

25

MICROCHIPS THAT SHOOK THE WORLD

In microchip design, as in life, small things sometimes add up to big things. Dream up a clever microcircuit, get it sculpted in a sliver of silicon, and your little creation may unleash a technological revolution. It happened with the Intel 8088 microprocessor. And the Mostek MK4096 4-kilobit DRAM. And the Texas Instruments TMS32010 digital signal processor.

Among the many great chips that have emerged from fabs during the half-century reign of the integrated circuit, a small group stands out. Their designs proved so cutting-edge, so out of the box, so ahead of their time, that we are left

groping for more technology clichés to describe them. Suffice it to say that they gave us the technology that made our brief, otherwise tedious existence in this universe worth living.

We've compiled here a list of 25 ICs that we think deserve the best spot on the mantelpiece of the house that Jack Kilby and Robert Noyce built. Some have become enduring objects of worship among the chiperati: the Signetics 555 timer, for example. Others, such as the Fairchild 741 operational amplifier, became textbook design examples. Some, like Microchip Technology's PIC microcontrollers, have sold billions, and are still doing so. A precious few, like Toshiba's flash memory, created whole new markets. And one, at least, became a geeky reference in popular culture. Question: What processor powers Bender, the alcoholic, chain-smoking, morally repre-

hensible robot in "Futurama"? Answer: MOS Technology's 6502.

What these chips have in common is that they're part of the reason why engineers don't get out enough.

Of course, lists like this are nothing if not contentious. Some may accuse us of capricious choices and blatant omissions (and, no, it won't be the first time). Why Intel's 8088 microprocessor and not the 4004 (the first) or the 8080 (the famed)? Where's the radiation-hardened military-grade RCA 1802 processor that was the brains of numerous spacecraft?

If you take only one thing away from this introduction, let it be this: Our list is what remained after weeks of raucous debate between the author, his trusted sources, and several editors of *IEEE Spectrum*. We never intended to compile an exhaustive reckoning of every chip that was a commercial success or a major technical advance. Nor could we include chips that were great but so obscure that only the five engineers who designed them would remember them. We focused on chips that proved unique, intriguing, awe-inspiring. We wanted chips of varied types, from both big and small companies, created long ago or more recently. Above all, we sought ICs that had an impact on the lives of lots of people—chips that became part of earth-shaking gadgets, symbolized technological trends, or simply delighted people.

For each chip, we describe how it came about and why it was innovative, with comments from the engineers and executives who architected it. And because we're not the *IEEE Annals of the History of Computing*, we didn't order the 25 chips chronologically or by type or importance; we arbitrarily scattered them on these pages in a way we think makes for a good read. History is messy, after all.

As a bonus, we asked eminent technologists about their favorite chips. Ever wonder which IC has a special place in the hearts of both Gordon Moore, of Intel, and Morris Chang, founder of Taiwan Semiconductor Manufacturing Company? (Hint: It's a DRAM chip.)

We also want to know what *you* think. Is there a chip whose absence from our list sent you into paroxysms of rage? Take a few deep breaths, have a nice cup of chamomile tea, and then go to <http://www.spectrum.ieee.org/may09/25chips>. There you can drop us a line and check out the runners-up that didn't make the list and more favorite picks by other luminaries.

BY BRIAN R. SANTO


Signetics NE555 Timer (1971)

IT WAS THE SUMMER of 1970, and chip designer Hans Camenzind could tell you a thing or two about Chinese restaurants: His small office was squeezed between two of them in downtown Sunnyvale, Calif. Camenzind was working as a consultant to Signetics, a local semiconductor firm. The economy was tanking. He was making less than US \$15 000 a year and had a wife and four children at home. He really needed to invent something good.

And so he did. One of the greatest chips of all time, in fact. The 555 was a simple IC that could function as a timer or an oscillator. It would become a best seller in analog semiconductors, winding up in kitchen appliances, toys, spacecraft, and a few thousand other things.

“And it almost didn’t get made,” recalls Camenzind, who at 75 is still designing chips, albeit nowhere near a Chinese restaurant.

The idea for the 555 came to him when he was working on a kind of system called a phase-locked loop. With some modifications, the circuit could work as a simple timer: You’d trigger it and it would run for a certain period. Simple as it may sound, there was nothing like that around.



At first, Signetics’s engineering department rejected the idea. The company was already selling components that customers could use to make timers. That could have been the end of it. But Camenzind insisted. He went to Art Fury, Signetics marketing manager. Fury liked it.

Camenzind spent nearly a year testing breadboard prototypes, drawing the circuit components on paper, and cutting sheets of Rubylith—a masking film. “It was all done by hand, no computer,” he says. His final design had 23 transistors, 16 resistors, and 2 diodes.

When the 555 hit the market in 1971, it was a sensation. In 1975 Signetics was absorbed by Philips Semiconductors, now NXP, which says that many billions have been sold. Engineers still use the 555 to create useful electronic modules—as well as less useful things like “Knight Rider”—style lights for car grilles.




