Intel® Itanium® Quad-Core Architecture for the Enterprise

Lambert Schaelicke & Eric DeLano
(lambert.schaelicke@intel.com, eric.delano@intel.com)
Intel Corporation

1. Introduction

Intel recently launched the Intel Itanium processor 9300 series, the latest generation processor in the Itanium family of EPIC architectures. This processor, code-named Tukwila, combines four execution cores and their respective cache hierarchies with an integrated crossbar router, two coherency controllers and two memory controllers. While the microarchitecture of the execution cores remains largely unchanged from the previous generation 9100 series, the integrated system interface is a marked departure from previous bus-based system interconnects. Tukwila is the first Itanium processor to implement Intel's Quick-Path Interconnect (QPI) as well as the point-to-point Intel Scalable Memory Interconnect (SMI). The six Intel QPI links provided by Tukwila allow vendors to interconnect multiple Tukwila sockets with each other and with the necessary I/O hubs in a variety of topologies. Tukwila implements a directory-based coherency protocol on Intel QPI, leading to significantly improved scalability. The two Intel SMI-based memory controllers use an external memory buffer (Intel 7500 Scalable Memory Buffer) to drive up to two pairs of DDR3 RDIMMs each. These memory buffers convert the high-speed link-based Intel SMI protocol into the standard DDR3 protocol, but more importantly greatly increase the fan-out and memory capacity that Tukwila support without sacrificing link speed due to electrical loading or routing distances. In addition to these architectural changes, Tukwila also provides dynamic voltage and frequency scaling to improve power efficiency and peak performance, as well as an extensive set of RAS capabilities for new levels of resilience against soft and hard errors.

The Intel Itanium processor 9300 series is the result of the hard work and dedication of a large number of teams and individuals at Intel, including design and validation teams, firmware and BIOS development teams, the Itanium performance team, and many other partner groups within Intel.

2. Performance Features

The four processing cores are to a large degree leveraged from the previous-generation 9100 series processor, but include a number of significant changes. Most importantly, Tukwila incorporates twice as many cores as its predecessor. The First-level translation look-aside buffer (TLB) now supports 16KB pages, bringing some of the benefit of large pages from the second-level TLB closer to the core pipeline and reducing stalls due to TLB misses. The switch-on-event multithreading algorithm has been enhanced with a switch on stall event. This new event initiates a thread switch after the main pipeline has been stalled for certain number of cycles, usually due to a secondary miss in the shadow of a prefetch instruction. As a result, the distance between prefetch instructions and the associated demand loads has become somewhat less critical when balancing single-threaded and multi-threaded application performance. Other important changes include reducing the sizes of the second-level instruction and the third-level unified cache to allow for doubling the number of cores. Reducing cache capacity can in some cases increase the bandwidth demand of a workload, but the dramatic bandwidth improvement provided by the integrated Intel QPI and Intel SMI system interconnects is in many cases able to compensate for this increased demand. Finally, core frequency has been increased to 1.86 GHz in boost mode and 1.73 GHz in base frequency mode.

The on-chip router connects the four cores with the two integrated directory and memory controllers, and also provides four full-width and two half-width links for

external connections. These external Intel QPI links facilitate the construction of four and eight-socket systems with up to two I/O hubs without the need for any additional routers or node controllers. Larger systems can be built using additional node controllers. On Tukwila, Intel QPI employs a directory-based coherency protocol, where the on-chip directory controllers track cache line owners and state. Compared to previous bus-based system interconnects, this Intel QPI implementation greatly improves system scalability by reducing the number of snoop requests sent to caching agents, and by avoiding the bottleneck of a centralized shared bus.

The on-chip memory controllers connect to external memory buffers via serial high-speed Intel SMI links. Each of the four memory buffers hosts up to four DDR3 RDIMMs, thus providing a significantly higher memory capacity compared to a direct-attached memory configuration. As a result of the integrated memory controllers, total memory bandwidth scales with the number of sockets in the system, while maintaining a constant latency to the memory local to the respective cores. Internal measurements show that one-hop remote memory latency is only 40% higher than local memory latency, reducing the need for extensive NUMA optimizations in many applications.

3. Power Management

Tukwila implements a new, fully deterministic frequency and power management scheme based on measured activity in various parts of the chip. Depending on the total power demand, the cores may run at frequencies of 1.86 or 1.73 GHz, thus allowing workloads with a lower core power demand to execute at a higher frequency.

4. Reliability, Availability and Serviceability

Tukwila improves the RAS capabilities of previous Itanium by extending the Intel Cache Safe technology from the third-level to the second level data caches. This technology

tracks transient errors of individual cache lines and maps out lines that show too many errors to protect against persistent errors. Memory protection supports double device data correction (DDDC) and patrol scrubbing, as well as memory mirroring and failover for the highest levels of reliability.

5. Summary

The Intel Itanium 9300 series processor provides significant advances over previous-generation Itanium processors. It doubles the number of processing cores compared to the 9100 series, implements numerous microarchitectural enhancements in the cores and integrates a scalable directory-based system interconnect and two memory controllers on the same die. The processor has been launched in February of 2010 and will be available in systems from a variety of vendors.