

Multi-Core Processors: A New Way Forward and Challenges

Abinash Roy, Jingye Xu and Masud H. Chowdhury

Department of Electrical and Computer Engineering, University of Illinois at Chicago
Chicago, IL 60607, USA

aroy5@uic.edu, jxu6@uic.edu, masudh@uic.edu

Abstract- Continuous effort to achieve higher performance without driving up the power consumption and thermal effects has led the researchers to look for alternative architectures for microprocessors. Like the parallel processing which is extensively used in today's all microprocessors, multi-core architecture which combines several independent microprocessor cores in a single die has currently become very popular in most high performance integrated circuits. Although multi-core processor offers excellent instruction execution speed with reduced power consumption, optimizing performance of individual processors and then incorporating them by interconnection on a single chip is a non-trivial task. This paper investigates the leading challenges associated with current high performance multi-core processor in terms of interfacing different cores, design automation and verification, software adaptability.

Index Terms— Power Consumption, Interconnect challenges, Design Automation, Software Adaptability, Multi-core Processor

I. INTRODUCTION

In order to track the continuing performance improvement following Moore's law, successive technologies have relied on scaling of different device and interconnect parameters. Historically, these performance gains have been accomplished by efficient exploitation of sophisticated process technology, innovative architecture or micro-architecture [2]. To keep on the circuit speed, focus is mainly given on increasing operating frequency. Device dimensions have been scaled to support higher integration density for greater functions. However, this scaling brings in several critical challenges in current sub-65nm technology. Although power supply has been scaled to keep the dynamic power at bay, due to the aggressive scaling of MOS geometry leakage power has been a major part of total power. Growing device components further aggravates the heat generation in a small chip. Therefore, thermal challenges have appeared as one of the major challenges to the successful advancement of CMOS technology [6].

A new approach called "parallel processing" was proposed in early 1990s to save power [9]. Since then, this method has gained wide acceptance among the architecture designers for which almost every processor now a days runs on this principle. However, this approach can't alone continue supporting growing speed of microprocessors. Clock frequency must be increased which in turn generates more heat in a processor. This is the main reason for all commercial processors to stick around 4GHz.

Considering the growing concern for power dissipation, the concept of multi-core processor is a new step forward and it has become the technology for current and next decades [11], [18]. A multi-core chip-level processor combines two or more independent cores into a single die. Thus a Dual-core processor contains two cores; a Quad-core processor contains four cores and so on. A multi-core processor implements multi-processing units on a single physical package. One basic difference between single processor and multi-core processor is that a single processor has a unique L1 cache along with a L2 cache where as each independent processor in a multi-core system has a common shared L2 cache in addition to an individual L1 cache. In a single core processor, 45nm technology is currently in production and next come 32nm, and 22nm and most likely 10nm node is the limiting technology node considering the strong quantum effects. Consequently, multi-core processor is a promising architecture technique.

IBM first introduced multi-core processor chip, Power4 in 2001 [15] through which designers were able to achieve much greater communication bandwidth and resulting performance. In mid-2006, Intel reached new levels of energy-efficient performance with their Intel Core™2 Duo processors using 65 nm technology and latest micro-architecture [2]. Although it has been a frequently used architecture, numerous challenges involve accordingly and they must be addressed by the researchers.

The rest of this paper is organized as follows. Section II briefly gives some major advantages of multi-core processor. It is concluded that multi-core processors become the standard for delivering greater performance, improved performance per watt and new capabilities across different electronic applications. Section III describes the leading interconnect challenges in multi-core

processors. Challenges incurred by design automation and verification and software adaptability have been briefly studied in section IV and V respectively. Finally section VI wraps up the paper.

II. PROMISING ASPECTS OF MULTI-CORE PROCESSORS

The key driving force to adopting multi-core processor architecture was to address power and cooling challenges. Figure 1 gives performance comparison between a single core and multi-core processor [2]. This analysis which is performed based on Intel tests using the SPECint2000 and SPECfp2000 benchmarks reports that multi-core processors perform much better than a single core processor and it is projected that relative advantage of multi-core system will enhance over the next couple of years.




Figure 1. Performance comparison between a single core and multi-core processor

A. Controlling power consumption by multi-core processor

Historically chip manufacturers have met the demand for increasing processor speed by boosting up the operating clock frequency along with the higher integration density. This approach has resulted in uncontrollable heat dissipation in current technology node. With heat rising incrementally faster than the rate at which clock signal propagates through the processors, it has prompted the processor designers for alternative methodologies. 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 power among them normally given to a single core. The result is a big performance increase over a single core processor. It can be observed that increasing clock frequency by 20% to a single core delivers a 13% performance gain, but requires 73% greater power. Conversely, decreasing clock frequency by 20% reduces power usage by 49%, but causes only 13% performance loss [2]. If a second core is added into the

single core architecture, it results in a dual-core processor that at 20% reduced clock frequency; it can effectively deliver 73% more performance while using approximately the same power as a single-core processor at maximum frequency.