Texas Instruments TMC0281 Speech Synthesizer (1978)

IF IT WEREN’T for the TMC0281, E.T. would’ve never been able to “phone home.” That’s because the TMC0281, the first single-chip speech synthesizer, was the heart (or should we say the mouth?) of Texas Instruments’ Speak & Spell

learning toy. In the Steven Spielberg movie, the flat-headed alien uses it to build his interplanetary communicator. (For the record, E.T. also uses a coat hanger, a coffee can, and a circular saw.)

The TMC0281 conveyed voice using a technique called linear predictive coding; the sound came out as a combination of buzzing, hissing, and popping. It was a surprising solution for something deemed “impossible to do in an integrated circuit,” says Gene A. Frantz, one of the four engineers who designed the toy and is still at TI. Variants of the chip were used in Atari arcade games and Chrysler’s K-cars. In 2001, TI sold its speech-synthesis chip line to Sensory, which discontinued it in late 2007. But if you ever need to place a long, very-long-distance phone call, you can find Speak & Spell units in excellent condition on eBay for about US \$50.



MOS Technology 6502 Microprocessor (1975)

WHEN THE CHUBBY-FACED geek stuck that chip on the computer and booted it up, the universe skipped a beat. The geek was Steve Wozniak, the computer was the Apple I, and the chip was the 6502, an 8-bit microprocessor developed by MOS Technology. The chip went on to become the main brains of ridiculously seminal computers like the Apple II, the Commodore PET, and the BBC Micro, not to mention game systems like the Nintendo and Atari. Chuck Peddle, one of the chip’s creators, recalls when they introduced the 6502 at a trade show in 1975. “We had two glass jars filled with chips,” he says, “and I had my wife sit there selling them.” Hordes showed up. The reason: The 6502 wasn’t just faster than its competitors—it was also way cheaper, selling for US \$25 while Intel’s 8080 and Motorola’s 6800 were both fetching nearly \$200.

The breakthrough, says Bill Mensch, who created the 6502 with Peddle, was a minimal instruction set combined with a fabrication process that “yielded 10 times as many good chips as the competition.” The 6502 almost single-handedly forced the price of processors to drop, helping launch the personal computer revolution. Some embedded systems still use the chip. More interesting perhaps, the 6502 is the electronic brain of Bender, the depraved robot in “Futurama,” as revealed in a 1999 episode.

[See “The Truth About Bender’s Brain,” in this issue, where David X. Cohen, the executive producer and head writer for “Futurama,” explains how the choice of the 6502 came about.]




Texas Instruments TMS32010 Digital Signal Processor (1983)

THE BIG STATE OF TEXAS has given us many great things, including the 10-gallon hat, chicken-fried steak, Dr Pepper, and perhaps less prominently, the TMS32010 digital signal processor chip. Created by Texas Instruments, the TMS32010 wasn't the first DSP (that'd be Western Electric's DSP-1, introduced in 1980), but it was surely the fastest. It could compute a multiply operation in 200 nanoseconds, a feat that made engineers all tingly. What's more, it could execute instructions from both on-chip ROM and off-chip RAM, whereas competing chips had only canned DSP functions. "That made program development [for the TMS32010] flexible, just like with microcontrollers and microprocessors," says Wanda Gass, a member of the DSP design team, who is still at TI. At US \$500 apiece, the chip sold about 1000 units the first year. Sales eventually ramped up, and the DSP became part of modems, medical devices, and military systems. Oh, and another application: Worlds of Wonder's Julie, a Chucky-style creepy doll that could sing and talk ("Are we making too much noise?"). The chip was the first in a large DSP family that made—and continues to make—TI's fortune.


Intersil ICL8038 Waveform Generator (circa 1983*)

CRITICS SCOFFED at the ICL8038's limited performance and propensity for behaving erratically. The chip, a generator of sine, square, triangular, sawtooth, and pulse waveforms, was indeed a bit temperamental. But engineers soon learned how to use the chip reliably, and the 8038 became a major hit, eventually selling into the hundreds of millions and finding its way into countless applications—like the famed Moog music synthesizers and the "blue boxes" that "phreakers" used to



Microchip Technology PIC 16C84 Microcontroller (1993)

BACK IN THE EARLY 1990s, the huge 8-bit microcontroller universe belonged to one company, the almighty Motorola. Then along came a small contender with a non-descript name, Microchip Technology. Microchip developed the PIC 16C84, which incorporated a type of memory called EEPROM, for electrically erasable programmable read-only memory. It didn't need UV light to be erased, as did its progenitor, EPROM. "Now users could change their code on the fly," says Rod Drake, the chip's lead designer and now a director at Microchip. Even better, the chip cost less than US \$5, or a quarter the cost of existing alternatives, most of them from, yes, Motorola. The 16C84 found use in smart cards,



remote controls, and wireless car keys. It was the beginning of a line of microcontrollers that became electronics superstars among Fortune 500 companies and weekend hobbyists alike. Some 6 billion have been sold, used in things like industrial controllers, unmanned aerial vehicles, digital pregnancy tests, chip-controlled fireworks, LED jewelry, and a septic-tank monitor named the Turd Alert.

