Second Generation Quad-Core Intel Xeon Processors Bring 45 nm Technology and a New Level of Performance to HPC Applications

Paweł Gepner, David L. Fraser, and Michał F. Kowalik

Intel Corporation

{pawel.gepner,david.l.fraser,michal.f.kowalik}@intel.com

Abstract. The second generation of Quad-Core Intel® Xeon® processors was launched on November 12th 2007. In this paper we take a look at what the new 45 nm based Quad-Core Intel Xeon Processor brings to high performance computing. We compare an Intel Xeon 5300 series based system with a server utilizing his successor the Intel Xeon 5400. We measure both CPU generations operating in dual socket platforms in typical HPC benchmark scenario using some common HPC benchmarks. The results presented clearly show that the new Intel Xeon processor 5400 family provides significant performance advantage on typical HPC workloads and would therefore be seen to be an appropriate choice for many of HPC installations.

Keywords: HPC, multi-core processors, quad-core processors, parallel processing, benchmarks.

1 Introduction

Today multi-core processors are becoming a standard for high performance computing. The second generation Quad-Core Intel Xeon processor not only represents a technology shrink to 45 nm process technology but also brings lot of new mechanisms which improve overall performance and power savings characteristics. Intel Xeon 5400 is based on the same micro-architecture as the Intel® Core™ Microarchitecture including some extensions which actually raise the performance.

This paper has been written by Intel employees, therefore competitive products were not taken into consideration. Intel Xeon processor family contains 3000, 5000 and 7000 series where each of them is dedicated to different platforms and applications:

- Intel Xeon 3000 family is optimized for single-socket solutions;
- Intel Xeon 5000 family is optimized for dual-socket solutions;
- Intel Xeon 7000 family is optimized for multi-socket system (4 + way).

All of them are based on the same microarchitecture principles and they are all used in HPC instalations, where Intel Xeon 5000 family is the most common therefore authors have focused on them.

Our strategy of building CPUs which do not drive performance via faster clock speeds but rather based on an energy-efficient architecture has changed the landscape of high-performance computing (HPC) completely. If we look on the 30th edition of the TOP500 list released (Nov. 12, 2007) at SC07, the international conference on high performance computing, networking, storage and analysis, in Reno (NV, US). We see 317 systems based on the Intel® Core™ Microarchitecture and 102 of them are systems based on Quad-Core Intel® Xeon® processor.

True performance is a combination of both clock frequency and Instruction Per Clock (IPC). This shows that the performance can be improved by increasing frequency and/or IPC. Frequency is a function of both the manufacturing process and the micro-architecture. Basically there are two micro-architecture approaches which somehow determinate CPU design: more IPC or higher frequency. The first approach uses very few transistors, but the path from start to finish is very long, the second is based on shorter path, but it uses many more transistors [1, 2].

From manufacturing process perspective a key consideration is reducing the size of the transistors which means reducing the distances between the transistors and reducing transistor switching times. These two things contribute significantly to faster processor clock frequencies. Unfortunately as processor frequencies rise, the total heat produced by the processor scales with it. Reducing the transistor size, because smaller transistors can operate on lower voltages, allows the chip to produce less heat this does not however solve the all of the issues and in fact generates another problem with electrons. Based on quantum mechanics principles small elements such as electrons are able to spontaneously flow, over short distances. The emitter and transistor base are now so close together that a considerable number of electrons can escape from one to the other, this effect is called leakage. Reducing operating voltages which effectively reduces the available voltage swing, that is the difference between logic 1 and logic 0 becomes too small so that the transistor will not operate properly. In addition we need to deal with the decreased transistor size, as the leakage current increases we require more complicated process technology. In conclusion, this complicated situation where the number of transistors per unit area needs to increase, but the operating frequency must go down, will likely increase the number of cores required and decrease (not increase so quickly) the clock frequency [3, 4].

Multi-core processor systems changed the dynamics of the market and enabled new innovative designs delivering high performance with optimized power characteristics [5]. They drive multithreading and parallelism at a higher than instruction level, and provide it to mainstream computing on a massive scale.

2 Processor Microarchitecture

The Architecture typically refers to the high level description of the Instruction Set Architecture (ISA). The Architecture generally defines an instruction set for software to write to and, in turn, the software could then be run on all processor implementations of that particular architecture.