B. Efficient usage of chip area by Caches and Memory modules

As stated earlier, each single processor core in a multi-core architecture has its unique L1 cache and all processors in the die share a common L2 cache. Therefore, number of caches and memories required become less than if single core processor is used for the equal number of jobs that need to be performed. For example, Intel Advanced Smart Cache works by sharing the L2 cache among cores so that data are stored in one place that each core can access. Sharing L2 cache enables each core to dynamically utilize even up to 100% of available L2 cache, thus optimizing cache resources [2].

Intel® Smart Memory Access improves system performance by optimizing available data bandwidth from the memory subsystem and hiding the latency of memory accesses through two techniques: a new adaptability called memory disambiguation, and an instruction pointer-based pre-fetcher that fetches memory contents before they are requested [2].

C. Performance enhancement by multi-threading technology

Along with parallel processing method, multi-threading technology is extensively used in single core processor. According to this approach, on a single processor, multithreading generally works on the principle of time-division multiplexing which is much similar to the parallel execution of multiple tasks where the processor switches between different threads [13]. This context switching happens so fast that it creates the illusion of simultaneity to an end user. On a multiprocessor system, threading can be achieved via multiprocessing, where as different threads and processes can run simultaneously on different processor cores. Threading a task in parallel processing machines thus not only increases the number of tasks executed per unit time but also enhances the accuracy of the task. Consequently, it is obvious that significant performance improvement can be achieved using multi-core systems coupled with advances in memory, I/O, and storage devices.

III. INTERCONNECT CHALLENGES IN MULTI-CORE PROCESSOR

Although device performance in a single processor has increased with the continuous scaling technology parameters over the generations, interconnect performance has degraded since interconnect scaling exhibit exactly

opposite trend. Therefore, the overall performance of a microprocessor is determined by interconnect characteristics [4].

The main bottleneck associated with the interconnection network of multi-core processor is the interfacing of different cores in a single die. Several interconnection mechanisms have been proposed in [10] to spice up the interconnect performance. Among them, a shared bus fabric (SBF) that provides connection to various modules with the capabilities of coherence source and sinking, a point to point link that connects two SBFs and a cross bar connection system are most commonly used. A shared bus fabric is a high speed link which can communicate data between processors, caches, I/O and memory in a multi-processor system. The effectiveness of such an approach depends on the probability that an L2 miss is serviced on a local cache (an L2 connected to the same SBF), rather than a cache on a remote SBF. It is not easy to maintain this condition in a complex multi-core system. Another problem is that the interconnect fabric itself is large and power-hungry, consuming resources that would otherwise be available for more cores and caches. Interconnect, even without the sharing of L2 caches, can occupy the area of three cores and the power of one. Each P2P link must be capable of transferring all kinds of transactions (request/response/data) in both directions. Each P2P link is terminated with multiple queues at each end. Therefore, this topology needs a queue and an arbiter for each kind of transaction.

In addition to this on-core interconnect challenges; interfacing several independent cores in a single core is also a non-trivial task. Different control logic circuits are required to properly route signals and data among the cores. Due to the huge latency of this routing path, in spite of excellent performance of all single cores, the overall performance of the multi-core architecture degrades.

Intra-chip interconnect coupling (capacitive and inductive coupling) is also significant in current single processors that form the multi-core system at the recent operating frequency of high performance processors [16], [17]. Although operating frequency in multi-core processor is being kept low (maximum that can be made without increasing power dissipation to threshold point), signal latency and signal distortion are still a major concern for the proper functioning of individual processors at this frequency range. In multi-core systems, if a specific processor shows higher latency to finish a task due to its internal signal degradation by noise coupling, then the overall system performance degrades because a complete instruction set is divided among several processors and therefore, its completion is dependent on all the processors involved. Each processor is not optimized equally [2], so due to different performance capabilities offered by processor cores, the system performance fluctuates. Considering inductive effect into the analysis of interconnect behavior is also

very important [16]. Higher inductive coupling leads to increasing signal propagation delay, crosstalk noise as well as delay and skew variations.

IV. DESIGN AUTOMATION CHALLENGES IN MULTI-CORE PROCESSOR

CAD tools are very important for the verification and optimization of high density integrated circuits from early phase of performance analysis to the final. The verification of a large block has been one of the major challenges in electronic design automation [2]. As the push for more functionality in a single core along with increasing power gating, variable frequencies and asynchronous communication, especially in ASIC design has grown over the years, the verification process has been more complex. Different functioning processors in a multi-core unit have different types of working principles. Therefore, their internal architecture is also different which causes onerous challenges to the verification teams to individually check on each processor. It complicates and creates additional delay to the chip production cycle [7].

Each core can be functioning at a different mixture of instructions leading to wide variations of dynamic power dissipation. Therefore, accurate analysis by a CAD tool becomes very hard to predict this thermal behavior.

In the early phases of processor design, chip layout, power analysis and packaging are divided among different teams. Considering all realistic assumptions about architecture, packaging and market requirements as well as alternatives, a performance model is formed. In physical architecture, each functioning cell is modeled as a block. A particular block generally acts as standard for reuse so many times in a single processor chip. Thus it saves design challenges and time. However, due to different performance requirements in different processor cores in a multi-core system, there is no standard method for minimizing layout optimization. A number of unique controlling blocks are required to properly design the whole multi-core system and it is not possible to implement all blocks using a common technology. This becomes more important when performance and power efficiency are key design criteria. In the final phase, a rigorous testing is performed on each processor to ensure the desired performance. There is no common test platform for multi-core processor systems.