Fairchild Semiconductor µA741 Op-Amp (1968)

OPERATIONAL AMPLIFIERS are the sliced bread of analog design. You can always use some, and you can slap them together with almost anything and get something satisfying. Designers use them to make audio and video preamplifiers, voltage comparators, precision rectifiers, and many other systems that are part of everyday electronics.

In 1963, a 26-year-old engineer named Robert Widlar designed the first monolithic op-amp IC, the µA702, at Fairchild Semiconductor. It sold for US \$300 a pop. Widlar followed up with an improved design, the µA709, cutting the cost to \$70 and making the chip a huge commercial success. The story goes that the freewheeling Widlar asked for a raise. When he didn't get it, he quit. National Semiconductor was

beat the phone companies in the 1980s. The part was so popular the company put out a document titled "Everything You Always Wanted to Know About the ICL8038." Sample question: "Why does connecting pin 7 to pin 8 give the best temperature performance?" Intersil discontinued the 8038 in 2002, but hobbyists still seek it today to make things like homemade function generators and theremins.


* Neither Intersil's PR department nor the company's last engineer working with the part knows the precise introduction date. Do you?

Western Digital WD1402A UART (1971)

GORDON BELL is famous for launching the PDP series of minicomputers at Digital Equipment Corp. in the 1960s. But he also invented a lesser known but no less significant piece of technology: the universal asynchronous receiver/transmitter, or UART. Bell needed some circuitry to connect a Teletype to a PDP-1, a task that required converting parallel signals into serial signals and vice versa. His implementation used some 50 discrete components. Western Digital, a small company making calculator chips, offered to create a single-chip UART. Western Digital founder Al Phillips still remembers when his vice president of engineering showed him the Rubylith sheets with the design, ready for fabrication. “I looked at it for a minute and spotted an open circuit,” Phillips says. “The VP got hysterical.” Western Digital introduced the WD1402A around 1971, and other versions soon followed. Now UARTs are widely used in modems, PC peripherals, and other equipment.

Acorn Computers ARM1 Processor (1985)

IN THE EARLY 1980s, Acorn Computers was a small company with a big product. The firm, based in Cambridge, England, had sold over 1.5 million BBC Micro desktop computers. It was now time to design a new model, and Acorn engineers decided to create their own 32-bit microprocessor. They called it the Acorn RISC Machine, or ARM. The engineers knew it wouldn't be easy; in fact, they half expected they'd encounter an insurmountable design hurdle and have to scrap the whole project. “The team was so small that every design decision had to favor simplicity—or we'd never finish it!” says codesigner Steve Furber, now a computer engineering professor at the University of Manchester. In the end, the simplicity made all the difference. The ARM was small, low power, and easy to program. Sophie Wilson, who designed the instruc-



tion set, still remembers when they first tested the chip on a computer. “We did ‘PRINT PI’ at the prompt, and it gave the right answer,” she says. “We cracked open the bottles of champagne.” In 1990, Acorn spun off its ARM division, and the ARM architecture went on to become the dominant 32-bit embedded processor. More than 10 billion ARM cores have been used in all sorts of gadgetry, including one of Apple's most humiliating flops, the Newton handheld, and one of its most glittering successes, the iPhone.

only too happy to hire a guy who was then helping establish the discipline of analog IC design. In 1967, Widlar created an ever better op-amp for National, the LM101.


While Fairchild managers fretted about the sudden competition, over at the company's R&D lab a recent hire, David Fullagar, scrutinized the LM101. He realized that the chip, however brilliant, had a couple of drawbacks. To avoid certain frequency distortions, engineers had to attach an external capacitor to the chip. What's more, the IC's input stage, the so-called front end, was for some chips overly sensitive to noise, because of quality variations in the semiconductors.

“The front end looked kind of kludgy,” he says.

Fullagar embarked on his own design. He stretched the

limits of semiconductor manufacturing processes at the time, incorporating a 30-picofarad capacitor into the chip. Now, how to improve the front end? The solution was profoundly simple—“it just came to me, I don't know, driving to Tahoe”—and consisted of a couple of extra transistors. That additional circuitry made the amplification smoother and consistent from chip to chip.

Fullagar took his design to the head of R&D at Fairchild, a guy named Gordon Moore, who sent it to the company's commercial division. The new chip, the μ A741, would become *the* standard for op-amps. The IC—and variants created by Fairchild's competitors—have sold in the hundreds of millions. Now, for \$300—the price tag of that primordial 702 op-amp—you can get about a thousand of today's 741 chips.




Kodak KAF-1300 Image Sensor (1986)

LAUNCHED IN 1991, the Kodak DCS 100 digital camera cost as much as US \$13 000 and required a 5-kilogram external data storage unit that users had to carry on a shoulder strap. The sight of a person lugging the contraption? *Not* a Kodak moment. Still, the camera's electronics—housed inside a Nikon F3 body—included one impressive piece of hardware: a thumbnail-size chip that

could capture images at a resolution of 1.3 megapixels, enough for sharp 5-by-7-inch prints. “At the time, 1 megapixel was a magic number,” says Eric Stevens, the chip's lead designer, who still works for Kodak. The chip—a true two-phase charge-coupled device—became the basis for future CCD sensors, helping to kick-start the digital photography revolution. What, by the way, was the very first photo made with the KAF-1300? “Uh,” says Stevens, “we just pointed the sensor at the wall of the laboratory.”



