White Paper

Intel® QuickPath Architecture

A new system architecture for unleashing the performance of future generations of Intel® multi-core microprocessors

Introduction

Through its rapid "tick-tock" cadence for microprocessor innovation, Intel introduces a new microarchitecture or a new process technology (that includes enhanced microarchitecture features) nearly every year. The result is a cycle of industry-leading microprocessor performance. At the same time, the company also continuously looks for other ways to boost system infrastructure to support increased microprocessor performance. The latest example is Intel's use of scalable shared memory, a memory access technique that has long been used in the high-end expandable server world to improve performance.

Starting in 2008 with its next generation microarchitectures—code named Nehalem and Tukwila—Intel is incorporating a scalable shared memory (also known as non-uniform shared access or NUMA). Intel's new system architecture and platform technology will be called Intel® QuickPath Technology. It features new system architecture that: 1) integrates a memory controller into each microprocessor, and 2) connects processors and other components with a new high-speed interconnect. Previously announced under the code name Common System Interface or CSI, the Intel® QuickPath Interconnect is designed to unleash the performance of Intel's future generations of multi-core processors, delivering the additional performance, bandwidth, and reliability required to support Intel's next generation of dynamically scalable processors and platforms.

The Importance of the System Interconnect

While a microprocessor and its microarchitecture provide the processing power for a computer or server, the system interconnect plays a vital role in connecting the microprocessor to the rest of the system components. These components include system memory, various I/O devices, and other microprocessors.

Since more powerful microprocessors execute instructions faster, a potential bottleneck can form any time a processor or its individual cores can't fetch instructions and data as fast as they're being executed. Whenever the processor or an execution core must wait for an instruction or data to be received, performance slows. Of particular importance to the performance of a system is the speed at which a microprocessor and its execution cores can access system memory (in addition to internal cache). In a multi-processor system not only is the actual access to data important, but also the multi-processor communication required to ensure memory coherency (also called snoop traffic).

For years Intel met this demand by using an external bi-directional data bus called a front-side bus (FSB). Most FSBs perform as a backbone between the processor cores and a chipset that contains the memory controller hub and serves as the connection point for all other buses (PCI, AGP, etc.) in the system. This enables data flow between connected devices.

Intel has continually increased FSB performance by increasing the frequency (speed) of the bus and adding multiple FSBs per system. To take pressure off the bus, Intel also used its advanced process technology to include large memory caches in its processors. This allowed more instructions and data to be brought into the processor itself. Intel's latest 45nm processors, for example, include up to 50-percent larger L2 caches than previous generation (65nm) Intel® processors, along with a higher degree of associativity, to further improve the hit rate, maximize utilization, and free up the FSB.

Their performance says it all. Intel's 45nm processors advance Intel's industry-leading performance yet again. Intel's second-generation quad-core processors, the Quad-Core Intel® Xeon® processor 5400 series built with 45nm enhanced Intel® Core™ microarchitecture, boost performance by up to 25 percent in existing platforms using the same technologies, software, and socket compatibility.[1] These 45nm quad-core processors deliver up to 2x better performance than the previous generation dual-core and up to 5x better than the prior single-core Intel Xeon processor series.[2]

That's Today, Let's Look at What's Next

Building on this success, Intel is looking ahead and seeking new ways to enhance performance. In its long-range planning, the company has long anticipated the development of a next generation platform architecture that includes a new system architecture integrating the memory controller into the microprocessor and using a high-speed packetized interconnect as the processor bus. The first processors employing this technology will be Intel's next generation microarchitectures (Nehalem and Tukwila).

Along with new microarchitectures, Intel will introduce its Intel® QuickPath Architecture featuring Intel's first system interconnect transition since the introduction of FSB. The market can look forward to:

- Best-in-class interconnect bandwidth that's designed to unleash the full performance of next generation Intel microarchitecture (Nehalem) and future generations of Intel multi-core, dynamically scalable microprocessors.
- Best-in-class memory performance and flexibility to support leading memory technologies.
- Tightly integrated interconnect reliability, availability, and serviceability (RAS) and design-scalable configurations for high reliability and optimal price/performance/energy efficiency.
- Strong industry support—numerous industry leaders have adopted Intel QuickPath Technology and are developing innovative products for it.

What Is Intel® QuickPath Architecture?

Intel QuickPath Architecture is a platform architecture that provides high-speed connections between microprocessors and external memory, and between microprocessors and the I/O hub. One of its biggest changes will be the implementation of scalable shared memory. Instead of using a single shared pool of memory connected to all the processors through FSBs and memory controller hubs, each processor will have its own dedicated memory that it accesses directly through an Integrated Memory Controller. In cases where a processor needs to access the dedicated memory of another processor, it can do so through a high-speed Intel QuickPath Interconnect that links all the processors.

A big advantage of the Intel QuickPath Interconnect is that it is point-to-point. There is no single bus that all the processors must use and contend with each other to reach memory and I/O. It also improves scalability, eliminating the competition between processors for bus bandwidth. Coupled with Intel's great cache memory, this technological achievement will enable the performance of servers and workstations to take another leap forward.

Intel QuickPath Architecture is not Intel's first implementation of scalable shared memory—a computer memory design with a single physical address space, but in which various parts of memory are faster to access than other parts. Intel has used it in other platforms, such as Intel® 8870 chipset-based servers. Nor is it the first time Intel has used an integrated memory controller. Next generation microarchitecture-based platforms will simply be the first to bring both scalable shared memory and integrated memory controllers together. With each processor having its own memory controller and dedicated memory, the local memory will always be the fastest to access. In other cases, when an instruction or data is located in another microprocessor's dedicated memory, the memory will take longer to access. But not much—Intel QuickPath Interconnect is extremely fast.