The Microarchitecture defines a specific means of implementing compatible hardware that supports the higher level architecture. New micro-architectures typically define improvements that ultimately increase the user benefits when running software that is compatible with the high-level architecture. Microarchitecture is enhanced with each processor generation, delivering improvements in performance, energy efficiency, and capabilities while still maintaining application-level compatibility. Microarchitecture refers to the implementation of the ISA in silicon, including cache memory design, execution units, and pipelining. In fact, benefits from many microarchitecture enhancements can be achieved without any modification or recompilation of code.

The 45 nm next generation Intel® Quad Core Xeon® processor family (Harpertown) is the next instance of Intel processors based on Intel 45 nm transistor technology, a new transistor breakthrough that allows for processors with nearly twice the transistor density and drastically reduced electrical leakage. New Intel Quad Core Xeon includes new instructions and microarchitecture enhancements that will deliver superior performance and energy-efficiency while maintaining compatibility to already existing applications.

Microarchitecture enhancements in Intel Quad Core Xeon 5400 processor family include:

- New set of instructions – Intel SSE4
- 50% larger L2 Cache
- Super Shuffle Engine and Fast Radix-16 Divider
- Enhanced Cache Line Split Load
- Deep Power Down Technology
- Enhanced Intel Dynamic Acceleration Technology

Intel SSE4 is a set of new instructions designed to improve the performance and energy efficiency of a broad range of applications. Intel SSE4 builds upon the Intel 64 Instruction Set Architecture (ISA), the most popular and broadly used computer architecture for developing 32-bit and 64-bit applications. Intel SSE4 consists of 54 instructions divided into two major categories: Vectorizing Compiler/Media Accelerators and Efficient Accelerated String/Text Processing. The new Intel Quad Core Xeon currently supports 47 of the Intel SSE4 instructions including the Vectorizing Compiler and Media Accelerator instructions. The remaining instructions will be available in future generations of Intel processors. Software will be able to use Vectorizing Compiler and Media Accelerators to provide high performance compiler primitives, such as packed (using multiple operands at the same time) integer and floating point operations, that allow for performance optimized code generation. It also includes highly optimized media-related operations such as sum absolute difference, floating point dot products, and memory loads. The Vectorizing Compiler and Media Accelerator instructions should improve the performance of audio, video, and image editing applications, video encoders, 3-D applications, and games.

50% larger L2 Cache (up to 12 MB in Quad Core implementation): Reduces the latencies for accessing instructions and data, improving application

performance (especially those that work on large data sets). 24 Way Set Associativity improves data access versus previous generation Intel® Xeon® (16 Way Set Associativity). Improved Store Forwarding maximizes data balance cache to memory.

Super Shuffle Engine and Fast Radix-16 Divider: 3X faster shuffles and 1.6X-2X faster divides. The Super Shuffle Engine will greatly improve the performance of Intel SSE4 and Supplemental Streaming SIMD Extensions 3 (SSSE3) instructions. New super shuffle engine performs 128 bit operation in a single cycle and does not require any software changes.

Enhanced Cache Line Split Load: Greatly improved performance on unaligned loads (those that span across cache boundaries) and optimized store and load operations. New 16 byte aligned load instruction on WC (write combining) memory improves read bandwidth from WC memory by reading cache-line size quantities. This Streaming Load routine gives 8X faster reading from WC Memory and improves the performance of memory-intensive applications.

Deep Power Down Technology: A new power state that dramatically reduces processor power consumption. This is ideal solution for developing energy efficient applications.

Enhanced Intel Dynamic Acceleration Technology: Improves energy efficiency by dynamically increasing the performance of active cores when not all cores are utilized. Conceptually it uses the power headroom of the idle cores to boost the performance of the non-idle core. When one core enters an idle power C-state (CC3 or deeper) and the OS requests a higher performance state on the running core, the non-idle core is boosted up to a higher voltage and higher frequency (EDAT freq) however the overall chip power envelope still remains within the specified Thermal Design Power (TDP).

How all of these innovations and changes in microarchitecture reflect to the overall system performance and accelerating high performance computing is described below. In the testing environment we have been testing single system performance in a typical HPC workload.

