TECHNOLOGY LAB —

Moore’s law really is dead this time

The chip industry is no longer going to treat Gordon Moore's law as the target to aim for.

PETER BRIGHT - 2/10/2016, 8:22 PM

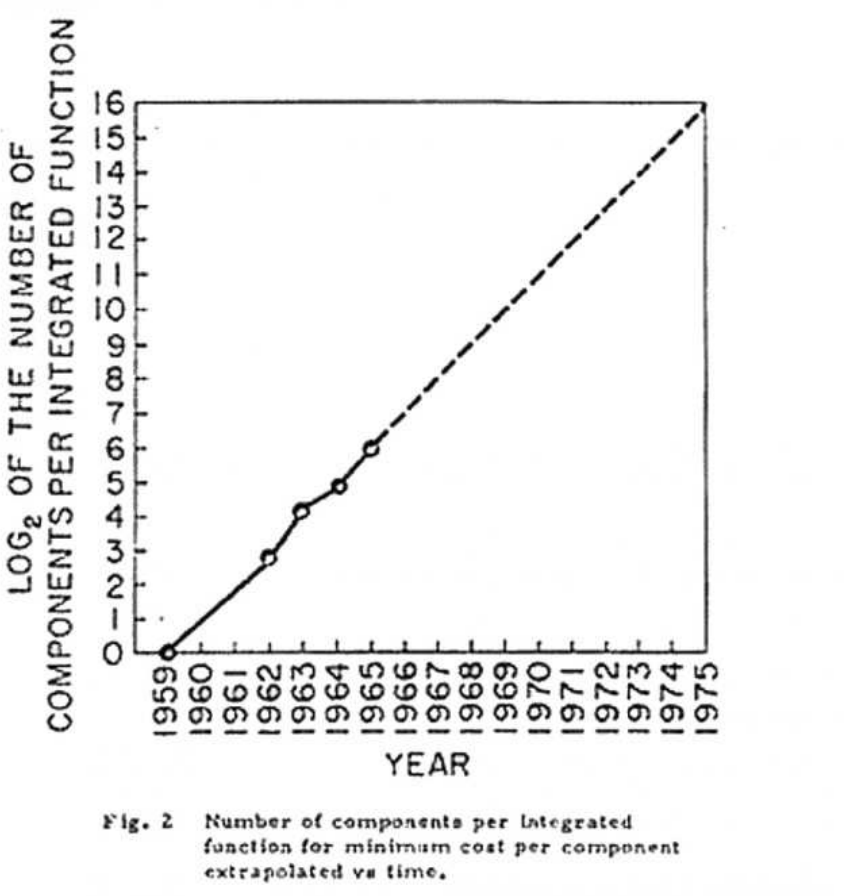

Intel

Gordon Moore's original graph, showing projected transistor counts, long before the term "Moore's law" was coined.

227

Moore's law has died at the age of 51 after an extended illness.

In 1965, Intel co-founder Gordon Moore made an observation that the number of components in integrated circuits was doubling every 12 months or so. Moreover, as this site wrote extensively about in 2003, that the number of transistors per chip that resulted in the lowest price per transistor was doubling every 12 months. In 1965, this meant that 50 transistors per chip offered the lowest per-transistor cost; Moore predicted that by 1970, this would rise to 1,000 components per chip, and that the price per transistor would drop by 90 percent.

With a little more data and some simplification, this observation became "Moore's law": the number of transistors per chip would double every 12 months.

Gordon Moore's observation was not driven by any particular scientific or engineering necessity. It was a reflection on just how things happened to turn out. The silicon chip industry took note and started using it not merely as a descriptive, predictive observation, but as a prescriptive, positive law: a target that the entire industry should hit.

Hitting this target didn't happen by accident. Building a silicon chip is a complex process, and it uses machinery, software, and raw materials that are sourced from a number of different companies. To ensure that all the different players are aligned and working on compatible timetables to preserve Moore's law, the industry has published roadmaps laying out the expected technologies and transitions that will be needed to preserve Moore's law. The Semiconductor Industry Association, a predominantly North American group that includes Intel, AMD, TSMC, GlobalFoundries, and IBM, started publishing roadmaps in 1992, and in 1998 the SIA joined up with similar organizations around the world to produce the International Technology Roadmap for Semiconductors. The most recent roadmap was published in 2013.

Problems with the original formulation of Moore's law became apparent at an early date. In 1975, with more empirical data available, Gordon Moore himself updated the law to have a doubling time of 24 months rather than the initial 12. Still, for three decades, simple geometric scaling—just making everything on a chip smaller—enabled steady shrinks and conformed with Moore's prediction.

In the 2000s, it was clear that this geometric scaling was at an end, but various technical measures were devised to keep pace of the Moore's law curves. At 90nm, strained silicon was introduced; at 45nm, new materials to increase the capacitance of each transistor layered on the silicon were introduced. At 22nm, tri-gate transistors maintained the scaling.

But even these new techniques were up against a wall. The photolithography process used to transfer the chip patterns to the silicon wafer has been under considerable pressure: currently, light with a 193 nanometre wavelength is used to create chips with features just 14 nanometres. The oversized light wavelength is not insurmountable but adds extra complexity and cost to the manufacturing process. It has long been hoped that extreme UV, with a 13.5nm wavelength, will ease this constraint, but production-ready EUV technology has proven difficult to engineer.

