Moshe Y. Vardi

DOI:10.1145/2600347

Moore's Law and the Sand-Heap Paradox

The Sand-Heap Paradox is an ancient Greek paradox, which considers the failure of induction. Imagine a heap of sand from which one grain is removed. Surely what is left

is a heap of sand. But if we repeat this process long enough, we are left with a single grain of sand, which is not a heap.

Many of us have lived with Moore's Law for all of our professional lives. We knew that Moore's Law—the doubling of the number of transistors on a chip every couple of years—cannot continue forever, but the end of Moore's Law always seemed to be beyond the horizon. No more. It is now becoming clear that we are witnessing the denouement of an extraordinary technical saga.

In fact, this denouement has been going on for the past decade. As I wrote in “Is Moore's Party Over?” (Nov. 2011), Dennard Scaling, described by IBM's Robert Dennard in 1974, which asserted that as transistors got smaller, power density stays constant, already broke down about 10 years ago. As I wrote then, that meant increased transistor density no longer automatically leads to improved computer performance. While the semiconductor industry has been able to continue innovating, reducing transistor size, and increasing transistor density, there are more and more signs that Moore's Law is in a serious trouble.

To start, the terminology used by the industry to describe its ongoing march to the drumbeat of Moore's Law has turned from physics to marketing. Back in the days of 0.35-micrometer chips, the number referred to transistor gate lengths. But today, as Intel starts production of 14-nanometer chips, it is not clear at all what this

number means other than a suggestive reference to the continuing increase in transistor density.

While the industry has been struggling to harness transistor density to deliver performance, it is also being challenged from the business side. After all, the real point of Moore's Law was not merely delivering improved performance, but delivering improved cost-performance, which meant we got improved performance at a fixed and even reduced cost. No more. The Linley Group, a semiconductor consultancy, pointed out last year that while in 2012 one could buy 20M 28-nanometer transistors per dollar, the forecast for 2015 is 19M 16-nanometer transistors per dollar. Such a rise in the cost of transistors is simply unprecedented.

Finally, at the current rate of progress we will reach the five-nanometer milestone within 10–15 years, and there are strong technical arguments why CMOS, the semiconductor technology that served us well for decades, cannot be scaled down further. Indeed, Robert Colwell, currently at DARPA and previously chief IA-32 architect at Intel, recently declared publicly that he expects Moore's Law to die around 2020. In a recent analysis, Andrew A. Chien and Vijay Karamcheti argued that when it comes to flash memories Moore's Law has already ended and increases in capacity will be accompanied by reduced reliability and performance.

While there are numerous alternatives to CMOS technology, it is doubtful any one of them will be mature enough to become the workhorse of the semiconductor industry in 10 years.

So the real question is not when precisely Moore's Law will die; one can say it is already a walking dead. The real question is what happens now, when the force that has been driving our field for the past 50 years is dissipating. In fact, Moore's Law has shaped much of the modern world we see around us. A recent McKinsey study ascribed “up to 40% of the global productivity growth achieved during the last two decades to the expansion of information and communication technologies made possible by semiconductor performance and cost improvements.” Indeed, the demise of Moore's Law is one reason some economists predict a “great stagnation” (see my Sept. 2013 column).

“Predictions are difficult,” it is said, “especially about the future.” The only safe bet is that the next 20 years will be “interesting times.” On one hand, since Moore's Law will not be handing us improved performance on a silver platter, we will have to deliver performance the hard way, by improved algorithms and systems. This is a great opportunity for computing research. On the other hand, it is possible that the industry would experience technological commoditization, leading to reduced profitability. Without healthy profit margins to plow into research and development, innovation may slow down and the transition to the post-CMOS world may be long, slow, and agonizing.

However things unfold, we must accept that Moore's Law is dying, and we are heading into an uncharted territory.

Follow me on Facebook, Google+, and Twitter.

Moshe Y. Vardi, EDITOR-IN-CHIEF

Copyright held by Author.