THE SLATER PERSPECTIVE

Living Without Moore's Law

Slowing of Semiconductor Process Advances Will Shift Focus of Innovation

I recently spent an interesting three days at the ACM97 conference, which celebrated the Association for Computing Machinery's 50th anniversary and looked forward to what the next 50 years will bring. Presentations from the dozen or so industry luminaries were inevitably uneven—several were outstanding—but most met the goal: to stimulate thinking about where computing is headed. (You can hear the presentations on the Web; see my Web page, listed at the end of this column, for links.)

Fifty years is an extremely long time when it comes to technology forecasting. Market research firms generally place their horizon five years out, and a few brave companies look ten years out. Imagine trying to forecast the state of computer technology today from the vantage point of fifty years ago. The few electronic computers that existed were room-sized and used almost entirely for scientific calculations. Transistors hadn't yet been invented. Forecasting today's desktop and notebook computers, much less the far-reaching impact of the Internet, was virtually impossible—and anyone who made such a forecast would have been dismissed as a kook. (Vannevar Bush's landmark 1945 paper, "As We May Think," was remarkably prescient but still limited in scope.) There is little reason to believe that a 50-year forecast today is likely to be any more accurate.

Although any hope of an accurate forecast would be folly, there are nevertheless a number of interesting issues to ponder. Four themes stood out: the limits of CMOS technology (and what might replace it), the significance of the Internet, the evolution of user interfaces, and the social issues raised by computer technology. In this column, I look at the first of these issues; I'll tackle the others in future columns.

The likelihood that Moore's law will someday cease to apply is of central interest for the microprocessor industry. The time when CMOS will run into a brick wall is clearly in sight: within 15 to 20 years, only a few electrons will fit in a single transistor, and the statistical electron behavior upon which CMOS depends will cease to apply.

At some point, a shift to a different kind of process, making use of the presence or absence of a single electron to indicate a one or a zero, will be needed to continue increasing semiconductor density. Such quantum transistors have been demonstrated in research labs, but enormous barriers remain before complex devices can be fabricated.

CMOS technology does appear capable of scaling for several more generations, but the investment required in new manufacturing technology could slow the doubling interval of Moore's law. Quantum technology has so far to go that it is not clear it will be ready to step in when CMOS runs out of gas. This could result in a period of stagnation in semiconductor density and performance improvements.

The microprocessor industry has always been driven by ever-advancing process technology. What would happen if semiconductor progress stagnates, as seems possible somewhere between 2010 and 2020? It would become harder to drive PC replacements, which are a major part of the PC economy. Embedded microprocessors might have fewer opportunities to expand into new markets.

It is reassuring, however, to consider where the technology will be when it approaches these limits. Microprocessors selling for a few hundred dollars will have hundreds of millions of transistors, if not billions. A single such chip could have dozens of processors, each with several times the complexity of today's most advanced devices, plus several megabytes of cache for each. Running at several gigahertz, the chip could include a video and 3D graphics system, peripheral controllers, network interface, modem, and so forth.

A system could be built with everything in the fastest workstation today—including memory—in a single chip. A $10 microcontroller will be faster than the fastest microprocessor today and have a full set of peripherals. A single 4-Gbit DRAM chip will provide 512 Mbytes of storage—the amount in today's enterprise servers. A dozen such chips together will match the capacity of today's largest disk drives.

In the early days of microprocessors, limited transistor budgets severely constrained what could be built. Today, the ability to design complex chips is as big a limitation as the silicon itself. In another decade, the canvas that CMOS will offer silicon designers will be so vast that it will take decades to explore the creative opportunity. The simple increases in cache size and clock rate may slow dramatically, but there will still be tremendous room for innovation in system (i.e., chip) architecture. Even raw transistor count may continue to increase through 3D fabrication techniques or larger die.

Neural computing offers a tantalizing prospect for replacing silicon technology someday. It may be "dry," using silicon for implementation, or "wet," using chemical or biological processes. Predicting when—or even if—this technology will be commercially viable is nearly impossible, but I suspect that most of us alive today won't live to see silicon displaced as the predominant computing technology.

See www.MDRonline.com/slater/silicon *for more on this subject. I welcome your feedback at* mslater@mdr.zd.com.