MULTI-CORE PROCESSORS

Abstract:

Multi-core processor architecture entails silicon design engineers placing two or more processor-based execution cores, or computational engines, within a single processor. A multi-core processor plugs directly into a single processor socket, but the operating system perceives each of its execution cores as a discrete logical processor, with all the associated execution resources. A processor equipped with thread-level parallelism can execute completely separate threads of code. This can mean one thread running from an application and a second thread running from an operating system, or parallel threads running from within a single application. 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 clock speed and incurring an increase in power consumption and heat. 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.

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.

Key Words: *Multi core processor, computational engines, multimedia applications, HT technology, nutshell, pc tasks.*

Conclusion: At present majority computers are in use with single processor, an attempt to multicore processor is in research. Indeed dual processors are presently used in some of the critical applications.

Core:

The processing part of a CPU chip minus the cache. It is made up of the control unit and the arithmetic logic unit (ALU).

Multi-Core:

"Multi Core" refers to - **two or more CPUs working** together **on one single chip (like AMD Athlon X2 or Intel Core Duo)** in contrast to DUAL CPU, which refers to two separate CPUs working together.

Advantages In Using Multi-Core Processors :

The proximity of multiple CPU cores on the same die allows the cache coherency circuitry to operate at a much higher clock rate than is possible if the signals have to travel off-chip. Combining equivalent CPUs on a single die significantly improves the performance of cache snoop (Bus snooping→**Bus sniffing** or **Bus snooping** is a technique used in distributed shared memory systems and multiprocessors aimed at achieving cache coherence. Every cache controller monitors the bus, awaiting for broadcasts which may cause it to invalidate its cache line.) operations. Put simply, this means that signals between different CPUs travel shorter distances, and therefore those signals degrade less. 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:

Implementing a multiprocessing system involves more than just buying a second processor. You'll need to outfit your SMP motherboard with a special chip that will manage the various functions of all processors simultaneously. The processors you use will also need to be SMP-friendly.

Some of the complex functions that multiprocessing necessitates are worth looking at. In order for two processors to not get ahead of each other and start generating data that renders either of them out-of-date, multiprocessing depends on something called cache coherency. As one processor digs into its cached memory to retrieve a piece of information, the other processor checks that it hasn't updated that particular item without the other processor's knowledge. If it has, then it supplies the latest data. This cache coherency prevents data from becoming corrupted as two or more processors work more or less separately on one task.

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 Hyperthreading:

Hyperthreading 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 Hyperthreading 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.

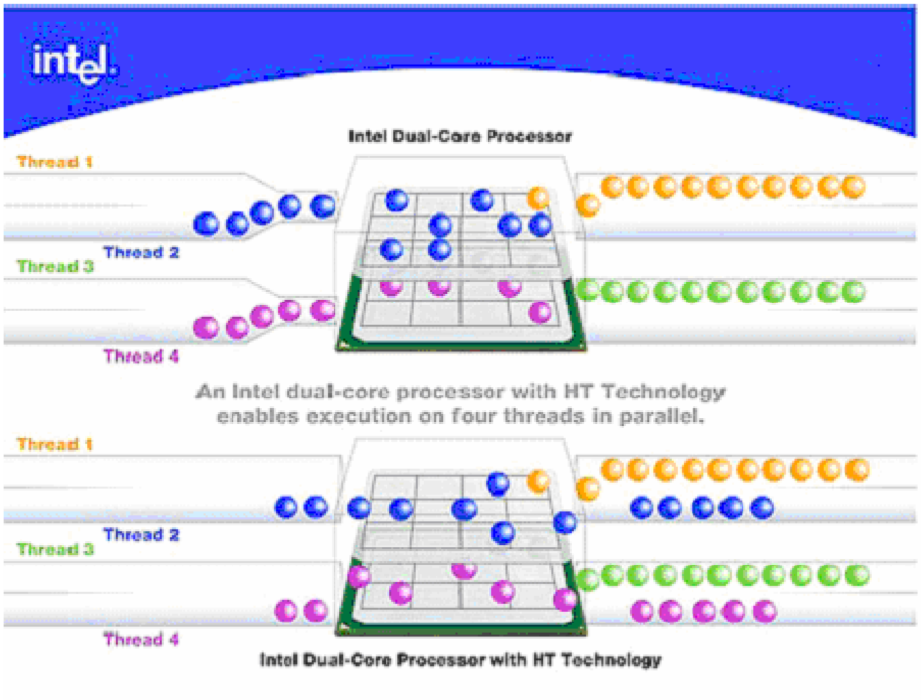

While Hyperthreaded CPUs appear as two logical processors to most operating systems, they are not comparable with true dual-core CPUs since each parallel pair of threads being worked on share the same execution pipeline and same set of L1 and L2 cache memory. Essentially, Hyperthreading is smoke-and-mirrors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