3 Processor Performance

In this section we have focused on processor performance to compare two generations of the Quad-Core Intel Xeon processors. A popular benchmark well-suited for parallel, core-limited workloads is the Linpack HPL benchmark. Linpack is a floating-point benchmark that solves a dense system of linear equations in parallel. The metric produced is Giga-FLOPS or billions of floating point operations per second. Linpack performs operations called LU Factorization. They are highly parallel and store most of their working data set in processor cache [6]. The processor operations it does perform are predominantly 64-bit floating-point vector operations and uses SSE instructions. This benchmark is used to determine the world's fastest computers published at the website [7].

In both cases each processor core has 3 functional units that are capable of generating 128-bit results per clock. In this case we may assume that a single processor core does two 64 bit floating-point ADD instructions and two 64 bit floating-point MUL instructions per clock. Theoretical performance, calculated as the product of MUL and ADD executed in each clock, multiplied by frequency in both cases are the same. Both implementations are based on the same microarchitectures. For Quad-Core Intel® Xeon® processor X5365 this gives 3 GHz x 4 operations per clock x 4 cores = 48 GFLOPS. Exactly the same theoretical performance we have for new Quad-Core Intel Xeon processor E5472 = 3 GHz x 4 operations per clock x 4 cores = 48 GFLOPS. This is theoretical performance only, it is interesting to observe how a 50% bigger cache and 20% faster Front Side Bus (FSB) of the Intel Xeon processor E5472 benefit Linpack and see if this is the good benchmark for CPU performance.

In the all performance tests we were using systems configured as follows, with a one exception for the Stream benchmark.

Configuration details

Quad-Core Intel Xeon processor X5365 based platform details: Intel preproduction server platform with two Quad-Core Intel Xeon processors X5365 3.00 GHz, 2x4 MB L2 cache, 1333 MHz FSB, 16 GB memory (8x2 GB FBDIMM 667 MHz), RedHat Enterprise Linux Server 5 Kernel 2.6.18-8.el5 on x86-64, Intel C/Fortran Compiler. Workload: Intel Optimized SMP LINPACK Benchmark 10 for Linux / LMBENCH 3.0 / Amber, Eclipse, Fluent, Gamess, Gromacs, Gaussian, LS-DYNA, Monte Carlo, PamCrash, Star-CD.

Quad-Core Intel Xeon processor E5472 based platform details: Intel preproduction server platform with two Quad-Core Intel Xeon processors E5472 3.00 GHz, 2x6 MB L2 cache, 1600 MHz FSB, 16 GB memory (8x2 GB FBDIMM 800 MHz), RedHat Enterprise Linux Server 5 Kernel 2.6.18-8.el5 on x86-64, Intel C/Fortran Compiler. Workload: Intel Optimized SMP LINPACK Benchmark 10 for Linux / LMBENCH 3.0 / Amber, Eclipse, Fluent, Gamess, Gromacs, Gaussian, LS-DYNA, Monte Carlo, PamCrash, Star-CD.

Using LINPACK HPL we see (Fig. 1) 5% performance improvement between the system based on the Quad-Core Intel Xeon processor X5365 and Quad-Core Intel Xeon processor E5472. It indicates that as we have expected based on the theoretical performance we will not see a big performance improvement on CPU intensive tasks. New Quad-Core Intel Xeon processor E5472 processor and its bigger cache and 1600 FSB do not play an important role in Linpack scenarios as it is not a memory intensive benchmark.

4 Memory Performance

In this section we illustrate the memory performance of the two generations of Quad-Core Intel Xeon processors. Memory performance is combination of two elements: latency and throughput. Each is appropriate to a different work load, and each tells a different story. Latency measures how long it takes to chase a

Fig. 1. LINPACK: Dense Floating-Point Operations

chain of pointers through memory. Only a single chain is tracked at a time. Each chain stores only one pointer in a cache line and each cache line is randomly selected from a pool of memory. The pool of memory simulates the working environment of an application. When the memory pool is small enough to be placed inside cache, the benchmark measures the latency required to fetch data from cache. By changing the size of the memory pool we can measure the latency of any specific level of cache, or to main memory, by creating the pool bigger than all levels of cache.

