


Mythology surrounds the law that seems to drive the semiconductor industry. What's the truth?

**Chris Edwards** investigates.



# the many lives of moore's law

ASK SOMEONE what Moore's Law means, and the chances are that they will tell you it means a doubling in the density of integrated circuits every 18 months. And they would be wrong.

The 18-month period is a myth with surprising endurance. Part of the reason is that Intel spent the latter half of the 1990s trying to convince people that they were seeing a doubling in processor capability every 18 months and used Moore's Law as the basis for their claim. But the law's creator, Intel founder

Gordon Moore, has claimed that he never said it was 18 months, even though people have been keen to say that he did.

Take a keynote speech that Moore delivered at Intel's developers' forum in 1997. The session's compere was confident he knew the story of how Moore came up with his projection: "And I happen to be privy to the insight of how this really happened. You see, Gordon is a deep sea fisherman, he loves to go fishing, and early back in 1965 or so, he was out fishing one day, and he was musing on the

fact that he had noticed over the last year or so that he was actually able to get a couple more transistors on that die.

And he thought, 'Hmm, I wonder how often these things are actually going to double in terms of the number of transistors I can get on the chip.'

"And he thought, 'I know; however many fish I catch today; that's how many years it's going to be.' He got 18 fish; it was 18 months."


It's a cute story. But wrong. In a later interview, Moore claimed that Intel's then marketing

manager Dave House came up with the number to factor in increases in both density and clock speed. Perhaps that is why, as an Intel executive, Moore did not point out in his speech that the story relayed on his behalf was complete baloney.

However, if you wind down the rate of change to two years or so, Moore's Law has turned out to be surprisingly accurate. For close to 30 years, the rate of growth in capacity of integrated circuits has stuck to more or less the same value: around 40 per cent per year. But the true long-term growth rate did not emerge until ten years after Moore made his original prediction.


When Moore wrote the original article in 1965 for the magazine *Electronics*, scaling was happening much more quickly than it is today. From the early 1960s through to the mid-1970s, the rate of growth was doubling every year, although Moore later explained that he was working with a very

In 1975, Gordon Moore changed the rate of growth from a doubling every year to a 100 per cent increase in chip capacity every two years. He made three assumptions for that projection (right). Although the assumptions were flawed, the rate of growth has stuck close to the two-year doubling as shown by the graph below, which, for the first time, brings together a much broader set of samples. The graph shows the role that die size plays in bringing the highest-capacity chips into production. The size of each point is proportional to area of the actual die.



'I blindly extrapolated for about ten years and said OK, in 1975 we'll have about 60,000 components on a chip. I had no idea this was going to be an accurate prediction.'

Gordon Moore



Meeting (IEDM) together with much greater detail on his assumptions. However, even though his overall prediction held, few of the assumptions held up over the following 30 years.

### EMBRACING THE LAW

In looking back at Moore's Law, we did not just take Intel's own figures but used them in combination with data culled from 30 years of papers presented at the International Solid-State Circuits Conference (ISSCC) and other conferences. They range from processing units for mainframe computers through system-on-chip devices for mobile phones. The much broader range of devices in the graph on the opposite page does not produce as neat a straight line as the conventional Intel-only data, but the data still supports a growth rate in the range of two to two-and-a-half years. It demonstrates why chipmakers took Moore's

Law to their hearts after the term was coined by Professor Carver Mead, the co-inventor of one of the design techniques that made massive digital integration feasible.

What is most surprising is the consistency of the growth rate since the mid-1970s given that all of the techniques that underpinned Moore's assumptions used to guide the projection have changed radically. The one that he considered to be the least important in 1975, as shown in Moore's 1975 graph above, turned out to be the main driver for more than 20 years.

In his 1975 IEDM speech, Moore claimed there were three main components to improvements in integration. He saw increases in die size as coming up with almost half the growth in transistor count with reductions in their dimensions contributing far less. However, taken together, the two factors were seen as providing ►

small set of samples. "We [had] made a few circuits and gotten up to 30 circuits on the most complex chips that were out there in the laboratory, we were working on about 60, and I looked and said, gee, in fact from the days of the original planar transistor, which was 1959, we had about doubled every year the amount of components we could put on a chip. So I blindly extrapolated for about ten years and said OK, in 1975 we'll have about 60,000 components on a chip... I had no idea this was going to be an accurate prediction."

