International Journal of Emerging Technology and Advanced Engineering
Website: www.ijetae.com (ISSN 2250-2459, ISO 9001:2008 Certified Journal, Volume 6, Issue 10, October 2016)

Considerations and Challenges before Moving Towards Multi-Core Processors

Raksha Pandey
Computer Science & Engineering Department, KNIT, Sultanpur, India

***Abstract*— The thrust to achieve higher performance without driving up power consumption and heat has become a concern for many IT organizations. Multi-core processor architectures are designed to boost performance and minimize heat by mounting two or more processor socket. This research paper introduces the multi-core concept and discusses the major factors that IT organization should consider to harness the real power of Multi-core technology.**

***Keywords*— Multi-core Technology, Multi-core challenges, High Performance Computing**

I. INTRODUCTION

According to Lizhe Wang et al. in [1], "A Multi-core processor is typically a single processor which contains several cores on a single chip". These multiple cores on a single chip combine to improve the performance of a single faster processor. The individual cores on a multi-Core processor don't necessarily run as fast as the highest performing single-Core processors, but they improve overall performance by handling more tasks in parallel [2].

The traditional way for improving CPU performance in single core architecture has lead dangerously close to the power consumption ceiling, which stops the performance progress in processor development. Ongoing progress in processor designs has enabled servers to continue delivering increased performance incurs a corresponding increase in processor power consumption and heat is a consequence of power use.

According to R.M. Ramanathan in [3] at Intel Corporation, presented a fundamental theorem of multi-core processors. According to this paper, Multi-core processors take advantage of a fundamental relationship between power and frequency. By incorporating multiple cores, each core is able to run at a lower frequency, dividing among them the power normally given to a single core. The result is a big performance increase over a single-core processor. The following illustration based on lab experiments with commonly used workloads— illustrates this key advantage.

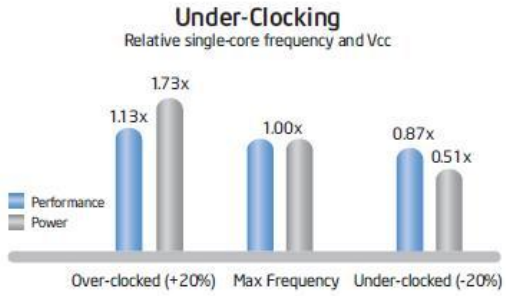

Figure.1 Performance relative to power consumption in single core

Figure.1 increasing clock frequency by 20 percent to a single core delivers a 13 percent performance gain, but requires 73 percent greater power. Conversely, decreasing clock frequency by 20 percent reduces power usage by 49 percent, but results in just a 13 percent performance loss.

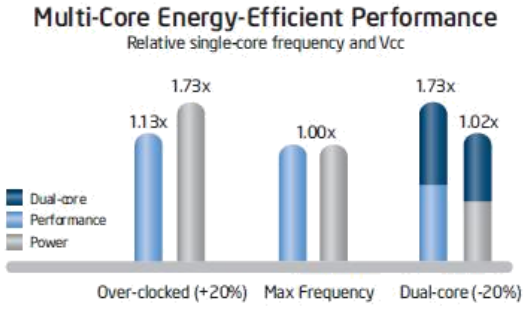

Figure.2 Performance relative to power consumption in multi-core

Figure.2: Here add a second core on the under-clocked example in Figure 1. This results in a dual-core processor that at 20 percent reduced clock frequency effectively delivers 73 percent more performance while using approximately the same power as a single-core processor at maximum frequency.

International Journal of Emerging Technology and Advanced Engineering
Website: www.ijetae.com (ISSN 2250-2459, ISO 9001:2008 Certified Journal, Volume 6, Issue 10, October 2016)

It was predicted that by placing an extra processing core in the same chip, there shall be enhancement in the performance, as well as lower production of heat, but the actual speed of the core was lower in comparison to the single core processor. The IEEE reviewed in September 2005 that the power consumption increases up to 60% with the use of every 400 MHz rise in clock speed, it also cited that we can get considerable improvement in performance through the means of dual-core approach [4].

Major chip manufacturers [5] have recently ramped up the development of massively multi-core processors for a variety of reasons including power consumption, memory speed mismatch, and instruction-level parallelism limits. For example, AMD has introduced a 12-core Opteron processor targeting the data center server market, while Intel has developed a 48-core single-chip computer for cloud computing. Projecting these trends into the future, chip manufacturers predict hundreds of processor cores per chip in the near future. These developments, however, demand a dramatic change in conventional programming paradigms and software models..

Kai Zheng et al. [6] Proposed coordination of multiple cores in a single processor to achieve better performance and lower real time power consumption for single multimedia streaming application, as opposed to multiple independent applications executing in a parallel fashion.