We measured latency using a 3.0 GHz Quad-Core Intel® Xeon® processor X5365 and a 3.0 GHz Quad-Core Intel Xeon processor E5472. The results of this experiment are shown in Fig. 2. Based on the different FSB as well as the bigger L2 cache the Quad-Core Intel Xeon processor E5472 would have slightly lower latencies to caches but the memory latencies comparing to his predecessor are much more noticeable, and from the result we can see that this is the case. The bigger L2 cache and faster FSB 1600 MHz benefit L2 cache latency as well reduce the access time to the main memory. All those elements in conjunction with the new Intel® 5400 chipset enabling 1600 MHz FSB helps random memory access.

The second important characteristic of the memory performance is the throughput for sequential memory accesses. The benchmark we have used to measure the throughput is Stream benchmark. The Stream benchmark is a synthetic benchmark program, written in standard Fortran 77. It estimates, both memory reads and memory writes (in contrast to the standard usage for bcopy). It measures the performance of four long vector operations. These operations are:

COPY: $a(i) = b(i)$

SCALE: $a(i) = q * b(i)$

SUM: $a(i) = b(i) + c(i)$

TRIAD: $a(i) = b(i) + q * c(i)$

These operations are representative of long vector operations and the array sizes are defined in that way so that each array is larger than the cache of the processors that are going to be tested. This gives us an indication of

Fig. 2. Memory latency

how effective the memory subsystem is in our implementation excluding caches As Fig. 3 shows we see huge memory bandwidth improvement with new Quad-Core Intel® Xeon® processor E5472 mainly due to 20% faster FSB (1600 MHz) as well improved functionality of Intel 5400 chipsets (“Seaburg”) capable to operate at 1600 MHz FSB. This 32% throughput improvement versus older generation Quad Core Xeon based system will be reflected in all memory intensive applications not only for Stream but also for other HPC data intensive workloads.

Configuration details

Quad-Core Intel Xeon processor X5355 based platform details: Intel preproduction server platform with two Quad-Core Intel Xeon processors X5355 2.66 GHz, 2x4 MB L2 cache, 1333 MHz FSB, 16 GB memory (8x2 GB FBDIMM 667 MHz), RedHat Enterprise Linux Server 5 Kernel 2.6.18-8.el5 on x86-64, Intel C/Fortran Compiler. Workload: Stream Benchmark.

Quad-Core Intel Xeon processor E5472 based platform details: Intel preproduction server platform with two Quad-Core Intel Xeon processors E5472 3.00 GHz, 2x6 MB L2 cache, 1600 MHz FSB, 16 GB memory (8x2 GB FBDIMM 800 MHz), RedHat Enterprise Linux Server 5 Kernel 2.6.18-8.el5 on x86-64, Intel C/Fortran Compiler. Workload: Stream Benchmark.

5 Application Performance

Linpack and Stream are synthetic benchmarks and they measure the performance of specific subsystems and do not deliver the full picture of system capability. Typical HPC applications use much more than a single subsystem and their nature is much more sophisticated. So to get a better understanding how the new Quad-Core Intel® Xeon® processor E5472 based platform benefits real applications we have selected a couple of real code examples. These application

Fig. 3. Stream benchmark – memory throughput

and benchmarks represent a broad spectrum of HPC workloads and seem to be a typical representation of a testing suite for this class of calculation.

Amber is a package of molecular simulation programs. The workload measures the number of problems solved per day (PS) using eight standard molecular dynamic simulations. See [8] for more information.

Eclipse from Schlumberger – reservoir simulation software for structure, geology, fluids and development scheme.

Fluent is a commercial engineering application used to model computational fluid dynamics. The benchmark consists of 9 standard workloads organized into small, medium and large models. These comparisons use all but the largest of the models which do not fit into the 8 GB of memory available on the platforms. The Rating, the default Fluent metric, was used in calculating the ratio of the platforms by taking a geometric mean of the 8 workload ratings measured.

GAMESS from Iowa State University, Inc. – a quantum chemistry program widely used to calculate energies, geometries, frequencies and properties of molecular systems.

Gaussian from Gaussian, Inc. – a quantum chemistry program widely used to calculate energies, geometries, frequencies and properties of molecular systems.