My Favorite **CHIP**



Gordon Moore *cofounder and chairman emeritus of Intel*

There were lots of great chips, but one that will always be dear to me was the **Intel 1103, the first commercial 1024-bit DRAM** [introduced in 1970]. It was the chip that really got Intel over the hump to profitability. It was not the most elegant design, having many of the problems that memory engineers had become familiar with in core memories. But this was comforting to such engineers—it meant that their expertise was not

going to be made obsolete by the new technology. Even today, when I look at my digital watch and see 11:03, I cannot help but remember this key product in Intel's history.

Lee Felsenstein
computer pioneer

The humble, ubiquitous, and quite inexpensive **Signetics 555 timer** had a big impact on my career, when I found myself using it to craft various pulse-sequencing circuits, baud-rate oscillators, and ramp generators for my



earliest consulting clients. With a latch triggered by two comparators, a high-current-drive output and a separate open-collector output for discharging capacitors, the 555 offered a wide range of uses with the addition of a few analog components. It suggested that it could provide more than it delivered at times, which helped me refine my understanding of the limits of a chip and the need to consider its operation in the analog domain.



Jeff Hawkins
founder of Palm and Numenta

My personal favorite—a chip that opened my eyes to what was possible—was the **Intel 2716 EPROM**, vintage late 1970s. The 2716 was nonvolatile, held 2 kilobytes of memory, and, unlike many other chips, used a single 5-volt power supply, so it could be used as storage in a small and practical package. Of course, its big disadvantage was that you needed a UV light to erase it. But with a little imagination you could see that one day it could all be done electrically. In some sense the 2716 is the great-grandparent of today's flash memory chips.




Morris Chang *founder of TSMC*

One of my favorite great microchips is the **Intel 1103, 1-kilobit DRAM**, circa 1970. Reasons: one, huge commercial success; two, started Intel on its way; three, demonstrated the power of MOS technology (versus bipolar); and four, opened up at least another 40 years of life for Moore's Law.

Carver Mead


professor of engineering and applied science, Caltech



My favorite chip contained only a single transistor, although a remarkable one: a **Shottky barrier gate field-effect transistor** made from GaAs (later called the MESFET and now, using more advanced semiconductor structures, the HEMT). I designed it over Thanksgiving break in 1965. The very high mobility of the III-IV materials, together with the absence of minority-carrier storage effects, made these devices far superior for a microwave power-output stage. Despite my efforts to interest American companies, the Japanese were the first to develop these devices. They have made microwave communications in satellites, cellphones, and many other systems possible for many decades.

Sophie Vandebroek

chief technology officer, Xerox



My selection is **Analog Devices' iMEMS accelerometer**, the first commercial chip to significantly integrate MEMS and logic circuitry. Commercialized in the early 1990s, it revolutionized the automotive air-bag industry—and saved lives! Today this type of accelerometer is used in a variety of applications, including the Nintendo Wii and the Apple iPhone. Other types of MEMS chips are more and more present in a broad variety of applications.

Steve Jurvetson

managing director of Draper Fisher Jurvetson



The **Motorola 68000** was a special one. I built a speaking computer with it. I also wrote a multitasker in assembler. That chip was the workhorse of learning back during my M.S.E.E. days. But wait—other chips come to mind. I also like ZettaCore's first 1-megabyte molecular memory chip, D-Wave's solid-state 128-qubit quantum computer processor, the Pentium (I've got an 8-inch wafer signed by Andy Grove in my office), the Canon EOS 5D 12.8-megapixel CMOS sensor, and the iTV 5-bit asynchronous processor running Forth and a custom OS—whew!

Vinod Khosla

cofounder of Sun Microsystems and partner at Kleiner, Perkins, Caufield & Byers



The **Motorola 68010** was my favorite chip. With its power and virtual memory support, it said to the world, “Microprocessors can stand and compete with the big boys—the minis and mainframes.”



David Ditzel

Intel's chief architect for hybrid parallel computing

My favorite chip: the **6502 from MOS Technology**. This 8-bit microprocessor was used by me and many hobbyists for building our own computers [above, me and my homemade PC in 1977]. I wrote a small operating system that fit in 4 kilobytes, and Tiny BASIC from Tom Pittman fit in 2 KB. The reason it was a great chip was that in the 8-bit era, you could use the first 256 bytes of memory as 128 16-bit pointers to index from, making this machine much easier to program than other 8-bit alternatives. It was a big deal when you had to enter your programs in hexadecimal binary form.

