Electronics 101.9: Integrated Circuits

Making things smaller

Dave Astels

daveastels.com

Dave's career started in the 8-bit days, with the Z80 and 6502, and he's been working with computers ever since. Check him out at: **daveastels.com** and **learn.adafruit.com**

We've seen several kinds of components so far. All of them are fairly big. Even if we consider the smallest surface mount resistors, capacitors, and so on, they're still quite large. This wasn't a 'big' problem back in the day, when circuits were fairly simple. As we started designing and building more complex circuits, they got physically larger, and connections got more complex. At the same time, demand for electronic devices grew. Building devices was time-consuming and error prone. **Figure 1** shows a simple logic circuit built from discrete components. Many modules must be connected, typically by using a wire-wrapped backplane – see **Figure 2**. There had to be a better way. There was: the integrated circuit [IC].

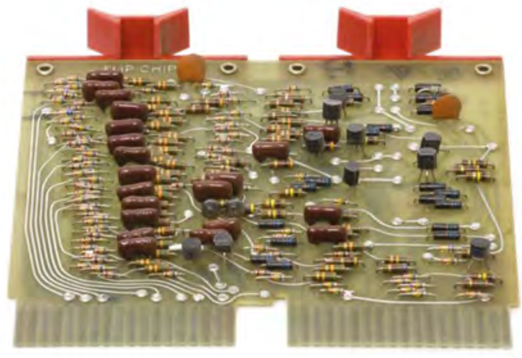

Figure 1
A transistor/diode-based logic module

THE CHIPS ARE DOWN

Jack Kilby, at Texas Instruments (TI), developed a way of putting an entire circuit on a single chip in the autumn of 1958. The firm announced it in March 1959. However, Robert Noyce (then at Fairchild) had previously sketched out ideas for doing this (which shows the value of keeping an engineering notebook). Kilby's approach used actual wires to connect components on the chip, while Noyce came up with the idea of depositing metal traces onto the chip to connect things, much like the traces on a printed circuit board. Two years later, Fairchild released its first commercial ICs, which were heavily used by the Apollo spacecraft guidance computer. TI sued Fairchild, since it had patented first. However, Fairchild's idea actually worked. Before long, TI was building chips using Fairchild's process.

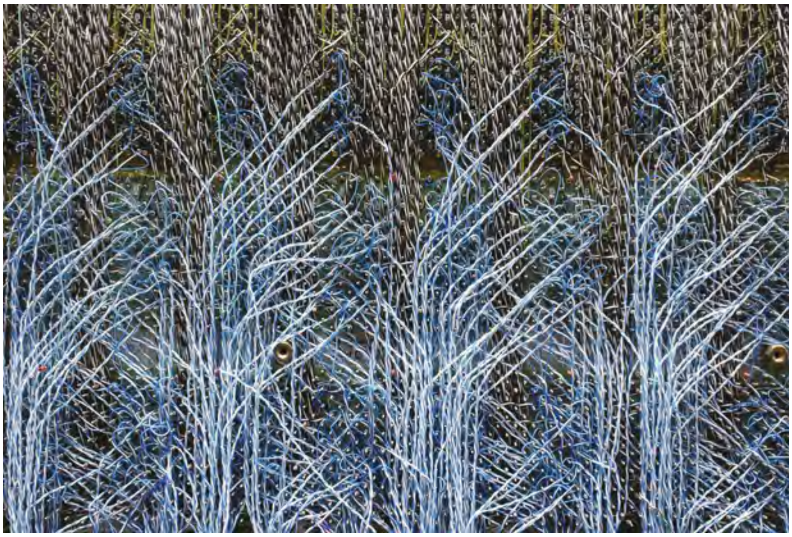

Figure 2
A wire-wrapped backplane that connects many modules

Credit
Dave Fischer CCA-SA 3.0

One of the first uses of integrated circuits was to take logic circuitry, like that shown in **Figure 1**, and put it on a single piece of silicon. **Figure 3** shows a single two-input NAND gate, which is one quarter of the classic 7400 IC. That takes a total of 16 transistors, four diodes, and 16 resistors, which would take up close to 16 cm^2 of circuit board, and puts it on a single piece of silicon that's only several square millimetres. This is shown in **Figure 4**. In order to be useful, this needs to be mounted in a package with leads that can be connected to it. **Figure 5** shows an example.