GROMACS (Groningen Machine for Chemical Simulations) from Groningen University is molecular dynamics program widely used to calculate energies, geometries, frequencies and properties of molecular systems.

LS-DYNA is a commercial engineering application used in finite element analysis such as a car collision. The workload used in these comparisons is called 3 Vehicle Collision and is publicly available from [9]. The metric for the benchmark is elapsed time in seconds.

Monte Carlo from Oxford Center for Computational Finance (OCCF) – financial simulation engine using Monte Carlo technique [10].

Fig. 4. Platform comparison across HPC selected workloads

PamCrash from ESI Group – an explicit finite-element program well suited for crash simulations.

Star-CD is a suite of test cases, selected to demonstrate the versatility and robustness of STAR-CD in computational fluid dynamic solutions. The metric produced is elapsed seconds converted to jobs per day. For more information go to [11].

All these selected workloads have been tested on two dual socket HPC optimized platforms. The Quad-Core Intel® Xeon® processor X5365 based platform has been used as the baseline to illustrate the improvement (Fig. 4) that the new platform is going to bring in different workloads in typical HPC scenario. As we can see the Quad-Core Intel Xeon processor E5472 based platform shows significant performance improvement up to 37%. The 12 MB L2 cache and new 1600 FSB help especially in the data intensive application and the workloads where data movement plays an important role. If the task is more CPU intensive the difference is around 10-15%.

6 Conclusion

The new Quad-Core Intel Xeon processors bring 45 nm technology to HPC with all the benefits of superior thermal characteristic and substantial production capability. Following the path of its predecessor these new products are continuing to demonstrate performance leadership. From the theoretical performance point of view both generations of Quad-Core Intel Xeon families deliver the same theoretical performance peak but we have seen that in a real life scenario the new architecture extensions bring a lot of performance improvement, even typically CPU intensive tasks show a better performance in the range of 7-15%. In the future this can be extended when the new set of SSE4 instructions become even more widely used and start providing additional headroom for performance improvement. The Vectorizing Compiler instructions should improve the performance of all those HPC applications which use multiple operands at the same time. These are the areas where the new Quad-Core Intel® Xeon® processors

bring the biggest improvement – the data intensive workloads. The 50% bigger L2 cache as well as 1600 MHz FSB in conjunction with the Intel 5400 chipsets made the new Quad-Core Intel Xeon processor based platform up 37% more effective when comparing to the old one – whilst operating at the same processor frequency.

We see a significant performance advantage, ranging from 20-37% and in addition observe significant performance per watt reduction, as the platform stays in the same power envelope. All of this drives significant improvements in user experience for the HPC environment and become a compelling choice for many of the new HPC installations.

References

1. Gepner, P., Kowalik, M.F.: Multi-Core Processors: New Way to Achieve High System Performance. In: PARELEC 2006, pp. 9–13 (2006)
2. Smith, J.E., Sohi, G.S.: The Microarchitecture of superscalar processors. *Proc. IEEE* 83, 1609–1624 (1995)
3. Ronen, R., Mendelson, A., Lai, K., Lu, S.-L., Pollack, F., Shen, J.P.: Coming challenges in microarchitecture and architecture. *Proc. IEEE* 89, 325–340 (2001)
4. Moshovos, A., Sohi, G.S.: Microarchitectural innovations: Boosting microprocessor performance beyond semiconductor technology scaling. *Proc. IEEE* 89, 1560–1575 (2001)
5. Ramanathan, R.M.: Intel Multi-Core Processors: Leading the Next Digital Revolution, *Technology @ Intel Magazine* (2005),
<http://www.intel.com/technology/magazine/computing/multi-core-0905.pdf>
6. Dongarra, J., Luszczek, P., Petit, A.: Linpack Benchmark: Past, Present, and Future, <http://www.cs.utk.edu/~luszczek/pubs/hplpaper.pdf>
7. <http://www.top500.org/>
8. <http://amber.ch.ic.ac.uk/amber8.bench1.html>
9. <http://www.topcrunch.org/>
10. <http://www.occf.ox.ac.uk/>
11. <http://www.cd-adapco.com/products/STAR-CD/>