MULTIMEDIA APPLICATIONS ARE ESPECIALLY CONDUCIVE TO THREAD-LEVEL PARALLELISM

Dr S.Kishore Reddy*

Dr. Syed Musthak Ahmed**

Abstract

Multimedia applications are especially conducive to thread-level parallelism because many of their operations can run in parallel. Although faster processors are one way to improve server performance, other approaches can help boost performance without increasing 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. In a technical nutshell, multi-core processing can support several key capabilities that will enhance the user experience, including the number of PC tasks a user can do at one time, do multiple bandwidth-intensive activities and increase the number of users utilizing the same PC.

* Associate Professor, Department of ECE, SR Engineering College, Warangal, AP, India

** Professor &head, Department of ECE, SR Engineering College, Warangal, AP, India

Introduction

These higher quality signals allow more data to be sent in a given time period since individual signals can be shorter and do not need to be repeated as often. Assuming that the die can fit into the package, physically, the multi-core CPU designs require much less Printed Circuit Board (PCB) space than multi-chip SMP designs. Also, a dual-core processor uses slightly less power than two coupled single-core processors, principally because of the increased power required to drive signals external to the chip and because the smaller silicon process geometry allows the cores to operate at lower voltages; such reduction reduces latency. Furthermore, the cores share some circuitry, like the L2 cache and the interface to the front side bus (FSB). In terms of competing technologies for the available silicon die area, multi-core design can make use of proven CPU core library designs and produce a product with lower risk of design error than devising a new wider core design. Also, adding more cache suffers from diminishing returns. Let's start by looking at some of the multi-processor technologies which have contributed to AMD and Intel's newest products.

SMP (Symmetric Multi-Processing)**:** SMP is the most common approach to creating a multi-processor system, in which two or more separate processors work together on the same motherboard. The processors co-ordinate and share information through the system bus and the processors arbitrate the workload amongst themselves with the help of the motherboard chipset and the operating system. The OS treats both processors more or less equally, assigning work as needed. Both AMD and Intel's new dual-core chips can be considered SMP capable (internally). AMD's dual-core Opteron server processors can be linked to other dual-core chips externally also, but this capability is not present in either company's desktop dual-core lines. The major limitations of SMP have to do with software and operating system support. Many operating systems (such as Windows XP Home) are not SMP capable and will not make use of the second physical processor. Also, most modern programs are single-threaded, meaning that there is only ever one current set of linked instructions and data for them. This means that only one processor can effectively work on them at a time. Multi-threaded programs do exist, and can take better advantage of the potential power of dual- or multi-CPU configurations, but are not as common as we might like. No other current mainstream desktop processors are SMP capable, as Intel and AMD tend to restrict cutting edge technology to the higher-end server processors such as the Opteron and Xeon. In the past though, mainstream processors have been SMP capable, most notably the later Intel Pentium 3 processors.

Implementing Multiprocessors

Another feature that makes multiprocessing work better is referred to as bus arbitration. In order to access memory that's deeper than the processor's cache -- i.e., the RAM or the computer's hard drive -- the processor must go through something called the system bus. When more than one processor is operating, however, a squabble can arise over which processor gets to use the bus first. Bus arbitration assigns different voltages to pins on the processors, which in turn determine which processor is active and which is inactive. The active processor gets dibs on the bus. And this active/inactive status switches randomly between the processors, so nobody gets upset. At the end of the day, multiprocessing is a neat concept, but for at-home computer-users, it hardly pays off. Doubling your processor does not necessarily double your speed, unless you're constantly using applications that benefit from dual processors. For most at-home users, this is not the case. For them, given the cost of all the necessary upgrades, you're better off investing in a typing course to speed up your keyboard skills than shelling out for this kind of glamorous upgrade.

Highway to hyper threading:

Hyper threading was Intel's pre-emptive take on multi-core CPUs. The company cloned the front end of its high-end Pentium 4 CPUs, allowing the Pentium 4-HT to begin two operations at once. Once in process, the twin operation 'threads' both share the same set of execution resources, but one thread can take advantage of sections left idle by the other. The idea of Hyper threading is to double the amount of activity in the chip in order to reduce the problem of 'missed' memory cache requests slowing down the operation of the processor. It also theoretically ensures that less of the processor's resources will be left idle at any given time.