Hyperthreading 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 Hyperthreading 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. Dual processor systems also gain from a general decline in latency. Simply put, while there is no current way to share the current operating system load evenly between two processors, the second processor can step in and keep the system running smoothly while the first is maxed out to 100% burning a CD or encoding a file (or from a software error).
Obviously, if dual-core systems become main stream, which it looks like they are going to, future operating systems and applications will be designed with the feature in mind, leading to better functiona4lity down the road.
Thirdly, and less obviously, AMD and Intel are desperate. Both companies have run into barriers when it comes to increasing the raw speed of processors, or decreasing the die size. Until these roadblocks are cleared or until the general buying public understands that GHz does not directly translate to performance, both companies will be scrambling to discover any new improvements that will improve processor performance... without actually boosting core speed. This is why the idea of dual-core processors is now a reality.

Dual Single-Core vs. Single Dual- Core:

AMD's Opteron chip is capable of SMP due to its multiple hyper transport links, so which is faster; a single dual-core chip or two single-core chips? On paper, dual Opterons should be faster than a single dual-core Opteron at equivalent clock speed for one major reason: Due to the built-in memory controller, each Opteron has exclusive access to its own set of system memory.

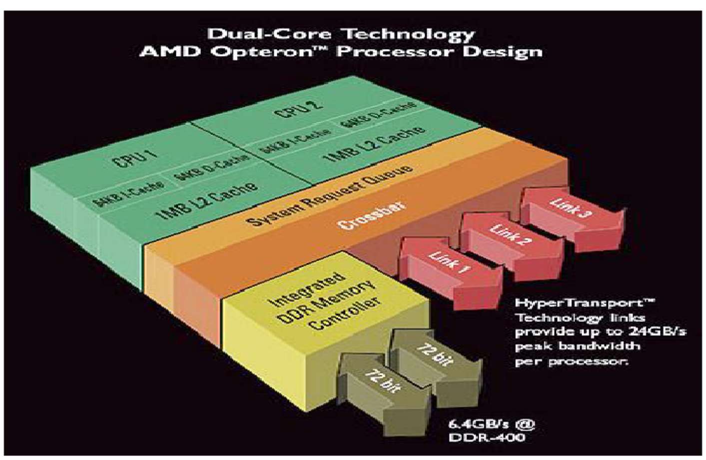

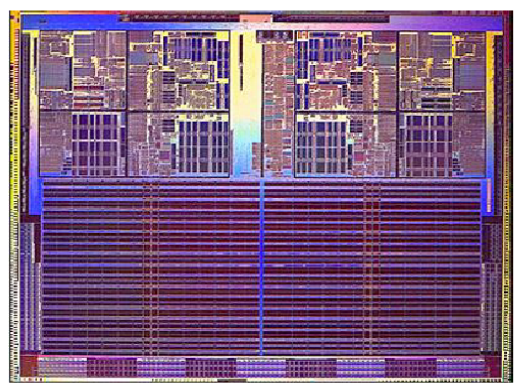

The dual-core designs have to share the memory controller, leading to competition for resources that will inevitably drag down comparative performance.
Intel SMP systems do not gain this advantage over dual-core siblings since they already share a single memory controller over the front-side bus of the motherboard. It's difficult to tell whether either design has any performance advantage in Intel's implementation.
The data has a shorter path to travel with the dual-core chips, but not so much as to make a radical difference. Certainly Intel dual-core chips should have a pricing advantage over SMP solutions, especially when you factor in the price premium that dual-socket motherboards demand.
It's time to talk money. At first glance, basic economics suggests that dual-core processors should be more affordable than buying a pair of single core processors. After all, the companies are integrating two cores into a single die, saving manufacturing effort.
Besides, there would be no point in charging extra money for the second core of a dual-core chip; no one would buy it, right? Maybe, but let's not forget what dual-core chips have to offer besides convenience. The picture is quite different for Intel as opposed to AMD, so let's run through each company's pricing strategies for these chips.

Quad-Core Processors:

In November 2006, Intel introduced the first quad-core microprocessors for the volume x86 markets. The quad-core chips were designed to offer better performance compared with the previous generation of single- and dual-core processors. A little less than a year later Advanced Micro Devices brought its quad-core Opteron to the market, showing that all four cores could be placed on a single piece of silicon. While chipmakers figure out their next chip movies, here are 10 things you should know about quad core processors.
While Intel came to market first with a quad-core processor, it did so essentially by tying two dual-core processors together on the same silicon package. AMD took more time to bring its quad-core chip to market but developed a manufacturing process that placed all four cores on the same piece of silicon.
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 high-end 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.