IBM Deep Blue 2 Chess Chip (1997)

ON ONE SIDE of the board, 1.5 kilograms of gray matter. On the other side, 480 chess chips. Humans finally fell to computers in 1997, when IBM's chess-playing computer, Deep Blue, beat the reigning world champion, Garry Kasparov. Each of Deep Blue's chips consisted of 1.5 million transistors arranged into specialized blocks—like a move-generator logic array—as well as some RAM and ROM. Together, the chips could churn through 200 million chess positions per second. That brute-force power, combined with clever game-evaluation functions, gave



Deep Blue decisive—Kasparov called them “uncomputerlike”—moves. “They exerted great psychological pressures,” recalls Deep Blue's mastermind, Feng-hsiung Hsu, now at Microsoft.




Transmeta Corp. Crusoe Processor (2000)

WITH GREAT POWER come great heat sinks. And short battery life. And crazy electricity consumption. Hence Transmeta's goal of designing a low-power processor that'd put those hogs offered by Intel and AMD to shame. The plan: Software would translate x86 instructions on the fly into Crusoe's own machine code, whose higher level of parallelism would save time and power. It was hyped as the greatest thing since sliced silicon, and for a while, it was. “Engineering Wizards Conjure Up Processor Gold” was how *IEEE Spectrum's* May 2000 cover put it. Crusoe and its successor, Efficeon, “proved that dynamic binary translation was commercially viable,” says David Ditzel, Transmeta's cofounder, now at Intel. Unfortunately, he adds, the chips arrived several years before the market for low-power computers took off. In the end, while Transmeta did not deliver on its promises, it did force Intel and AMD—through licenses and lawsuits—to chill out.

FROM TOP: ADAM NADEL/AP PHOTO; IEEE SPECTRUM; TEXAS INSTRUMENTS

Texas Instruments Digital Micromirror Device (1987)

ON 18 JUNE 1999, Larry Hornbeck took his wife, Laura, on a date. They went to watch *Star Wars: Episode 1—The Phantom Menace* at a theater in Burbank, Calif. Not that the graying engineer was an avid Jedi fan. The reason they were there was actually the projector. It used a chip—the digital micromirror device—that Hornbeck had invented at Texas Instruments. The chip uses millions of hinged microscopic mirrors to direct light through a projection lens. The screening was “the first digital exhibition of a major motion picture,” says Hornbeck, a TI Fellow. Now movie projectors using this digital light-processing technology—or DLP, as TI branded it—are used in thousands of theaters. It's also used in rear-projection TVs, office projectors, and tiny projectors for cellphones. “To paraphrase Houdini,” Hornbeck says, “micromirrors, gentlemen. The effect is created with micromirrors.”





Intel 8088 Microprocessor (1979)

WAS THERE ANY one chip that propelled Intel into the Fortune 500? Intel says there was: the 8088. This was the 16-bit CPU that IBM chose for its original PC line, which went on to dominate the desktop computer market.

In an odd twist of fate, the chip that established what would become known as the *x86* architecture didn't have a name appended with an "86." The 8088 was basically just a slightly modified 8086, Intel's first 16-bit CPU. Or as Intel engineer Stephen Morse once put it, the 8088 was "a castrated version of the 8086." That's because the new chip's main innovation wasn't exactly a step forward in technical terms: The 8088 processed data in 16-bit words, but it used an 8-bit external data bus.

Intel managers kept the 8088 project under wraps until the 8086 design was mostly complete. "Management didn't want to delay the 8086 by even a day by even telling us they had the 8088 variant in mind," says Peter A. Stoll, a lead engineer for the 8086 project who did some work on the 8088—a "one-day task force to fix a microcode bug that took three days."


It was only after the first functional 8086 came out that Intel shipped the 8086 artwork and documentation to a design unit in Haifa, Israel, where two engineers, Rafi Retter and Dany Star, altered the chip to an 8-bit bus.

The modification proved to be one of Intel's best decisions. The 29 000-transistor 8088 CPU required fewer, less expensive support chips than the 8086 and had "full compatibility with 8-bit hardware, while also providing faster processing and a smooth transition to 16-bit processors," as Intel's Robert Noyce and Ted Hoff wrote in a 1981 article for *IEEE Micro* magazine.

The first PC to use the 8088 was IBM's Model 5150, a monochrome machine that cost US \$3000. Now almost all the world's PCs are built around CPUs that can claim the 8088 as an ancestor. Not bad for a castrated chip.

Micronas Semiconductor MAS3507 MP3 Decoder (1997)

BEFORE THE IPOD, there was the Diamond Rio PMP300. Not that you'd remember. Launched in 1998, the PMP300 became an instant hit, but then the hype faded faster than Milli Vanilli. One thing, though, was notable about the player. It carried the MAS3507 MP3 decoder chip—a RISC-based digital signal processor with an instruction set optimized for audio compression and decompression. The chip, developed by Micronas, let the Rio squeeze about a dozen songs onto its flash memory—laughable today but at the time just enough to compete with portable CD players. Quaint, huh? The Rio and its successors paved the way for the iPod, and now you can carry thousands of songs—and all of Milli Vanilli's albums and music videos—in your pocket.



