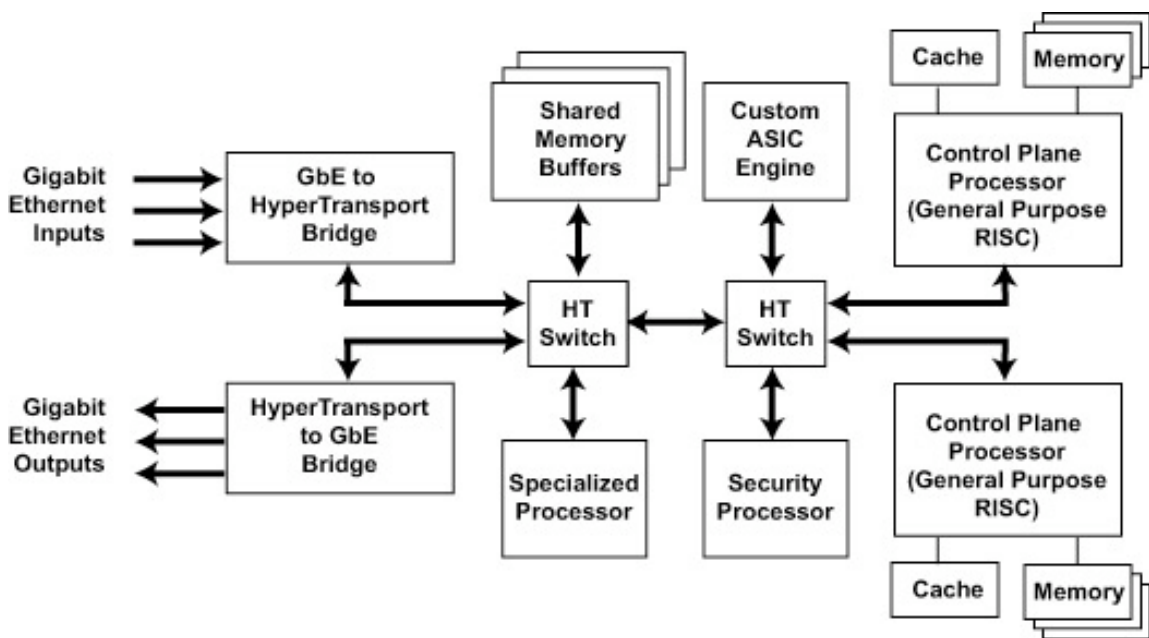


Figure 10 – HyperTransport-enabled multiprocessors eliminate bottlenecks in control plane applications in network equipment such as high-end routers. Core 32- and 64-bit RISC processors can use HyperTransport links to connect to each other and to move large data streams through the data plane from high-speed 10Gb Ethernet and SPI-4 communications links into the processor complex.

Using HyperTransport as the primary I/O channel through the router simplifies the design of high performance systems. Multiple HyperTransport switches connect processor complexes, shared memory buffers and network data pipes. Bridges to high-speed communications technologies such as Gigabit Ethernet, 10 GbE, SPI-4 and InfiniBand simplify network integration with emerging high-speed technologies. Because of the software transparent nature of HyperTransport, there is little software development work required in the control plane processor, again simplifying the system upgrade process.

HyperTransport provides a single internal high bandwidth data path with enough flexibility to support both legacy and emerging communications and I/O protocols while reducing costs and simplifying the integration task of adding multiple general purpose and custom ASIC processing units.

Embedded applications benefit from HyperTransport technology integration and low-cost implementations

It is not just high-end massively compute intensive systems that benefit from HyperTransport technology. Embedded systems of all types can benefit from the system integration functionality that comes with HyperTransport technology and the lower-cost of implementing narrow width chip-to-chip interconnects.

For example, printing and imaging systems typically utilize 32-bit RISC processors, local memory and PCI-compatible I/O subsystems. If they are network-ready, they also utilize Ethernet network interface cards. Using HyperTransport, next generation systems can be implemented that provide greater performance, cost less to manufacture and fit in more compact form factors.

Summary – HyperTransport technology delivers significant advantages to a wide range of key applications making it an ideal universal chip-to-chip interconnect technology

Unlike other protocols that attempt to solve all chip-to-chip, board-to-board, and box-to-box I/O issues with a single, but many-faceted specification, HyperTransport technology is tightly focused on becoming the universal chip-to-chip interconnect standard.

HyperTransport has won widespread industry support because of its many benefits to existing and emerging applications. Its high bandwidth characteristics are balanced by a low cost of implementation, ease of implementation, legacy support for older I/O technologies, compatibility with new high-speed I/O technologies and its scalability.

There are over 50 members of the HyperTransport Consortium group that maintains and guides the royalty-free HyperTransport technology. Leading providers of processors, chipsets, BIOSs, graphics engines, PCI and PCI-X bridges, FPGAs, silicon IP, test equipment and development tools have already announced and many are shipping HyperTransport-enabled products. This broad

HyperTransport ecosystem is fueling the rapid development of new HyperTransport-enabled products that are in turn defining the next generation of products in the embedded space, the personal computer market, high performance servers and high bandwidth network equipment.

For more information on HyperTransport technology please visit the HyperTransport Consortium website at www.hypertransport.org where additional white papers, detailed specifications and information on becoming a member of the HyperTransport Consortium is available.

A low-cost membership in the Consortium enables member companies to have royalty-free access to HyperTransport IP and to participate in technical working groups that define specifications and guide new additions to extend the technology to support future industry protocols.

END OF WHITE PAPER