Logic ICs made several things possible:

1. High-quality mass production of circuits (although that took a while to get right).

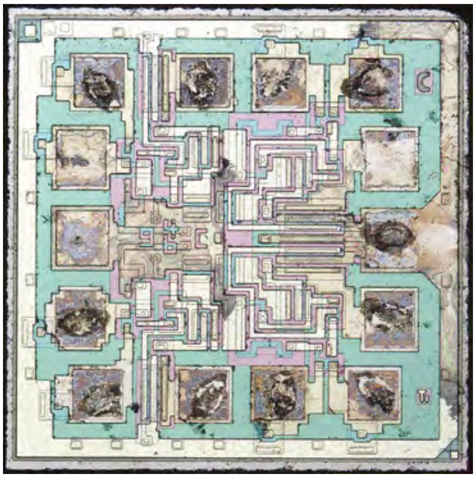

Figure 4
The 54HC00 quad-NAND gate die

Credit
Robert.Baruch CC-BY

2. Along with mass production came lower cost.

3. Taking those discrete components, and placing them on a single chip, drastically reduced the size of circuitry. Replacing vacuum tubes with transistors let a room-sized computer fit inside a half dozen or so refrigerators. Moving to ICs let the same computer fit in a bar-fridge. With each advance, power requirements were reduced as well, which meant less heat was generated. That means less space required for power circuitry and cooling equipment.

Of course, things didn't stop there. We figured out how to make the circuitry on silicon chips smaller, fitting more in the same space. At the same time, we figured out how to reliably make those chips larger.

> **"We figured out how to make the circuitry on silicon chips smaller"**

As progress was made on each front, the power per square millimetre grew. Eventually, the entire CPU of a computer could be fitted onto a single chip: the microprocessor. That changed everything.

SILICON BRAINS

The first microprocessor was the 4-bit 4004, released by Intel in 1971. It was designed for making calculators. Having reprogrammable circuits meant not having to design new circuitry for each →