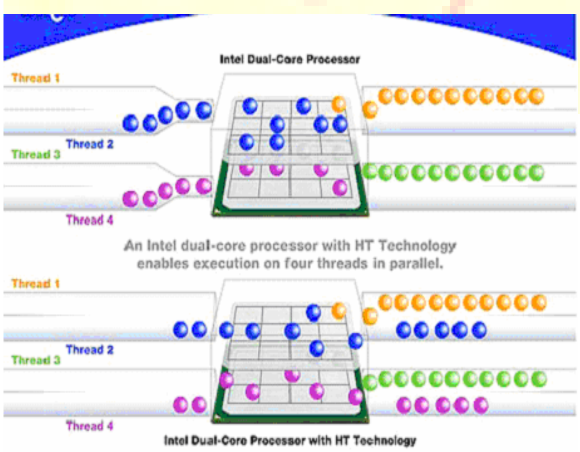

While Hyper threaded CPUs appear as two logical processors to most operating systems, they are not comparable with true dual-core CPUs since each parallel pair of threads being worked on share the same execution pipeline and same set of L1 and L2 cache memory. Essentially, Hyper threading is smoke-and-mirrors multitasking, since a single Hyper threaded processor cannot actually perform two identical actions at the same time.

Hyper threading does speed up certain operations which would be multi-processor capable, but never as much as a true multi-processor system, dual core or not. One of Intel's new dual-core chips, the higher-end Pentium Extreme Edition 840 processor, also support Hyper threading within each core, meaning that to an operating system it would appear as four logical processors on a single die. How this will work out remains to be seen.

Dual Core Processors

A dual-core processor is a **single chip** that contains two distinct processors or "execution cores" in the same integrated circuit.

Two Chips on One Die… Why?

Several reasons; first of all; competition, competition, competition!!! The race to superiority between AMD and INTEL (the precursors of Processor technology). AMD built the potential for dual-core capability into its 64-bit processors right from the start. The necessary I/O structure for the second core already exists, even on single core chips. Neither company can afford to let the other get much of an edge, and AMD has already stolen way too much attention for Intel's comfort with its incredibly successful line of 64-bit processors. It is imperative for Intel to launch a 'pre-emptive strike' and get its own dual-core technology to market quickly, lest market share flutter away. As for why dual core processors are being developed in the first place, read on to reason number three. Secondly, performance. Certain 'multi-threaded' applications can already benefit greatly by allowing more than one processor to work on them at once.

Quad Cores Are Better for Virtualization, as virtualization becomes more important and widely implemented businesses will need servers with quad core processors to help support more than one workload on a system and to supply the computing power to run multiple applications or operating systems on a single server. Within an enterprise, servers remain the ideal platform for quad-core processors because the chips are designed to take advantage of the multithreaded software that runs within most data centers. While quad-core processors provide a performance boost for digital content-creation and allow users to run certain tasks simultaneously, most desktop systems--outside of gaming PCs and highend models-- don't need quad-core chips. With more processing cores in each server, quad-core chips can help IT managers cut down on the number of systems each data centers needs. A quad-core Opteron processor will offer, on average, a 66 percent performance increase compared with a dual-core Opteron processor, according to benchmarking results released by AMD.

References

- Tensilica Inc., “Configurable processors: What, why, how?” Tensilica Xtensa LX2 White Papers, 2009[Online].Available:http://www.tensilica.com/products/literature-docs/white-papers/configurable-processors.htm
- L. Shimpi and D. Wilson, “NVIDIA’s 1.4 billion transistor GPU: GT200 arrives as the GeForce GTX 280 & 260,” Anandtech Web site, 2008 [Online].Available:http://www.anandtech.com/video/showdoc.aspx?i=3334
- M. Gschwind, H. P. Hofstee, B. Flachs, M. Hopkins, Y. Watanabe, and T. Yamazaki, “Synergistic processing in cell’s multicore architecture,” IEEE Micro, vol.
- 26, no. 2, pp. 10–24, 2006.
- J. Andrews and N. Baker, “Xbox 360 system architecture,” IEEE Micro, vol. 26,
- no. 2, pp. 25–37, 2006.
- Advanced Micro Devices Inc. “Key architectural features—AMD Phenom II processors,” AMD Product Information, 2008 [Online]. Available: http://www.amd.com/usen/Processors/ProductInformation/0,,30_118_15331_15917%5E15919,00.html
- Texas Instruments Inc., “OMAP 4: Mobile applications platform,” 2009 [Online]. Available: http://focus.ti.com/lit/ml/swpt034/swpt034.pdf