Mostek MK4096 4-Kilobit DRAM (1973)

MOSTEK WASN'T THE FIRST to put out a DRAM. Intel was. But Mostek's 4-kilobit DRAM chip brought about a key innovation, a circuitry trick called address multiplexing, concocted by Mostek cofounder Bob Proebsting. Basically, the chip used the same pins to access the memory's rows and columns by multiplexing the addressing signals. As a result, the chip wouldn't require more pins as memory density increased and could be made for less money. There was just a little compatibility problem. The 4096 used 16 pins, whereas the memories made by Texas Instruments, Intel, and Motorola had 22 pins.



What followed was one of the most epic face-offs in DRAM history. With Mostek betting its future on the chip, its executives set out to proselytize customers, partners, the press, and even its staff. Fred K. Beckhusen, who as a recent hire was drafted to test the 4096 devices, recalls when Proebsting and chief executive L.J. Sevin came to his night shift to give a seminar—at 2 a.m. "They boldly predicted that in six months no one would hear or care about 22-pin DRAM," Beckhusen says. They were right. The 4096 and its successors became the dominant DRAM for years.

Xilinx XC2064 FPGA (1985)

BACK IN THE EARLY 1980s, chip designers tried to get the most out of each and every transistor on their circuits. But then Ross Freeman had a pretty radical idea. He

came up with a chip packed with transistors that formed loosely organized logic blocks that in turn could be configured and reconfigured with software. Sometimes a bunch

Zilog Z80 Microprocessor (1976)


FEDERICO FAGGIN knew well the kind of money and man-hours it took to market a microprocessor. While at Intel, he had contributed to the designs of two seminal specimens: the primordial 4004, and the 8080, of Altair fame. So when he founded Zilog with former Intel colleague Ralph Ungermann, they decided to start with something simpler: a single-chip microcontroller.

Faggin and Ungermann

Masatoshi Shima, another ex-Intel engineer, worked 80-hour weeks hunched over tables, drawing the Z80 circuits. Faggin soon learned that when it comes to microchips, small is beautiful—but it can hurt your eyes.

“By the end I had to get glasses,” he says. “I became nearsighted.”

The team toiled through 1975 and into 1976. In March of that year, they finally had a prototype chip. The Z80 was a contemporary of MOS Technology’s 6502, and like that chip, it stood out not



rented an office in downtown Los Altos, Calif., drafted a business plan, and went looking for venture capital. They ate lunch at a nearby Safeway supermarket—“Camembert cheese and crackers,” he recalls.

But the engineers soon realized that the microcontroller market was crowded with very good chips. Even if theirs was better than the others, they’d see only slim profits—and continue lunching on cheese and crackers. Zilog had to aim higher on the food chain, so to speak, and the Z80 microprocessor project was born.


The goal was to outperform the 8080 and also offer full compatibility with 8080 software, to lure customers away from Intel. For months, Faggin, Ungermann, and

only for its elegant design but also for being dirt cheap (about US \$25). Still, getting the product out the door took a lot of convincing. “It was just an intense time,” says Faggin, who developed an ulcer as well.

But the sales eventually came through. The Z80 ended up in thousands of products, including the Osborne I (the first portable, or “luggable,” computer), and the Radio Shack TRS-80 and MSX home computers, as well as printers, fax machines, photocopiers, modems, and satellites. Zilog still makes the Z80, which is popular in some embedded systems. In a basic configuration today it costs \$5.73—not even as much as a cheese-and-crackers lunch.

of transistors wouldn’t be used—heresy!—but Freeman was betting that Moore’s Law would eventually make transistors really cheap. It did. To market the chip, called a field-programmable gate array, or FPGA, Freeman cofounded

Xilinx. (Apparently, a weird concept called for a weird company name.) When the company’s first product, the XC2064, came out in 1985, employees were given an assignment: They had to draw by hand an example circuit



Sun Microsystems SPARC Processor (1987)

THERE WAS A TIME, long ago (the early 1980s), when people wore neon-colored leg warmers and watched “Dallas,” and microprocessor architects sought to increase the complexity of CPU instructions as a way of getting more accomplished in each compute cycle. But then a group at the University of California, Berkeley, always a bastion of counterculture, called for the opposite: Simplify the instruction set, they said, and you’ll process instructions at a rate so fast you’ll more than compensate for doing less each cycle. The Berkeley group, led by David Patterson, called their approach RISC, for reduced-instruction-set computing.

As an academic study, RISC sounded great. But was it marketable? Sun Microsystems bet on it. In 1984, a small team of Sun engineers set out to develop a 32-bit RISC processor called SPARC (for Scalable Processor Architecture). The idea was to use the chips in a new line of workstations. One day, Scott McNealy, then Sun’s CEO, showed up at the SPARC development lab. “He said that SPARC would take Sun from a \$500-million-a-year company to a billion-dollar-a-year company,” recalls Patterson, a consultant to the SPARC project.

If that weren’t pressure enough, many outside Sun had expressed doubt the company could pull it off. Worse still, Sun’s marketing team had had a terrifying realization: SPARC spelled backward was...CRAPS! Team members had to swear they would not utter that word to anyone even inside Sun—lest the word get out to arch-rival MIPS Technologies, which was also exploring the RISC concept.