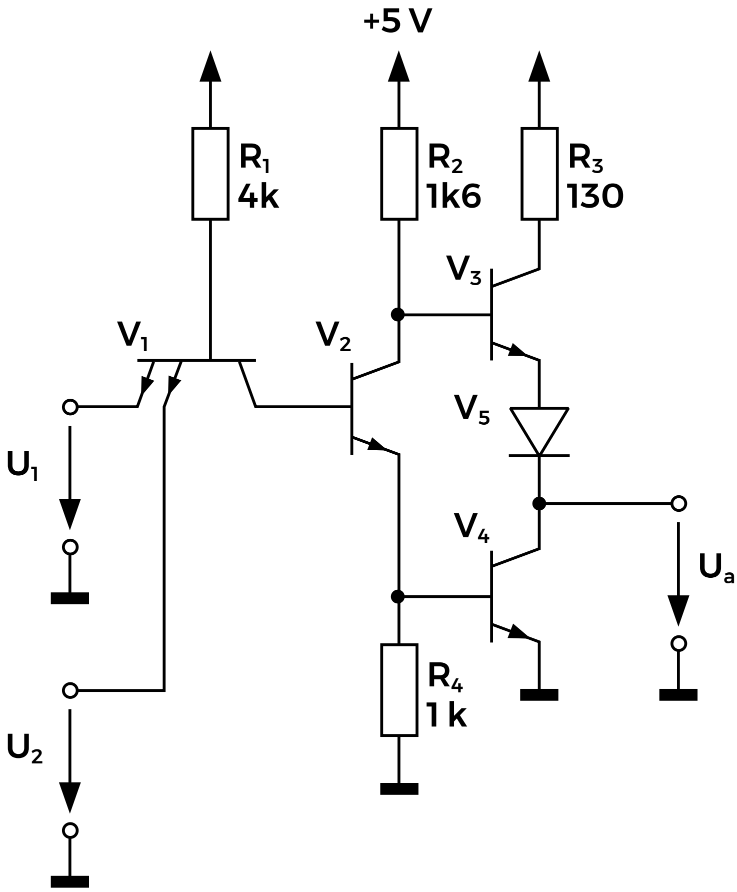

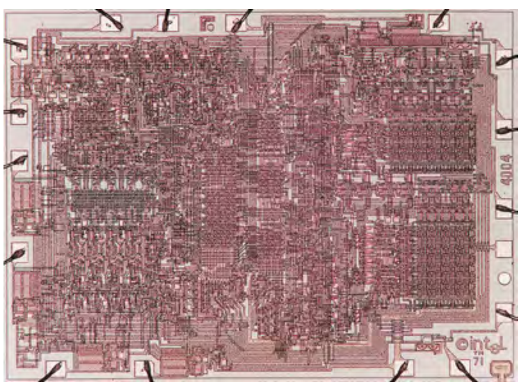

Figure 3
The circuit for a single two-input NAND gate

Figure 5
A chip mounted in a package, and connected to external leads

Figure 6
The 4004 die

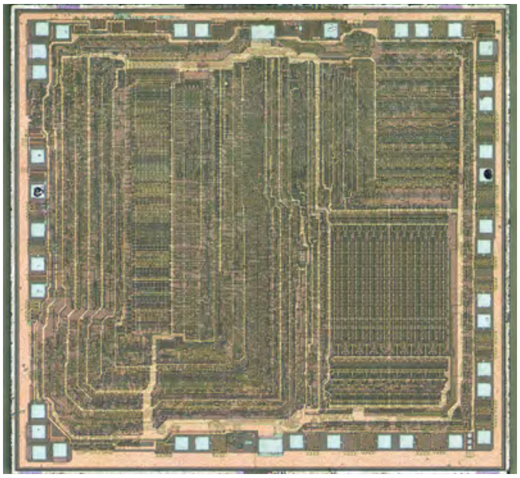

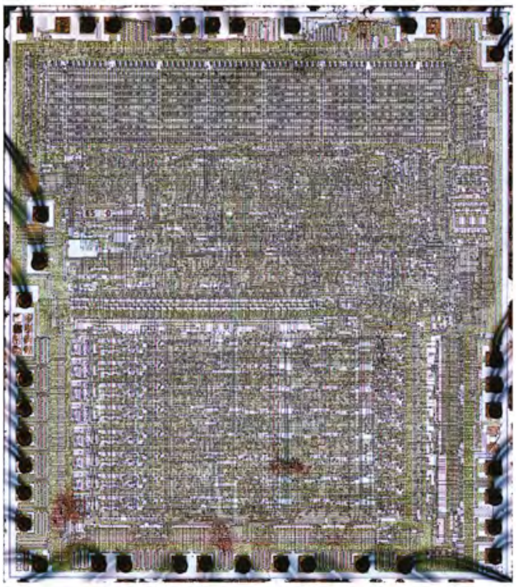

Figure 7
The Z80 die

Credit
Zeptobars CC-BY 3.0

Figure 8
The 6502 die

Figure 9
Comparison of 40- and 64-pin DIP packages

Figure 10
The 68000 die

Credit
Pauli Rautakorpi CC-BY-3.0

new product, you could simply write a new program instead. The 4004 die is shown in **Figure 6**, and contains 2250 transistors in 12 mm^2.

Chips kept growing in density and capability. Eventually, we had several families of 8-bit microprocessors. Intel's own 8080 became very popular. It contained 6000 transistors on a 20 mm^2 die. Another alumnus of Fairchild and Intel (where he led the work on the 4004 and 8080), Federico Faggin started Zilog and developed the Z80, which went on to be incredibly popular. It was used in the Sinclair ZX80, ZX81, and Spectrum computers. The Z80 die is shown in **Figure 7**. It had 8500 transistors in a die that was 18 mm^2.

Another extremely popular microprocessor was the 6502, shown in **Figure 8**. With 3510 transistors on a 21 mm^2 die, it was simple and spacious, in comparison to some. As a result, it was relatively easy to work with, and low in price. It was used in the Apple II line, Commodore's early computers (most famously, the C64), and Acorn's BBC Micro.

Eventually, 8-bit processors gave way to 16-bit ones. An early example is the 68000, famously used in the pre-PowerPC Macintosh models, the Commodore Amiga, and Atari ST line. A version of the 68000, the 68008, was used in the Sinclair QL. This was a much bigger chip: a 64-pin DIP package instead of the 8-bit processors' 40-pin (see **Figure 9**). The die is shown in **Figure 10**. The chip had 68,000 transistors on a 44 mm^2 die.