Even with EUV, it's unclear just how much further scaling is even possible; at 2nm, transistors would be just 10 atoms wide, and it's unlikely that they'd operate reliably at such a small scale. Even if these problems were resolved, the specter of power usage and dissipation looms large: as the transistors are packed ever tighter, dissipating the energy that they use becomes ever harder.

The new techniques, such as strained silicon and tri-gate transistors, took more than a decade to put in production. EUV has been talked about for longer still. There's also a significant cost factor. There's a kind of undesired counterpart to Moore's law, Rock's law, which observes that the cost of a chip fabrication plant doubles every 4 years. Technology may provide ways to further increase the number of transistors packed into a chip, but the manufacturing facilities to build these chips may be prohibitively expensive—a situation compounded by the growing use of smaller, cheaper processors.

We've recently seen these factors cause real problems for chip companies. Intel originally planned to switch to 10nm in 2016 with the Cannonlake processor, a shrunk version of the 14nm Skylakes shipping today. In July last year, the company changed this plan. An extra processor generation, Kaby Lake, will be released in 2016, still using the 14nm process. Cannonlake and 10nm are still planned but are not due until the second half of 2017.

Compounding all this is that all these extra transistors have become increasingly hard to use. In the 1980s and 1990s the value of the extra transistors was obvious: the Pentium was much faster than the 486, the Pentium II much faster than the Pentium, and so on and so forth. Existing workloads gained substantial speed-ups just from processor upgrades, thanks to a combination of better processors (going from simple in-order processors to complex superscalar out-of-order processors) and higher clockspeeds. Those easy improvements stopped coming in the 2000s. Constrained by heat, clock speeds have largely stood still, and the performance of each individual processor core has increased only incrementally. What we see instead are multiple

processor cores within a single chip. This increases the overall theoretical performance of a processor, but it can be difficult to actually exploit this improvement in software.

These difficulties mean that the Moore's law-driven roadmap is now at an end. ITRS decided in 2014 that its next roadmap would no longer be beholden to Moore's "law," and *Nature* writes that the next ITRS roadmap, published next month, will instead take a different approach.

Rather than focus on the technology used in the chips, the new roadmap will take an approach it describes as "More than Moore." The growth of smartphones and Internet of Things, for example, means that a diverse array of sensors and low power processors are now of great importance to chip companies. The highly integrated chips used in these devices mean that it's desirable to build processors that aren't just logic and cache, but which also include RAM, power regulation, analog components for GPS, cellular, and Wi-Fi radios, or even microelectromechanical components such as gyroscopes and accelerometers.

These different kinds of component traditionally use different manufacturing processes to handle their different needs, and the new roadmap will outline plans for bringing them together. Integrating the different manufacturing processes and handling the different materials will need new processes and supporting technology. For manufacturers building chips for these new markets, addressing this kind of problem is arguably more relevant than slavishly doubling the number of logic transistors.

There will also be a focus on new technology beyond the silicon CMOS process currently used. Intel has already announced that it will be dropping silicon at 7nm. Indium antimonide (InSb) and indium gallium arsenide (InGaAs) have both shown promise, and both offer much higher switching speeds at much lower power than silicon. Carbon, both in its nanotube and graphene forms, continues to be investigated and may prove better still.

While a lesser priority, scaling is not off the roadmap entirely. Beyond tri-gate transistors, perhaps around 2020, are "gate all around" transistors and nanowires. The mid-2020s could bring monolithic 3D chips, where a single piece of silicon has multiple layers of components that are built up on a single die.

As for the future, massive scaling isn't off the cards completely. The use of alternative materials, different quantum effects, or even more exotic techniques such as superconducting may provide a way to bring back the easy scaling that was enjoyed for decades, or even the more complex scaling of the last fifteen years. A big enough boost could even reinvigorate the demand for processors that are just plain *faster*, rather than smaller or lower power.

But for now, lawbreaking is going to be the new normal. Moore's law's time as a guide of what will come next, and as a rule to be followed, is at an end.

PETER BRIGHT
Peter is Technology Editor at Ars. He covers Microsoft, programming and software development, Web technology and browsers, and security. He is based in Brooklyn, NY.

EMAIL peter.bright@arstechnica.com // **TWITTER** @drpizza

READER COMMENTS 227 SHARE THIS STORY

← PREVIOUS STORY NEXT STORY →

Related Stories

Intel confirms tick-tock-shattering Kaby Lake processor as Moore’s Law falters

Intel retires “tick-tock” development model, extending the life of each process

Intel’s Xeon brand makes its first foray into SoC space with Xeon D

Skylake users given 18 months to upgrade to Windows 10

Today on Ars

The little yellow box that has made thousands of operations safer

Drug maker pushes pause on $89K drug after outcry

Apple waits for LG’s fix before selling any more 5K Thunderbolt displays

Now sites can fingerprint you online even when you use multiple browsers

How a 500-year-old engineering innovation led to selfies

How to not do presidential opsec: Crisis management over dinner in public

Potential risks of going gluten-free—high arsenic and mercury levels

T-Mobile answers Verizon by adding HD video and hotspot to unlimited plan