

Moore's Law gets a new lease on life



Don Rippert
Chief Technology Officer
Accenture

In 1965, three years before cofounding Intel, Gordon Moore wrote an article for *Electronics* magazine in which he predicted what would happen in the semiconductor industry during the next 10 years. “Cramming More Components Onto Integrated Circuits” explained how the industry would enjoy reduced costs and increased performance in the coming decade. These observations—which eventually became known as Moore’s Law—have proved to be remarkably durable for the past 40 years.

Moore, at the time the head of R&D at Fairchild Semiconductor, predicted that integrated circuit manufacturers would be able to double the performance of computers every 12 months. Moore later revised the time frame to 24 months to account for the increasing complexity of chips, although the popular expression of Moore’s Law is currently 18 months.

The three underlying trends behind Moore’s observation have held up. First, the numbers of transistors on a chip have increased by five orders of magnitude, from 2,300 in 1971 to more than 169 million today. Second, clock speeds (clocks regulate the rate of processing in a chip) have increased by four orders of magnitude, from 108 kHz in 1971 to 3.8 GHz today. Last, transistor sizes have shrunk by three orders of magnitude.

But the laws of physics have finally caught up with Moore’s Law. The fact is, chips are fighting an increasingly uphill battle against a byproduct of their own processing: heat.

In the mid-1990s, the heat density of the fastest PC chips was equivalent to that of an electric hot plate. As chips decreased in size and increased in speed, their heat density also increased. Unchecked, this growing heat could melt the processor’s adjacent transistors—literally frying the computer’s circuits. Although chip manufacturers have mitigated the problem through recent advances in cooling technologies, the heat barrier continues to be an impediment to the “more, smaller, faster” approach to processor design. Some see that as a death knell for Moore’s Law.

Multi-core revolution

However, manufacturers have started designing chips that continue to increase in performance while sidestepping the heat barriers associated with higher clock speeds. Recently, both Intel and AMD began shipping chips that have multiple processors, or “cores.”

In most cases, each core is slightly slower than the fastest single-core processor. But overall, the multi-core chip is more capable of executing complex computing tasks. The cores share a single clock and the same memory storage, while a micro-kernel on the chip coordinates multiple activity streams—the very feature that enables multi-core processors to achieve better performance without generating excessive heat. Chip manufacturers have essentially redefined the rules of the game through a conceptual work-around of the laws of physics.

At Accenture Technology Labs, we believe these multi-core processing

architectures will give Moore's Law a new lease on life. We see computing power increasing for the foreseeable future—at or even beyond the rate predicted by the Law—through the use of dual-core or multi-core processors on a single chip. The change is already under way, with the major chip manufacturers expected to shift entirely to dual- or multi-core designs in the next few years. This development will have a number of implications for businesses and governments.

We think multi-core systems will be one of the key drivers of server and application service virtualization, which enables greater IT resource utilization and flexibility by allowing virtual machines to run side by side on the same hardware. Enterprises seeking to reduce capital and operating expenditures will consolidate existing infrastructure onto fewer multi-core servers.

Beyond server virtualization, enterprises can expect to see a big boost in performance, especially in key areas such as databases, application servers and data-mining algorithms. Some early adopters have already been experimenting with using multi-core graphics processors to perform business calculations, such as options pricing, with promising results.

Finally, software vendors will be faced with a decision about pricing schemes: either continue pricing by the chip or shift to pricing by the core. We think many will eventually settle on a hybrid solution that accounts for the actual performance benefits that a given piece

of software derives from each additional core.

Today, Accenture is working with major chip manufacturers and software vendors to quantify the expected performance benefits that multi-core architectures will bring to enterprise applications and to evaluate the impact they will have on IT infrastructures. When completed, the results of this research will be discussed in this column and other Accenture venues.

Don Rippert is based in Reston, Virginia.

donald.j.rippert@accenture.com