The first version of the minimalist SPARC consisted of a “20 000-gate-array processor without even integer multiply/divide instructions,” says Robert Garner, the lead SPARC architect and now an IBM researcher. Yet, at 10 million instructions per second, it ran about three times as fast as the complex-instruction-set computer (CISC) processors of the day.

Sun would use SPARC to power profitable workstations and servers for years to come. The first SPARC-based product, introduced in 1987, was the Sun-4 line of workstations, which quickly dominated the market and helped propel the company’s revenues past the billion-dollar mark—just as McNealy had prophesied.

using XC2064’s logic blocks, just as Xilinx customers would. Bill Carter, a former chief technology officer, recalls being approached by CEO Bernie Vonderschmitt, who said he was having “a little difficulty doing his homework.” Carter

was only too happy to help the boss. “There we were,” he says, “with paper and colored pencils, working on Bernie’s assignment!” Today FPGAs—sold by Xilinx and others—are used in just too many things to list here. Go reconfigure!

LEFT: CPU-WORLD.COM; RIGHT: ROBERT GARNER




Tripath Technology TA2020 Audio Amplifier (1998)

THERE'S A SUBSET of audiophiles who insist that vacuum tube-based amplifiers produce the best sound and always will. So when some in the audio community claimed that a solid-state class-D amp concocted by a Silicon Valley company called Tripath Technology delivered sound as warm and vibrant as tube amps, it was a big deal. Tripath's trick was to use a 50-megahertz sampling system to drive the amplifier. The company boasted that its TA2020 performed better and cost much less than any comparable solid-state amp. To show off the chip at trade shows, "we'd play that song—that very romantic one from *Titanic*," says Adya Tripathi, Tripath's founder. Like most class-D amps, the 2020 was very power efficient; it didn't require a heat sink and could use a compact package. Tripath's low-end, 15-watt version of the TA2020 sold for US \$3 and was used in boom boxes and ministeereos. Other versions—the most powerful had a 1000-W output—were used in home theaters, high-end audio systems, and TV sets by Sony, Sharp, Toshiba, and others. Eventually, the big semiconductor companies caught up, creating similar chips and sending Tripath into oblivion. Its chips, however, developed a devoted cult following. Audio-amp kits and products based on the TA2020 are still available from such companies as 41 Hz Audio, Sure Electronics, and Winsome Labs.

Amati Communications Overture ADSL Chip Set (1994)

REMEMBER WHEN DSL came along and you chucked that pathetic 56.6-kilobit-per-second modem into the trash? You and the two-thirds of the world's broadband users who use DSL should thank Amati Communications, a start-up out of Stanford University. In the 1990s, it came up with a DSL modulation approach called discrete multi-tone, or DMT. It's basically a way of making one phone line look like hundreds of subchannels and improving transmission using an inverse Robin Hood strategy. "Bits are robbed from the poorest channels and given to the wealthiest channels," says John M. Cioffi, a cofounder of Amati and now an engineering professor at Stanford. DMT beat competing approaches—including ones from giants like AT&T—and became a global standard for DSL. In the mid-1990s, Amati's DSL chip set (one analog, two digital) sold in modest quantities, but by 2000, volume had increased to millions. In the early 2000s, sales exceeded 100 million chips per year. Texas Instruments bought Amati in 1997.




Motorola MC68000 Microprocessor (1979)

MOTOROLA was late to the 16-bit microprocessor party, so it decided to arrive in style. The hybrid 16-bit/32-bit MC68000 packed in 68 000 transistors, more than double the number of Intel's 8086. It had internal 32-bit registers, but a 32-bit bus would have made it prohibitively expensive, so the 68000 used 24-bit address and 16-bit data lines. The 68000 seems to have been the last major processor designed using pencil and paper. "I circulated reduced-size copies of flowcharts, execution-unit resources, decoders, and control logic to other project members," says Nick Tredennick, who designed the 68000's logic. The copies were small and difficult to read, and his bleary-eyed colleagues found a way to make that clear. "One day I came into my office to

find a credit-card-size copy of the flowcharts sitting on my desk," Tredennick recalls. The 68000 found its way into all the early Macintosh computers, as well as the Amiga and the Atari ST. Big sales numbers came from embedded applications in laser printers, arcade games, and industrial controllers. But the 68000 was also the subject of one of history's greatest near misses, right up there with Pete Best losing his spot as a drummer for the Beatles. IBM wanted to use the 68000 in its PC line, but the company went with Intel's 8088 because, among other things, the 68000 was still relatively scarce. As one observer later reflected, had Motorola prevailed, the Windows-Intel duopoly known as Wintel might have been Winola instead.