V. CHALLENGES IN SOFTWARE ADAPTABILITY IN MULTI-CORE PROCESSOR

Proper synchronization of software with the hardware is very important for the accurate and smooth operation of the whole system. Multi-core processor continues to exert significant impact on software evolution. Before the advent of multi-core processor technology, software was

optimized for single core processor. However, developing software for a growing large multi-core processor is a major challenge.

Energy efficiency, throughput and multi-tasking capability of multi-core system will be efficiently realized when instruction codes are threaded with parallel processing approach.

As the multi-core processor based systems more evolve, the software vendors must optimize their software with the system requirement [15]. The main problem in this case is that most industries don't disclose the details of their hardware and collaborate only with some specified software vendors they do business with. Therefore, there are not so many versatilities of software for hardware. It causes big problem to the consumers. When it comes to multi-core processor system, this restriction of software usability gets worse since a specific multi-core system consists of different cores which are implementing different tasks. So supporting these numerous types of hardware modules with the unique software for each hardware appear as a big challenge for the software engineers.

VI. CONCLUSION

It is worth mentioning that introduction of multi-core processors has opened many new doors for high performance integrated circuits. While device scaling is approaching its limit, multi-core processors have continued improving the performance of CMOS technology by splitting the instruction executions among different functioning cores with the addition to parallel processing and multithreading technique. However, like any technology, with the advancement to the successive generations, numerous challenges such as interfacing different cores, innovative CAD tools and software adaptability for new architecture will be more complex and these issues must be addressed for ensuring the optimum performance of multi-core processors. Besides, single core chips will continue to compete since they have well established and inexpensive technology for manufacture and therefore, they will even be popular for low- priced PCs and servers.

REFERENCES

- [1] J. Cong, L. He, C-K. Koh, and P. Madden, "Performance optimization of VLSI interconnect," *The VLSI Journal Integration*, Vol. 21, pp. 1-94, November 1996.
- [2] www.intel.com
- [3] International Technology Roadmap for semiconductors, Semiconductor Research Corporation, 2008
- [4] Deepak C Sekar and James D Meindl, "The impact of multi-core architectures on design of chip-level interconnect networks," *International Interconnect Technology Conference*, pp. 123-125, June 2007
- [5] Isci, C.; Buyuktosunoglu, A.; Chen, C.-Y.; Bose, P.; Martonosi, M, "An Analysis of Efficient Multi-Core Global Power Management Policies: Maximizing Performance for a Given Power Budget" *Annual IEEE/ACM International Symposium on microelectronics*, pp. 347-358, December 2006
- [6] Rajarshi MukherjeeT and Seda Ogresci Memik "Physical Aware Frequency Selection for Dynamic Thermal Management in Multi-Core Systems" *International conference on Computer Aided Design*, pp. 347-352, November 2006
- [7] John A Darringer, "Multi-Core Design Automation Challenges," *Design automation conference*, pp. 760-764, June 2007
- [8] H. Peter Hofstee, "Future Microprocessors and Off-Chip SOP Interconnect," *IEEE Tran. On advanced packaging*, Vol. 27, No. 2, May 2004
- [9] A Chandrakasan, S Sheng and R Brodersen, "Low power CMOS digital design," *Journal of Solid state Cir*, April 1992
- [10] Rakesh Kumar, Victor Zyuban, Dean M. Tulsen, "Interconnections in multi-core architecture: understanding mechanisms, overheads and scaling," *International Symposium on Computer Architecture*, pp. 1-11, 2005
- [11] Pawel Gepner and Michal F. Kowalik, "Multi-core processors: New way to achieve high system performance " *International Conference on Parallel Computing in Electrical Engineering*, 2006
- [12] J. Parkhurst, et al, "From single core to multi-core: preparing for a new exponential" *International conference on Computer Aided Design*, November 2006
- [13] Chunqing Wu; Xiangquan Shi; Xuejun Yang; Jinshu Su, "Multi-threading in high performance processor," *International Conference Grid and Cooperative Computing*, pp. 236-240, October 2006
- [14] Donald, James; Martonosi, Margaret, "Power Efficiency for Variation-Tolerant Multicore Processors" *Low Power Electronics and Design*, pp. 304-309, October 2006
- [15] www.ibm.com
- [16] Abinash Roy, Noha Mahmoud and M. H. Chowdhury, "Effects of Coupling Capacitance and Inductance on Delay Uncertainty and Clock Skew", *Design Automation Conference*, pp. 184-187, June 2007
- [17] K. T. Tang and E. G. Friedman, "Interconnect Coupling Noise in CMOS VLSI Circuits," *Proceedings of the ACM/IEEE International Symposium on Physical Design*, pp. 48--53, April 1999
- [18] Shekhar Borkar, "Thousand Core Chips—A Technology Perspective" *Design Automation Conference*, pp. 746-749, June 2007