### SPANNER IN THE WORKS

But, the industry was headed for wrenching changes as the ten-year period drew to a close. A crushing recession in 1974 of the kind that would not be seen for more than 25 years put a brake on development. The pace of development slowed dramatically and, naturally, had an effect on the rate at which chip-level electronic circuits shrank. Moore presented a revised projection in 1975 at the International Electron Device



Intel co-founder Gordon Moore talks about his law at the company's developers forum last year

### truth & lies

#### WAGGING THE DOG






That people can go through their careers believing that Moore's Law means a doubling in capacity every 18 months when all the evidence points to the contrary is an indication that the number itself is not all that important. What matters is that people believe that, to stay competitive, they need to keep pushing at the density to keep up with the competition.

Looked upon this way, Moore's Law is not all that different from what happens in other industries. This is a point that Mentor chairman and CEO Wally Rhines has made in several of his speeches at semiconductor and electronics-design conferences.

Rhines' contention is that technologies that exhibit any form of learning curve benefit massively from a lot of people doing the learning. Semiconductor production is one of those technologies. Part of that is through the sharing of knowledge, some of it deliberate, some of it unwitting. But there is also the effect of seeing that something is possible: it lets others know that there is an answer available even if they cannot simply go and ask for the solution from the person who has first-mover advantage.



The part of Moore's Law that chipmakers have learned to break is the growth in die size. On the right, a Pentium from 1993 sits on top of a wafer of the latest PENryn processors, noticeably smaller than their predecessor. Die sizes for DRAMs and processors grew at a rate close to Moore's 1975 prediction up to the mid-1990s. The graph below shows the growth in die size for DRAMs. The size of each point is proportional to the area of the basic storage cell and they are colour-coded by capacity. The graph at the bottom shows the growth in the size of Intel's processors. The large orange blocks are the company's recent Xeon and Itanium server processors. Although these are large, Intel has, wherever possible, strived to constrain the size of its desktop processors to the 100-200mm<sup>2</sup> range.



◀ two-thirds of the growth needed for what was to become Moore's Law. The remainder lay in the "contribution of device and circuit cleverness" – these were the architectural changes being made to circuits to improve overall density.

During the 1970s, increases in die size did a lot to keep manufacturers close to Moore's predictions. At the IEDM conference, he told delegates

that, based on historical growth in die size, they could expect to see, within five years, devices measuring 90,000mil<sup>2</sup> – using the then US standard measurement of a thousandth of an inch – or about 58mm<sup>2</sup>. In fact, the industry did a little better although it was not on the parts that Moore expected. Originally, he felt memories would do better. At the time, charge-coupled devices (CCD)

looked as though they would make low-cost, high-density memories. But they had problems and were ultimately replaced by dynamic random-access memory (DRAM) devices. The switch meant that memory capacity slightly lagged Moore's 1975 projection.

It is possible that there were bigger chips around, but the largest presented at the International Solid-State Circuits Conference (ISSCC) in 1980 – a Fujitsu microprocessor – clocked in at 111mm<sup>2</sup>. The memories, however, were way smaller, struggling to get above 40mm<sup>2</sup>.

### COMPLEX CHIP DESIGN

Circuit cleverness was a major component of Moore's Law in the early days: engineers took a number of years to learn the ways to pack more transistors into smaller and smaller spaces. By 1975, there did not seem, to Moore, a great deal of cleverness left to find. So, he predicted a slowdown in the rate of growth in chip density.

When it comes to logic chips, you could argue that circuit density went into reverse. Moore pointed to the problem in 1975: that it would get harder and harder to design chips as they became more complex. The response from the industry was to employ more automation, such as the techniques pioneered by Lynn Conway and Carver Mead. The techniques needed to make that automation possible generally led to less space-efficient circuitry – a brake on Moore's Law.

Microprocessors stayed closer to the predicted trend because they used a higher proportion of hand-drawn circuits compared with chips intended for things that would ship in lower volume, such as network equipment or industrial controls. You can see the consequence of this in the Moore's Law graph – a 'broadening' of the mass of points from the late 1980s until the late 1990s, when most companies moved to use the same design techniques and when on-chip memory became much more important to overall density.

Into the 1980s and 1990s, die size continued to grow. But, it became apparent that it would prove difficult to scale much beyond 100mm<sup>2</sup> and still maintain high yield. Companies were able to introduce larger and larger parts but at significant cost. The first two generations of the Pentium were on the line projected by Moore, but they were expensive.

Intel found that it needed to go to more advanced processes to produce 'shrinks' of the originals to be able to offer the processors at reasonable cost. It was a similar situation for memory. At introduction, a DRAM of the early 1980s, such as a 256Kb device from 1982, might cost \$51. That was higher than the \$20 charged for a 1974 device providing 4Kb. But it was a snip compared with the \$575 charged for the first 256Mb memories made in 1998. The skills of chipmakers in designing bigger chips was outstripping their ability to maintain yield. The trend in die size was set to level off dramatically.