MODERN DESIGNS

As processors got more complex, moving to 32 and 64 bits, they started bundling more peripheral functions on the chip itself. For example, processors meant for desktop and mobile applications started incorporating

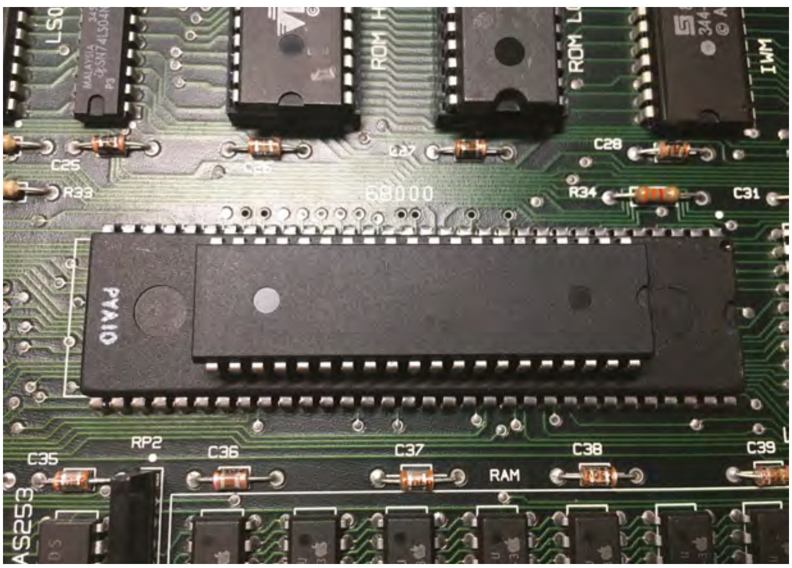

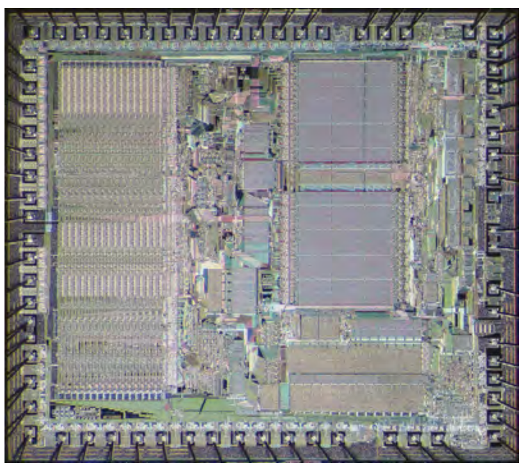

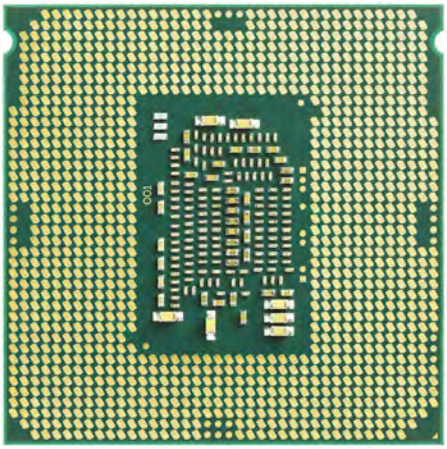

Figure 11
The bottom of an Intel Core i7 Skylake processor

Credit
Eric Gaba CC-BY 3.0

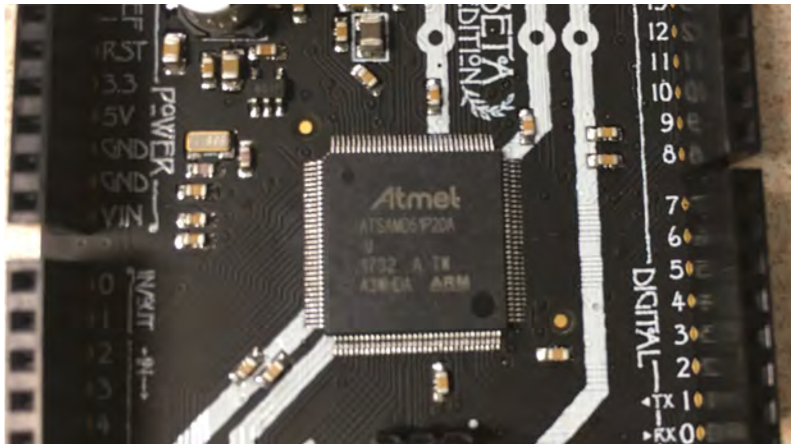

Figure 12
A SAMD51 microcontroller

Figure 13
A pure silicon cylinder ready to be sliced

Credit
Stahlkocher CC-BY 3.0

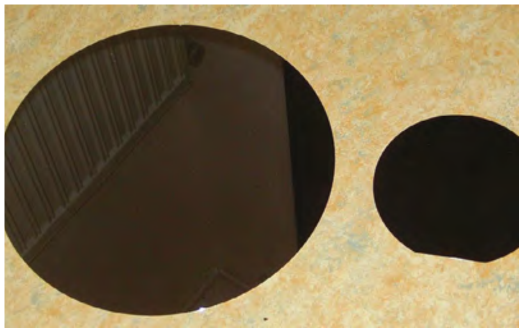

Figure 14
Polished wafers

graphics hardware. Floating point and cryptographic hardware became standard. Processors for embedded applications bundle memory and assorted input/output components (I/O ports, series ports, and so on). All this meant that more external connections were required. The DIP package was no longer adequate. At the same time, printed circuit board production and assembly techniques advanced. Package variety blossomed: square chips with pins on all four sides, grids of pins under the chips (**Figure 11** shows the underside of

> **An integrated circuit starts out as a** cylinder of very pure silicon

an Intel Core i7). Surface-mount technology became standard as well, increasing board density even more as the physical chip packages became smaller. For example, **Figure 12** shows an Atmel SAMD51: a 32-bit ARM CPU core, 1MB of flash memory (for code), 256kB of RAM (for data), and an abundance of interface hardware.

GRAINS OF SAND

