POINT OF VIEW

The Legacy of Moore's Law

BY FAWWAZ T. ULABY
Editor-in-Chief,
Proceedings of the IEEE

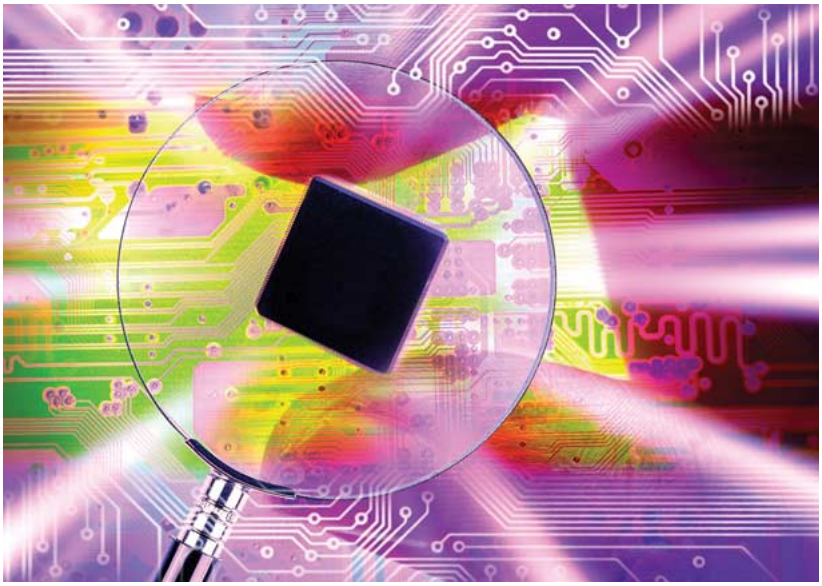

Following an erroneous report of his untimely passing, Mark Twain once said, "The report of my death was an exaggeration." The same might be said of Moore's Law. But while there is still life remaining in Moore's famous axiom, there is no question that it is nearing the end of its run. The question is what happens then.

In April 1965, Gordon Moore, cofounder of Intel, predicted that the number of transistors on a chip of silicon would double every year, an estimate he later changed to every two years. That prediction, known as "Moore's Law," subsequently became the beat to which the industry marched for the next 40 years. In that time, silicon chips diminished in scale even as the number of transistors aboard them doubled with predicted regularity.

Somewhere down the road, however, Moore's Law must face the physical realities posed by the dual processes of diminishing and doubling. The most obvious one is that of scale, the fact that there is no room left on the raft for another doubling of passengers, that "this is as far as you go."

Other factors, though, could spell the end of Moore's Law before the space frontier is reached. In 2004, for example, silicon chips measured a spare 90 nm square. Last year, Intel reduced that to 65 nm and predicted that the ultimate chip would probably measure 5 nm. Below that size, transistors would be approaching the atomic level, raising fundamental barriers involving cost, energy, performance, and manufacturability.

Chips, for instance, may reach a size at which they can no longer dissipate the heat that accumulates when components operate shoulder-to-shoulder in a constricted area. Then there is the difficulty of lithographing circuit patterns on diminishing microchips. Smaller devices require shorter wavelengths for etching. At some point they will reach a size where they exhaust the available spectrum—where the light, in effect, goes out.

Tunneling is another concern. Below 5 nm, the electron source and its drain become so close that the control gate between them can no longer regulate the electron flow. Once that happens, electrons tunneling under the gate generate spurious data transmissions that compromise the system's reliability.

There are nonphysical factors to be considered, as well: the supply of computational power may in time exceed demand, production costs may become prohibitive, or, finally, existing manufacturing methods may prove unable to produce chips below a certain dimension.

Despite these concerns, Moore believes that his law of diminishing returns can probably serve the industry for another 20 years. That draws a line on the horizon. And with that line now in view, many in the industry are asking what is to become of us. They are finding that "the end of the world as we know it" is not the answer.

Digital Object Identifier: 10.1109/JPROC.2006.876941

Over the years, Moore's Law became less of a prediction and more of a self-fulfilling prophecy, a goal that chip makers felt bound to achieve. That they have done so is a tribute to their dedication and engineering skill and has contributed substantially to the great strides made in the field of information processing. But that success has masked the fact that shrinking microprocessors to make them work faster was never an end in itself; it was only the means. The true end was to increase the operating efficiency of integrated circuits, a goal that can be approached by other avenues. Chip makers, alert to the future, are even now heading down those avenues, exploring innovative ways to enhance the performance and profitability of their products within existing speed limits.

One promising approach is the use of multicore processors. In this configuration, two or more processors, operating off a single chip, process data in tandem. As one core handles a calculation, another might summon data, while a third forwards instructions to the operating system. By sharing the workload, this arrangement increases processor performance, using less energy and generating less heat, without the need for running at warp speed.

IBM, Sun, and Hewlett-Packard already sell high-end computers powered by their own multicore chips, and Intel is focusing on delivering a new generation of microprocessors built on that premise. Intel released its first dual-core processor last April and forecasts that by the end of 2006, more than 80% of its server processors and 70% of its personal computers will be multicore. In the meantime, the company is exploring the feasibility of integrating dozens, possibly hundreds, of processors on a single chip as a means of enhancing processor capabilities and performance.

Moore's Law has reached 40 and is still going, but its eventual passing, whether in 10 or 20 years, will not mark the end of an era in semiconductor manufacture. It will signal, instead, an orderly transition to a new one. As Moore himself has said, "It will not be like we hit a brick wall and stop."

If nothing else, Moore's Law demonstrates that technology is not exhaustible but replenishable. As existing technologies reach the limits of their potential, scientists and engineers traditionally seize the opportunity to create alternatives that enable them, not to scrap the achievements of their predecessors, but to build on them. That is the legacy of Moore's Law. ■