Chipmakers were able to scale up the size of chips to keep pace with Moore's Law but they were losing more to scrap because larger dice have a higher probability of attracting serious faults. The die size of a 1998 DRAM that could carry 256Mb was more than four times bigger than the 256Kb chip of 1982. The 1998 device worked out to be more than seven times more expensive, after adjusting for inflation, than the smaller chip made on more primitive, lower-yielding fab lines. That is an indication of how yield plays a major part in chip cost.


Although chipmakers pumped out research devices that broke through the 1Gb barrier before the end of the 1990s, they were way too big to be produced economically. The result was that memory makers had to wait for improvements in transistor density to introduce the first commercial 1Gb devices. They did not appear until 2003 and, in an indication of how die size stopped scaling, they were smaller than the first

## THE AGE OF METAL

One of the rarely mentioned aspects of chip design is just how much space is wasted when implementing logic. Estimates indicate that, on most digital chips around 30 per cent is 'white space'. Part of the reason is that the overall density is limited by congestion in the wiring layers. Increasing the number of metal

layers helps alleviate congestion and increase overall density. As chipmakers moved away from increasing die size as a way of increasing transistor integration, they started to increase the number of metal layers as a way of improving device density. The increase in metal layers was helped by the development of chemi-

cal-mechanical polishing during the early 1990s. Before then, only two or three metal layers were feasible. With planarisation, it is now possible to build chips with as many as 12 layers of metal. As physics places limits on just how small some features can get, chipmakers will look more at what they can do in the third dimension.



256Mb chips. They were still costly as, measuring some 170mm<sup>2</sup>, they were above the range that memory makers like to stay in. It is not until they get closer to the 100mm<sup>2</sup> point that modern memories become mainstream, high-volume products. Similarly, mass-production microprocessors tend to occupy a zone between 100mm<sup>2</sup> and 200mm<sup>2</sup>, shown as orange in the graph of Intel processors.

Companies such as Intel were able to push die size but only for products destined to be used in comparatively low-volume applications, such as compute servers. Examples are the Itanium processors made in the early part of this decade. They approached the maximum size possible for any chip – the size of the reticle used to hold the masks that produce the images which need to be printed onto the silicon surface. For a brief while, it made it possible for Intel to beat the long-term average. But such growth had to come to an end.

### TICK TOCK STRATEGY

What took over? It was Moore's least favourite option in 1975: scaling in device size. And, to some extent, circuit cleverness. When die size growth stalled from about the mid-1990s, chipmakers had to find other ways to improve integration. One response was to increase the rate at which they introduced new, finer-geometry processes. In the 1970s, it took many years to introduce a new process. Slowly, the rate increased to a new generation about every three years. Now, we are seeing a new process come in every two years. This is what Intel president Paul Otellini called the 'tick-tock strategy'.

However, device scaling on its own is not enough to drive a doubling in chip capacity every two years. Fortunately for most chipmakers, the rise in on-chip memory has provided them with a way of increasing capacity beyond what is supported by linear scaling. The manufacturers have been subtly altering SRAM cells to improve their

density. As the amount of SRAM has soared, chip density has kept pace with projections. Then there is the contribution from using more layers to connect devices together, as shown in the graph above.

The result is that, as an example, Intel's latest processor, the 45nm Penryn, lies squarely on the Moore's Law line. Despite all the changes to the underlying assumptions, the law has been good for 30 years.

Even though people now believe SRAM scaling is reaching its limit, the law could still remain valid for another ten years thanks to further changes in the way that chips are designed. ■

PowerTools are renowned for being a long term investment because of their open and modular architecture.

## LEADING through Technology



Great Britain  
Lauterbach Ltd.  
[www.lauterbach.co.uk](http://www.lauterbach.co.uk)

Germany  
Lauterbach GmbH  
[www.lauterbach.de](http://www.lauterbach.de)

Italy  
Lauterbach Srl  
[www.lauterbach.it](http://www.lauterbach.it)

USA  
Lauterbach Inc.  
[www.lauterbach.com](http://www.lauterbach.com)

China  
Suzhou Lauterbach  
Technologies Co., Ltd.  
[www.lauterbach.cn](http://www.lauterbach.cn)

Japan  
Lauterbach Japan Ltd.  
[www.lauterbach.co.jp](http://www.lauterbach.co.jp)

### PowerTrace II

- 4 GByte of trace memory to record program and data flow
- More than 40 processor architectures supported by the trace port
- Trace port speed up to 500 MHz
- Sophisticated analysis methods to perform comprehensive performance analysis and quality assurance tests



### PowerDebug II

- A Gigabit ethernet interface guarantees immediate display and rapid analysis of the trace information.

**LAUTERBACH**