Parallelism and Next Generation Applications

By Dan Kusnetzky, Principal Analyst
Sponsored by Intel© Corporation

IS MOORE'S LAW STILL GUIDING PROCESSOR DEVELOPMENT?

We've all heard references to "Moore's Law" during conference presentations and in journal articles that are focused upon developments in the world of information technology. What is Moore's Law and does it still drive the world of processor technology as it once did? In sort, the answer is "Yes, but…"

Gordon Moore, one of the founders of Intel Corporation, observed in a presentation in 1965, that the number transistors that could be placed on a chip using low cost manufacturing processes would double every 2 years. During that presentation, he also suggested that this trend would, in all probability, continue into the future. Over time it began to be called "Moore's Law" as others in the industry cited the observation and it proved to be quite accurate. Over time, Moore's Law evolved to mean the doubling of microprocessor performance every two years.

What has this meant over the years? Suppliers of microprocessors have offered amazing advances offering more features, greater performance and a vastly reduced size and cost per transistor.

While faster microprocessors and cost reductions are good, this rapid march of technology presents us with a nasty side effect. As new microprocessors do more and do it faster, the power consumed and heat generated has increased. The law of physics just won't let you keep pushing the chips faster and faster.

Even if the laws of physics and the speed of light won't let the chips run faster, you still have Moore's law increasing the density of transistors. So what can microprocessor developers do with all those transistors to keep delivering performance improvements? The answer can be described in a single word, parallelism. (More information on Intel's view of parallelism can be found here: http://www.intel.com/platforms/parallel.htm.)

PARALLELISM IS THE WAY OF THE FUTURE FOR PROCESSOR ARCHITECTURES

Since power consumption and the resultant production of heat was becoming a serious impediment to future advances in microprocessor design, developers were forced to seek out new ways to increase performance and reduce overall cost. The approach the industry is increasing relying upon is chip design with multiple execution units or cores. Each of these cores run at a lower clock speed, produce far less heat, and still work together "in parallel" to offer improved overall performance at a lower cost.

What does this mean to software developers? It simply means that they no longer can assume that, at any given moment, a single function or thread is being executed. Commercial developers must use parallel computing techniques that have been a mainstay for developers of high performance computing software and developers of operating systems. Commercial developers must learn new techniques, such as "multi-threaded programming", that is designing

their applications so that multiple functions can be executing simultaneously, to make the most of today's multi-core microprocessors.

Intel Corporation and other members of the industry have been investing a great deal of resources to develop tools and processes that simplify the task of allowing today's applications to make better use of today's parallel machines, such as those deploying technology of the Intel® Core™2 processor family. They are also working on development tools and development methodologies that will make good use of future, more highly parallel architectures as they become available. Some of Intel's thoughts on this transformation can be found here: ftp://download.intel.com/technology/computing/archinnov/platform2015/download/Parallelism.pdf

THERE'S NO SUCH THING AS FREE LUNCH

There is no free lunch here. Developers must find better algorithms in order to enjoy optimal performance on today's and tomorrow's parallel systems.

There is no easy way to convert applications designed for serial processing into a parallel application. Developers must understand and use parallelism to make optimal use of these processors. As developers move towards "service oriented architectures" and "multi-tier" application architectures, they're already moving down the right path.

THERE'S NO MAGIC WAY TO CREATE PARALLEL SOFTWARE

Some suppliers of development tools and programming frameworks have hinted that a commercial developer merely needs to acquire one of their products and presto; one of today's serial applications will magically become a multithreaded, parallel application. While these tools *may* improve an application's performance level, the truth is that no automatic tool or process can really understand enough about an application's architecture to magically convert a serial application into a highly optimized parallel application.

The only magic available that can successfully accomplish this feat is a knowledgeable developer who can creatively use the newest tools to develop a better, multithreaded approach to his/her organization's problems. This is the reason that Intel Corporation has invested so heavily in development tools and libraries. More information can be found here:
http://www.intel.com/cd/software/products/asmo-na/eng/index.htm

SUMMARY

Getting the most out of today's dual and quad core microprocessors is an imperative facing developers of commercial applications. This means that their success is tied to understanding parallel programming and how to make the best of these next-generation platforms.