Dynamic power consumption of multi-core processors in [7] is estimated. Due to the variable behaviour of the different components of a processor, the modelling process decomposed into the following three component levels: i) processor's chip: these are components located out of dies (e.g. on-chip voltage regulator), ii) die: these are components within a die (e.g. off-chip cache) and iii) core: these are components within a core (e.g. control unit and on-chip cache). With the emergence of energy-saving mechanisms, the behaviour of power dissipation is different than that without such mechanisms.

According to Constantinos Christofi et al. In [8], Efficient and scalable Parallelism is the key to achieving high performance in Multi/Many-core systems. Sequential (control-flow) computing, also know and as von Neumann computing, has been the de-facto model of programming and computation since the advent of digital computers. The Power and Memory walls have forced the switch to multiple cores per chip thus elevating concurrency as the key challenge in achieving high performance.

New parallel programming and computation models are needed in order to fully exploit the ever increasing number of cores per chip. Their study investigated whether Data-Driven Multithreading (DDM) can achieve high performance in HPC applications.

DDM combines Data-Flow concurrency with the efficiency of sequential execution. DDM has been implemented and evaluated in both Heterogeneous (CELL) and Homogeneous Multi-core systems [9].

Reducing Peak Power Consumption in Multi-Core Systems without Violating Real Time Constraints in [10] is proposed by minimizing chip-level peak power consumption at design time without violating any timing requirements. With this goal achieved, one can optimize the chip design/packaging process and the power related components according to the minimized chip-level peak power consumption (guaranteed at design time), and then reduce the cost and/or size of the process and components by avoiding/minimizing their over-design.

II. Cooling And Power Management Of Multi-Core Processors

Significant power consumption and cooling requirements arise due to high server density in the data center. A multi-core architecture can help address the environmental challenges created by high-clock-speed, single-core processors. Heat is a function of several factors, two of which are processor density and clock speed. Other drivers include cache size and the size of the core itself. In traditional architectures, heat generated by each new generation of processors has increased at a greater rate than clock speed. In contrast, by using a shared cache rather than separate dedicated caches for each processor core, and low-clock-speed processors, multi-core processors may help minimize heat while maintaining high overall performance. This capability may help make future multi-core processors attractive for IT deployments in which density is a key factor, such as high-performance computing (HPC) clusters, Web farms, and large clustered applications. Environments in which 1U servers or blade servers are being deployed today could be enhanced by potential power savings and potential heat reductions from multi-core processors. Currently, technologies such as demand-based switching (DBS) are beginning to enter the mainstream, helping organizations reduce the utility power and cooling costs of computing. DBS allows a processor to reduce power consumption (by lowering frequency and voltage) during periods of low computing demand. In addition to potential performance advances, multi-core designs also hold great promise for reducing the power and cooling costs of computing, given DBS technology. DBS is available in single-core processors today, and its inclusion in multi-core processors may add capabilities for managing power consumption and, ultimately, heat output.

International Journal of Emerging Technology and Advanced Engineering
Website: www.ijetae.com (ISSN 2250-2459, ISO 9001:2008 Certified Journal, Volume 6, Issue 10, October 2016)

This potential utility cost savings could help accelerate the movement from proprietary platforms to energy-efficient industry-standard platforms.

III. Software Evolution Towards Multi-Core Technology

Multi-core processing continues to exert a significant impact on software evolution. Before the advent of multi-core processor technology, both SMP systems and HT Technology motivated many OS and application vendors to design software that could take advantage of multithreading capabilities. As multi-core processor based systems enter the mainstream and evolve, it is likely that OS and application vendors will optimize their offerings for multi-core architectures, resulting in potential performance increases over time through enhanced software efficiency. Most application vendors will likely continue to develop on industry standard processor platforms, considering the power, flexibility, and huge installed base of these systems. Currently, 64-bit Intel Xeon processors have the capability to run both 32-bit applications and 64-bit applications through the use of Intel Extended Memory 64 Technology (EM64T). The industry is gradually making the transition from a 32-bit standard to a 64-bit standard, and similarly, software can be expected to make the transition to take advantage of multi-core processors over time. Applications that are designed for a multiprocessor or multithreaded environment can currently take advantage of multi-core processor architectures. However, as software becomes optimized for multi-core processors, organizations can expect to see overall application performance enhancements deriving from software innovations that take advantage of multi-core-processor–based system architecture instead of increased clock speed. In addition, compilers and application development tools will likely become available to optimize software code for multi core processors, enabling long-term optimization and enhanced efficiency for multi-core processors which also may help realize performance improvements through highly tuned software design rather than a brute-force increase in clock speed. Intel is working toward introducing software tools and compilers to help optimize threading performance for both single-core and multi-core architectures. Organizations that begin to optimize their software today for multi-core system architecture may gain significant business advantages as these systems become mainstream over the next few years. For instance, today's dual Intel Xeon processor–based system with HT Technology can support four concurrent threads (two per processor).

With the advent of dual-core Intel Xeon processors with HT Technology, these four threads would double to eight. An OS would then have eight concurrent threads to distribute and manage workloads, leading to potential performance increases in processor utilization and processing efficiency.