Thanks to the pioneering work done by companies such as Sequent* (IBM*) and others, scalable shared memory has been used in the high-end server space for years and most modern operating systems (OSes) are optimized for it. This means they schedule processes and allocate memory to take advantage of local physical memory and improve execution performance. Most virtualization software is also written to take advantage of scalable shared memory, pinning a virtual machine to a specific execution microprocessor and its dedicated memory.

Intel® QuickPath Interconnect Advantages

Intel QuickPath Interconnect will deliver leading microprocessor interconnect bandwidth and RAS for Intel's next generation of server and workstation platforms. Key advantages include:

- Best interconnect performance in the mainstream server/workstation segment.
  - Intel QuickPath Interconnect uses up to 6.4 Gigatranfers/second links, delivering bandwidth up to 25 Gigabytes/second (GB/s) of total bandwidth—up to 300 percent greater than other interconnect solutions used today.[3] (Gigatransfer refers to the number of data transfers.)

- Efficient architecture improves interconnect performance.
  - Intel QuickPath Interconnect reduces the amount of communication required in the interface of multi-processor systems to deliver faster payloads.
  - The dense packet and lane structure allow more data transfers in less time, improving overall system performance.

- Tightly integrated RAS features for high reliability.
  - Implicit Cyclic Redundancy Check (CRC) with link-level retry ensures data quality and performance by providing CRC without the performance penalty of additional cycles. The link level retry retransmits data to make certain the transmission is completed without loss of data integrity.
  - For advanced servers which require the highest level of RAS features, some processors include additional features including the following: self-healing links that avoid persistent errors by re-configuring themselves to use the good parts of the link; clock fail-over to automatically re-route clock function to a data lane in the event of clock-pin failure; and hot-plug capability to enable hot-plugging of nodes, such as processor cards.

- Major vendors have already adopted Intel QuickPath Technology.
  - Vendors are already developing advanced chipsets for highly scalable servers.
  - Tool vendors are enabling robust development of silicon and boards that use the technology.
  - Other innovative products in the works include hardware accelerators for compute-intensive modeling in financial markets, oil exploration, and other fields.

Integrated Memory Controller Advantages

The Integrated Memory Controller is specially designed for servers and high-end clients to take full advantage of the Intel QuickPath Architecture with its scalable shared memory architecture. The independent high-bandwidth, low-latency memory controllers are paired with the high-bandwidth, low-latency Intel QuickPath Interconnects enabling fast, efficient access to remote memory controllers. The Integrated Memory Controller has the significant advantage of being coupled with large high-performance caches. This relieves pressure on the memory subsystem and lowers overall latency. The Integrated Memory Controller also continues Intel's legacy of best-in-class scalability and RAS features, plus of course takes advantage of next generation Intel® microarchitecture (Nehalem and Tukwila) and Hi-k 45nm process technology.

Key advantages include:

- Integrating the memory controller into the silicon die improves memory access latency (reduces communication delays in transferring data to and from memory) compared to traditional memory access through dedicated bus interface.
- Available memory bandwidth scales with the number of processors added.
- Full support for leading memory technologies allowing solutions optimized for different market segments.

Appearing Soon in a Server or Workstation near You

Intel QuickPath Architecture is a key platform ingredient in new upcoming Intel Xeon and Intel® Itanium® processor-based server and workstation platforms. This will mark the first time that these two server platforms share the same architecture and interconnect. This will simplify system design for the original equipment manufacturers (OEMs).

Server and workstation customers can look forward to another jump in performance as Intel introduces both a new microarchitecture and a new platform architecture with a faster system interconnect. Future desktop and mobile products will also include Intel QuickPath Architecture.

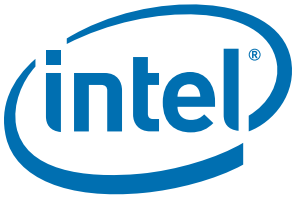

www.intel.com

Learn More

For more on Intel QuickPath Technology, see : **www.intel.com/technology/quickpath**

1. Quad-Core Intel® Xeon® processor x5460 series delivers up to 25% (1.25x) higher performance when compared to Quad-Core Intel® Xeon® processor 5365 series as published/measured using SPECjbb2005* in October 2007. For complete performance information, see: www.intel.com/products/processor/xeon5000/index.htm
2. Quad-Core Intel® Xeon® processor x5460 series delivers up to 119% (2.19x) higher performance when compared to Dual-Core Intel® Xeon® processor 5160 series as published/measured using SPECjbb2005* in November 12th, 2007. Quad-Core Intel Xeon processor x5460 series delivers up to 443% (5.43x) when compared to single-core 64-bit Intel® Xeon® processor 3.80 GHz as published/measured using SPECint*_rate_base2006 in November 12th, 2007.
3. Source: Intel estimates based on internal measurements March 2008.

   For complete performance information, see: www.intel.com/products/processor/xeon5000/index.htm

Copyright © 2008 Intel Corporation. All rights reserved. Intel, the Intel logo, the Intel. Leap ahead. logo, Intel Xeon, Itanium, and Intel Core are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

Printed in the United States. 0408/VP/HBD/PDF 319725-001US