

HYPER-ACTIVITY IN INTERCONNECT

HyperTransport is a low latency, high-performance, fully scalable, packet-based interconnect technology. CAROLINE HAYES looks at how a new version extends the interconnect technology

The era in which we live has been dubbed the 'click and get' economy, as the computer is a consumer's access to a world of consumables. While end-users enjoy the increase in computing power, it creates headaches for systems designers, combining multiple computing platforms, combating increased latency values and trying to effectively use the available bandwidth.

Add to this, the rise in processor performance with the adoption across the industry of multi-core technology and the urgency in providing interconnect solutions that increase the speed and bandwidth of signal transmission across devices and systems becomes apparent.

Standard types

PCI Express, for example, is best-suited to generic, peripheral interconnects in chip-to-chip applications. RapidIO is a multi-master mode technology, where every peripheral can act as a master. This capability lends the technology to embedded and networking applications.

HyperTransport is another interconnect technology, and is software compatible with legacy PCI, PCI-X and PCI Express

technologies. It is natively embedded in multiple CPUs available from AMD, Broadcom, IBM, NetLogic Microsystems, PMC-Sierra, Raza Microelectronics and Transmeta. It is also embedded in semiconductors and IP cores. It is created and defined by the HyperTransport Consortium, (www.hypertransport.org) presided over by AMD's David Rich, with PMC Sierra's Brian Holden as vice president and Mario Cavalli, a Consortium executive since 2004, acting as general manager.

Today there are over 40million HyperTransport technology-enabled systems. It can be found in PCs, workstations, servers, supercomputers and clusters from Apple, Cisco, Cray, Hewlett Packard, Fujitsu-Siemens, IBM, Microsoft, Sharp, Sun and those based on AMD Sempron, Athlon and Opteron, Turion 64 Mobile Technology and Transmeta Efficeon processors.

Specification 3.0

To cope with the next generation of processors used to build systems, the HyperTransport Consortium has finalised a third version of the standard specification. It nearly doubles the bandwidth and speed of HyperTransport 2.0. Version 3.0 increases clock support to over 40Gbyte, but without pin changes to the previous versions. It extends the DDR (dual data rate) maximum clock rate to 1.8, 2.0, 2.4 and 2.6GHz to deliver a maximum aggregate bandwidth of 41.6Gbyte/sec. This is an 86 per cent increase over the previous specification.

There is also a new, optional, AC coupling mode, through AC coupling capacitors. This is an alternative to enhancing the DC on the motherboard in order to support the increased clock speeds of these processors. The mode introduces AC/DC auto-sensing, which is transparent to system software, and auto-configuration to

extend maximum signal transmission or performance degradation. This technique adds packet overhead but without latency issues; it is still predictable and consistent, according to Cavalli. HyperTransport can now extend maximum signal transmission distance to 1m at the maximum specified clock speed without signal transmission or performance degradation.

Backplanes and cables can accommodate multi-boards of up to 32 interconnects to move the standard from chip-to-chip and board-to-board into chassis-to-chassis applications. Though similar to Ethernet, these systems are not intended to replace the network as it uses a cable to connect systems, which are not coupled side-by-side.

Feature enhancements

This specification of the standard can deliver intelligence to end products to optimise the power consumption of each HyperTransport component and module. This dynamic, auto-sensing, hardware-based, self-configuring power management is in real-time and operates transparently to the operating system and the application software.

Functionality has been added for link detection and link negotiation to make the standard truly hot-pluggable.

HyperTransport-enabled devices can be added or removed from the HyperTransport fabric without disrupting system operations.

The HyperTransport Consortium has also released the HyperTransport HTX connector specification. Using this specification, high-performance peripheral sub-systems can be linked directly to the system's CPU or multiple CPUs where applicable, using low-latency HyperTransport links.

The connector can be used to enable peripheral functions to perform co-processing functions such as server clustering, network security, real-time data analysis and routing storage management, math algorithms acceleration, encryption/decryption, 3D rendering and live video processing and medical imaging. **EPD**

High Performance Interconnect Evolution