An integrated circuit starts out as a cylinder of very pure silicon (see **Figure 13**) which is sliced into thin (no more than a tenth of a millimetre thick) discs. One side of these (where the circuitry will be built) are then polished (**Figure 14**), and coated with an insulating layer of silicon dioxide.

Now, they are ready to have components added. This is done in a similar way to the photographic process of making printed circuit boards. First, a photo-mask is made, which has transparent areas and opaque areas. The disc is coated with a photo-resist, the mask is placed on the disc, and it is exposed to ultraviolet light. This cures the resist where light reaches it. The mask is removed, and the uncured resist is dissolved, leaving a pattern of cured resist on the disc while the rest of the surface is exposed. Depending on what is needed, a small amount of the exposed surface can be etched away chemically, to become N- or P-type silicon, or can have a metallic layer deposited. The cured resist is then removed, and the process continues for the next mark and layer. Insulating layers of silicon dioxide are deposited at times, as required.

A final layer of silicon dioxide is added to seal and protect the circuit. This has contact areas etched away, to which metal contact pads are added which are used to connect to external pins via very fine wires.

This layering process builds a three-dimensional circuit. **Figure 15** (overleaf) shows a visualisation of this structure, while **Figure 16** shows a cross-section of a chip.

AN INTEGRATED CIRCUIT IS A THREE-DIMENSIONAL CIRCUIT

Each slice of silicon contains hundreds of chips, shown in **Figure 17**. Once the above process is complete, the individual chips are tested, and →