Chips & Technologies AT Chip Set (1985)

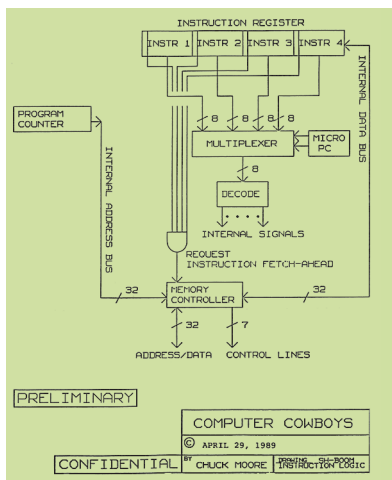
BY 1984, when IBM introduced its 80286 AT line of PCs, the company was already emerging as the clear winner in desktop computers—and it intended to maintain its dominance. But Big Blue's plans were foiled by a tiny company called Chips & Technologies, in San Jose, Calif. C&T developed five chips that duplicated the functionality of the AT motherboard, which used some 100 chips. To make sure the chip set was compatible with the IBM PC, the C&T engineers figured there was just one thing to do. "We had



Computer Cowboys Sh-Boom Processor (1988)

TWO CHIP DESIGNERS walk into a bar. They are Russell H. Fish III and Chuck H. Moore, and the bar is called Sh-Boom. No, this is not the beginning of a joke. It's actually part of a technology tale filled with discord and lawsuits, lots of lawsuits. It all started in 1988 when Fish and Moore created a bizarre processor called Sh-Boom. The chip was so streamlined it could run faster than the clock on the circuit board that drove the rest of the computer. So the two designers found a way to have the processor run its own superfast internal clock while still staying synchronized with the rest of the computer. Sh-Boom was never a commercial success, and after patenting its innovative parts, Moore and Fish moved on. Fish later sold his patent rights to a Carlsbad, Calif.-based firm, Patriot Scientific, which remained a profitless speck of a company until its executives had a revelation: In the years since Sh-Boom's invention, the speed of processors had by far surpassed that of motherboards, and so practically every maker of computers and consumer electronics wound up

using a solution just like the one Fish and Moore had patented. *Ka-ching!* Patriot fired a barrage of lawsuits against U.S. and Japanese companies. Whether these companies' chips depend on the Sh-Boom ideas is a matter of controversy. But since 2006, Patriot and Moore have reaped over US \$125 million in licensing fees from Intel, AMD, Sony, Olympus, and others. As for the name Sh-Boom, Moore, now at IntellaSys, in Cupertino, Calif., says: "It supposedly derived from the name of a bar where Fish and I drank bourbon and scribbled on napkins. There's little truth in that. But I did like the name he suggested."



LEFT: CHUCK MOORE; RIGHT: FUJIO MASUOKA



the nerve-racking but admittedly entertaining task of playing games for weeks," says Ravi Bhatnagar, the chip-set lead designer and now a vice president at Altierre Corp., in San Jose, Calif. The C&T chips enabled manufacturers like Taiwan's Acer to make cheaper PCs and launch the invasion of the PC clones. Intel bought C&T in 1997.

Toshiba NAND Flash Memory (1989)

THE SAGA THAT is the invention of flash memory began when a Toshiba factory manager named Fujio Masuoka decided he'd reinvent semiconductor memory. We'll get to that in a minute. First, a bit (groan) of history is in order.

Before flash memory came along, the only way to store what then passed for large amounts of data was to use magnetic tapes, floppy disks, and hard disks. Many companies were trying to create solid-state alternatives, but the choices, such as EPROM (or erasable programmable read-only memory, which required ultraviolet light to erase the data) and EEPROM (the extra E stands for "electrically," doing away with the UV) couldn't store much data economically.

Enter Masuoka-san at Toshiba. In 1980, he recruited four engineers to a semisecret project aimed at designing a memory chip that could store lots of data and would be affordable. Their strategy was simple. "We knew the cost of the chip would keep going down as long as transistors shrank in size," says Masuoka, now CTO of Unisantis Electronics, in Tokyo.

Masuoka's team came up with a variation of EEPROM that featured a memory cell consisting of a single transistor. At the time, conventional EEPROM needed two transistors per cell. It was a seemingly small difference that had a huge impact on cost.

In search of a catchy name, they settled on "flash" because of the chip's ultrafast erasing capability. Now, if you're thinking Toshiba rushed the


invention into production and watched as the money poured in, you don't know much about how huge corporations typically exploit internal innovations. As it turned out, Masuoka's bosses at Toshiba told him to, well, erase the idea.

He didn't, of course. In 1984 he presented a paper on his memory design at the IEEE International Electron Devices Meeting, in San Francisco. That prompted Intel to begin development of a type of flash memory based on NOR logic gates. In 1988, the company introduced a 256-kilobit chip that found use in vehicles, computers, and other mass-market items, creating a nice new business for Intel.

That's all it took for Toshiba to finally decide to market Masuoka's invention. His flash chip was based on NAND technology, which offered greater storage densities but proved trickier to manufacture. Success came in 1989, when Toshiba's first NAND flash hit the market. And just as Masuoka had predicted, prices kept falling.

Digital photography gave flash a big boost in the late 1990s, and Toshiba became one of the biggest players in a multibillion-dollar market. At the same time, though, Masuoka's relationship with other executives soured, and he quit Toshiba. (He later sued for a share of the vast profits and won a cash payment.)

Now NAND flash is a key piece of every gadget—cellphones, cameras, music players, and of course, the USB drives that techies love to wear around their necks. "Mine has 4 gigabytes," Masuoka says. □



With additional reporting by Sally Adee, Erico Guizzo, and Samuel K. Moore