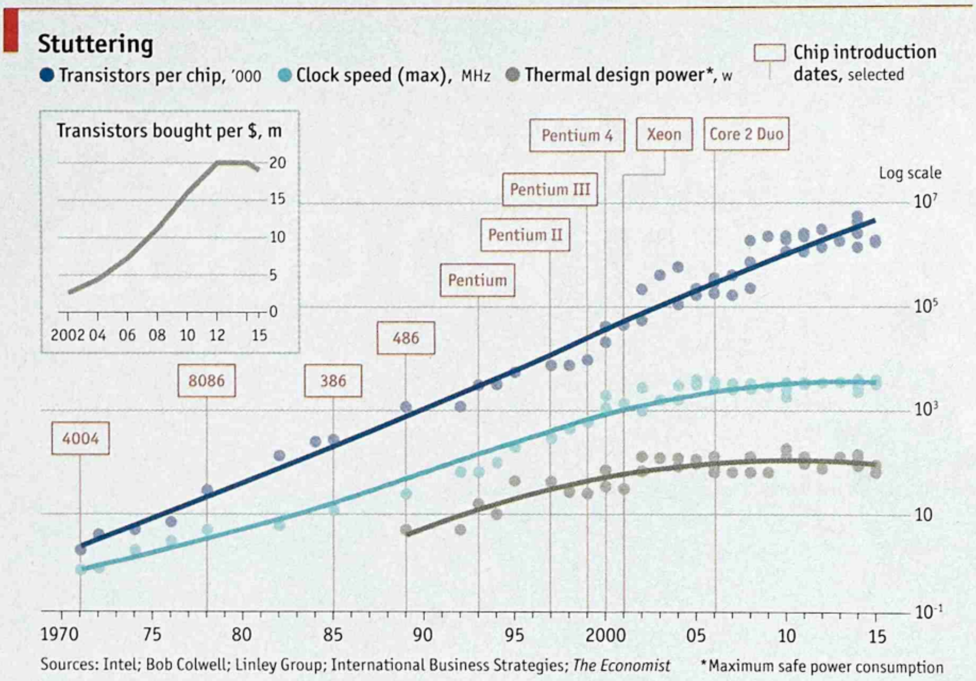

been the leading maker of microprocessors, and the rest of the industry turned it into a self-fulfilling prophecy. That fulfilment was made possible largely because transistors have the unusual quality of getting better as they get smaller; a small transistor can be turned on and off with less power and at greater speeds than a larger one. This meant that you could use more and faster transistors without needing more power or generating more waste heat, and thus that chips could get bigger as well as better.

Making chips bigger and transistors smaller was not easy; semiconductor companies have for decades spent heavily on R&D, and the facilities—"fabs"—in which the chips have been made have become much more expensive. But each time transistors shrank, and the chips made out of them became faster and more capable, the market for them grew, allowing the makers to recoup their R&D costs and reinvest in yet more research to make their products still tinier.

The demise of this virtuous circle has been predicted many times. "There's a law about Moore's law," jokes Peter Lee, a vice-president at Microsoft Research: "The number of people predicting the death of Moore's law doubles every two years." But now the computer industry is increasingly aware that the jig will soon be up. For some time, making transistors smaller has no longer been making them more energy-efficient; as a result, the operating speed of high-end chips has been on a plateau since the mid-2000s (see chart). And while the benefits of making things smaller have been decreasing, the costs have been rising.

This is in large part because the components are approaching a fundamental limit of smallness: the atom. A Skylake transistor is around 100 atoms across, and the fewer atoms you have, the harder it becomes to store and manipulate electronic 1s and 0s. Smaller transistors now need trickier designs and extra materials. And as chips get harder to make, fabs get ever more expensive.

Handel Jones, the CEO of International Business Strategies, reckons that a fab for state-of-the-art microprocessors now costs around $7 billion. He thinks that by the time the industry produces 5nm chips (which at past rates of progress might be in the early 2020s), this could rise to over $16 billion, or nearly a third of Intel's current annual revenue. In 2015 that revenue, at $55.4 billion, was only 2% more than in 2011. Such slow increases in revenue and big increases in cost seem to point to an obvious conclusion. "From an economic standpoint, Moore's law is over," says Linley Gwennap, who runs the Linley Group, a firm of Silicon Valley analysts.

The pace of advance has been slowing for a while. Marc Snir, a supercomputing expert at Argonne National Laboratory, Illinois, points out that the industry's International Technology Roadmap for Semiconductors, a collaborative document that tries to forecast the near future of chipmaking, has been over-optimistic for a decade. Promised manufacturing innovations have proved more difficult than expected, arriving years late or not at all.

Brian Krzanich, Intel's boss, has publicly admitted that the firm's rate of progress has slowed. Intel has a biennial "tick-tock" strategy: in one year it will bring out a chip featuring smaller transistors ("tick"); the following year it tweaks that chip's design ("tock") and prepares to shrink the transistors again in the following year. But when its first 14nm chips, codenamed Broadwell, ticked their way to market in 2014 they were nearly a year behind schedule. The tick to 10nm that was meant to follow the tock of the Skylakes has slipped too; Intel has said such products will not now arrive until 2017. Analysts reckon that because of technological problems the company is now on a "tick-tock-tock" cycle. Other big chipmakers have had similar problems.

Moore's law has not hit a brick wall. Chipmakers are spending billions on new designs and materials that may make transistors amenable to a bit more shrinkage and allow another few turns of the exponential crank. They are also exploring ways in which performance can be improved with customised designs and cleverer programming. In the past the relentless doubling and redoubling of computing power meant there was less of an incentive to experiment with other sorts of improvement.

Try a different route

More radically, some hope to redefine the computer itself. One idea is to harness quantum mechanics to perform certain calculations much faster than any classical computer could ever hope to do. Another is to emulate biological brains, which perform impressive feats using very little energy. Yet another is to diffuse computer power rather than concentrating it, spreading the ability to calculate and communicate across an ever greater range of everyday objects in the nascent internet of things.

Moore's law provided an unprecedented combination of blistering progress and certainty about the near future. As that certainty wanes, the effects could be felt far beyond the chipmakers faced with new challenges and costs. In a world where so many things—from the cruising speed of airliners to the median wage—seem to change little from decade to decade, the exponential growth in computing power underlies the future plans of technology providers working on everything from augmented-reality headsets to self-driving cars. More important, it has come to stand in the imagination for progress itself. If something like it cannot be salvaged, the world would look a grimmer place.

One idea is to redefine the computer itself

At the same time, some see benefits in a less predictable world that gives all sorts of new computing technologies an opportunity to come into their own. "The end of Moore's law could be an inflection point," says Microsoft's Dr Lee. "It's full of challenges—but it's also a chance to strike out in different directions, and to really shake things up." ■