IV. Shift In Focus Toward Multi-Core Technology

In order to exploit a multi-core processor at full capacity the applications run on the system must be multithreaded. There are relatively very few applications (and more importantly few programmers with the know-how) written with any level of parallelism. Multi-core processors represent the future of server architecture, which is expected to enhance application performance and platform power with thermal efficiency. By combining multiple logical processing units within a single processor package, multi-core processors have the potential to provide superior performance and scalability without power consumption and heat, as would be the case by simply increasing the clock speed of existing single-core processor designs. Until recently the focus was on processor count but now it has shifted to the number of processor cores per socket.

V. Challenges

Every new technology advancement brings new challenge. As discussed earlier there are a number of problems and challenges that should be considered while shifting towards energy-efficient industry-standard platforms (Multi-core platform). Temperature and Power management are two major concerns that my increase exponentially with multiple cores. Memory/cache coherence is also another issue that should be considered. And finally, using a multi-core processor to harness its full potential is another challenge. If programmers don't write such applications that may take advantage of multiple cores there is no gain. Applications need to be designed in such a way that different parts can be run concurrently and take full advantage of different cores.

Multi-core processors are designed to stick reasonable power consumption, heat dissipation. To fight against unnecessary power consumption many designs also incorporate a power control unit that has the authority to shut down unused cores.

Cache coherence is a concern in a multi-core environment because of distributed L1 and L2 cache. Since each core has its own cache, the copy of the data in that cache may not always be the most up-to-date version.

For example, imagine a dual-core processor where each core brought a block of memory into its private cache. One core writes a value to a specific location; when the second core attempts to read that value from its cache it will not have the updated copy unless its cache entry is invalidated and a cache miss occurs. This cache miss forces the second cores cache entry to be updated. If this coherence policy is not in place garbage value would be read and invalid results can be produced. Finally, the most important issue is using multithreading or other parallel processing techniques to get the real power of the multi-core processor. Rewriting applications to be multithreaded means a complete rework by programmers in most cases. Programmers have to write applications with threads that are able to be run in different cores, meaning that data dependencies will have to be resolved. If one or two cores are being used much more than another, the programmer is not taking full advantage of the multi-core system. Some companies like Microsoft and Apples have designed new products with multi-core capabilities whose newest operating systems can run on up to 4 cores.

VI. Conclusion

In this paper, the basic concept of multi-core processors is introduced. Adding multiple cores within a processor gave the technological advancements, but added interesting new problems. The Multi-core processors are developed to enhance performance of computing. The utilization of the processor can be 100% only when the applications being executed are designed so well.

However, before adopting this emerging multi-core architecture, IT organizations need to carefully evaluate the software ramifications of migrating applications to multi-core processor technology. This consideration can enable enterprises to benefit from the higher performance and lower power consumption as compared to single-core processor architecture.

Multi-core processors are the next innovation for computing architectures and platforms. Multi-core platforms will empower the development of new applications that will enable wide-ranging advances in everything from medicine to IT, from the digital office to the digital home, from mobility solutions to the latest games.

REFERENCES

[1] Wang Lizhe, et al., "Multicores in Cloud Computing: Research Challenges for Applications", Journal of Computers, Vol 5, No 6 (2010)

[2] Geer D., "Chip Makers Turn to Multi-core Processors," Computer, vol. 38, page(s): 11-13, 05, 2005.

[3] Ramanathan R.M., "Making the move to quad-core and beyond" Intel multi-core processors Technology,Intel Magazine, Dec 2006.

[4] Freuhe John, "Planning Considerations for Multi-core Processor Technology", Dell Power Solutions, May 2005.

[5] Lakshmanan Karthik, Kato Shinpei, Ragunathan (Raj) Rajkumar "Scheduling Parallel Real-Time Tasks on Multi-core Processors", Dec 2010, Page(s):259-268.

[6] Zheng Kai, Zhu Yongxin, Xu Jun "Evaluation of partitioning methods for stream applications on a heterogeneous multi-core processor simulator", Dec 2008, Page(s):486 – 491.

[7] Basmadjian Robert, Meer Hermann de "Evaluating and Modeling Power Consumption of Multi-Core Processors", May 2012, Page(s):1-10.

[8] Christofi Constantinos et al."Exploring HPC Parallelism with Data-Driven Multithreading", Sept 2012, Page(s):10-17.

[9] Arandi Samer and Evripidou Paraskevas, "the datadriven multithreading virtual machine for the cell processor", DDM-VMc: "In Proceedings of the 6th International Conference on High Performance and Embedded Architectures and Compilers (HiPEAC '11). ACM, New York, NY, USA, 2011, page(s):25-34.

[10] Lee Jinkyu et al.,"Reducing Peak Power Consumption in Multi-Core Systems without Violating Real-Time Constraints", IEEE transactions on Parallel and Distributed Systems, Vol. 25, No:4, APRIL 2014.