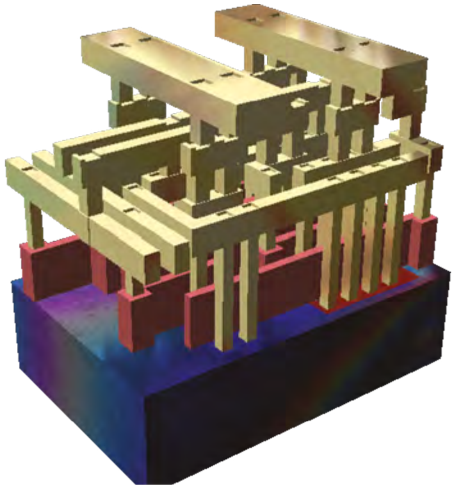

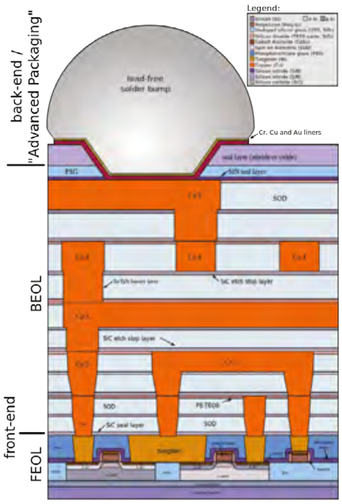

Figure 15 A visualisation of the 3D structure of a portion of an integrated circuit. Silicon is removed for illustrative purposes

Figure 16 Cross-section of an IC

Credit Cepheiden CC-BY 2.5

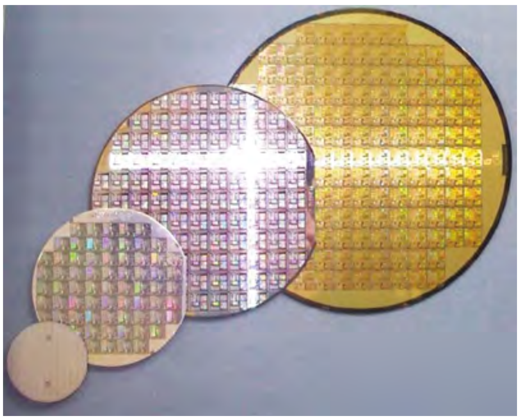

Figure 17 Several wafers, ready to be separated into chips

Credit Guillom CC-BY-3.0

EXPLORING FURTHER

If you want to explore ICs further, check out the blog of Ken Shirriff who delights in exploring (and writing about) the inner workings of chips: **righto.com**.

A great documentary on the early semiconductor industry in the US, including the initial development of the integrated circuit and microprocessor, is the 'Silicon Valley' episode of the PBS series *American Experience*: **hsmag.cc/ouOUPG**.

Micro Men is a British docudrama, showing the early years of Sinclair and Acorn.

non-functional ones are noted, to be discarded later. A diamond cutter is used to score the disc between chips, which are then separated by, essentially, snapping them apart. The non-functional ones are discarded, as are any that were damaged by the separation process.

MORE POWER

In 1965, Gordon Moore (of Shockley Labs, Fairchild Semiconductor, and Intel) observed that the number of transistors in a dense, integrated circuit doubled about every year. In 1975, he revised this to a doubling every two years. See **Figure 18.**

Moore's law is reaching its limit. Moore, himself, commented in 2015 that it would only continue to about 2025. To pack more transistors in the same space, everything needs to be smaller. And therein lies the problem: how small can components be made and still function? As they get smaller and smaller, quantum-level effects can start creeping in and disrupting a transistor's operation. There are many research efforts underway to discover ways to make transistors smaller. Some include: a junctionless transistor, at Tyndall National Institute in Cork, Ireland; a single-electron transistor (that routes single electrons rather than electron flow) at the University of Pittsburgh; and a single-atom transistor from the University of New South Wales. Other research concerns itself with alternatives to silicon, quantum computing, and biological computing.

> **"The number of transistors in a dense IC doubled about every year"**

We've explored the basic ideas of electricity and electronics, basic electronic components, and now we've discussed taking those ideas and making them tiny, enabling entire circuits to become a single component. Next issue, we'll look at a very common and useful type of integrated circuit: the operational amplifier. It's time to get the breadboard out again!

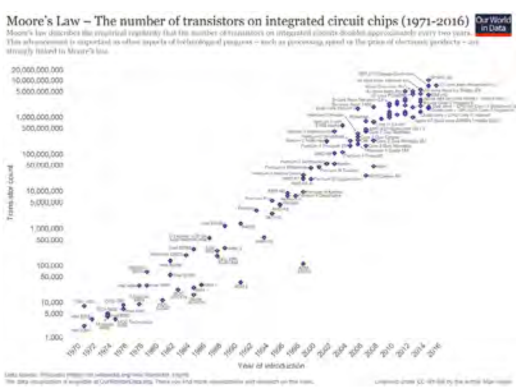

Figure 18 You can explore how chip density has changed over time at hsmag.cc/SCCGGA