



Part 2 Chips (IC's & MPU's) by Dr Jeff Drobman Dr Jeff Software

Index



Part 2: Chips

History of Silicon Valley & Chips Microprocessors (MPU/MCU) □ Microprocessor Timeline (Exhaustive) Early RISC MPU's AMD 29K, Intel i960, MIPS R2/3/4000 Advanced RISC MPU's Apple, Intel Core, AMD Zen, Mobile SoC Chips & Wafer Fabs (see separate file) Moore's Law Memories Logic (& Bit-Slice MPU) Debug/Test, JTAG



Chips



Transistors

The Transistor







size = ~1 inch

1947 ushered in the era of *Microelectronics*

A transistor is a semiconductor device used to amplify or switch electronic signals and electrical power. It is composed of semiconductor material usually with at least three terminals for connection to an external circuit. A voltage or current applied to one pair of the transistor's terminal



1947- Bipolar point/junction
1959- Planar bipolar [10]*

*1964- MOS (P-channel) [100]

✤ 1972- MOS (N-channel) [1,000]

✤ 1978- CMOS [4,000]

✤ 1990- sub-micron [10,000]

*2000-100 nm [100,000]

- ✤ 2011- FinFET [1,000,000]
- 2022- 5nm [50,000,000,000] *no. of transistors

Transistors have been shrunk every 2 years according to *Moore's Law*



• yields \rightarrow ~1M devices per cm²

Viewing Transistors



How are billions of transistors compressed into a single chip? Can a transistor in the chip be seen with a microscope?



Quora

Jeff Drobman

Works at Dr Jeff Software \cdot Just now \cdot (\$)

once a chip is complete, only the top few layers of metal interconnect are visible (unless etched away). a single transistor is way too small to be readily identified, as they are now as small as about 25–40 nm in overall size. note that a "5nm" node means that only the channel length is that small. the overall transistor size, and pitch, is more like 25–40 nm.

Transistor Atoms





Al Kordesch, Semiconductor Device Modeling Answered Feb 1, 2019

How many atoms are in a typical transistor in a chip? Short Answer: 49,000 atoms!

Apple's iPhone XS uses 7 nanometer transistors. So let's estimate how many atoms are in one of them. Excluding the connecting wires and other parts, I'm just going to calculate the size of the active part, the "channel" under the gate. The volume of the channel is about (7 nm long) x (7 nm deep) x (20 nm wide). The atomic density of silicon is 5E+28 atoms per cubic meter. So let's go!

Number of atoms n = volume x density

n = (980E-27) x (5E+28) = 49,000 atoms.

Silicon atomic radius = .111nm \rightarrow 4.5 atoms/nm

Cubic: 4.5³ = **91** atoms/cu nm

Channel volume @5n: 5x5x18 = **450** cu nm

Cubic: 91*450 = **40,950** atoms/channel

Bipolar Transistors







Bipolar/MOSFET Transistors



- Current flows opposite electrons (C->E, S->D)
- ✤ B, G are inputs (H/L)
- B, G voltages turn transistor ON/OFF
- Outputs (not shown) are tied to C, D



Chips



Chip History

Levels of Integration







1st Flip-Flop Chips







Fairchild Semiconductor 1961 Micrologic Family announcement brochure

1st Flip-Flop Chips





DR JEFF SOFTWARE Jeff Drobman ©2016-23

The first Fairchild integrated circuit contained four transistors. Photo: Fairchild Semiconductor

Old Computer Tech



Discretes R/T

This is a circuit board from an IBM 7040, built in 1963.



There are no small components there. Those horizontal cylinders are maybe half an inch (12mm) long.

Discrete Transistors

Old Computer Tech



MSI/LSI IC's

This is a serial interface board for a DEC PDP-11



The integrated circuits were very simple with only a few dozen transistors in, so they didn't need very precise machines to make them. The boards could still be assembled by hand.

Old Computer Tech



LSI IC's

This is the second computer I assembled, the Acorn Atom



New Computer Tech



VLSI IC's

12010

J10 ETHERNET E GP10 💊 Rospberry Pi. 4 Model B C Rospberry Pi 2018 10 m H information for a GLOBAL EN RUN 0001 FCC ID: 2ABCB-RPI4B 20953-RP148 Made in the UK

IBM PC Chips





IBM PC – Large Motherboard Grobman



The memory is organised in four banks in the bottom right corner of the motherboard – in this case there are four 64KB banks, adding up to a total of 256KB

1947-68 Early Semiconductors & IC'S Jeff Drobman ©2016-23

MILESTONES

trans	1947	*	Transistor invented (Bell Labs' Bardeen, Brattain, Shockley) - point contact form
	1951	*	Bipolar junction transistor (BJT) invented by Wm Shockley
	1956	*	Shockley Semiconductor Laboratory founded as a division of Beckman Instruments
			Shockley hires his PhD students Robert Noyce, Gordon Moore, et al.
Fairchil	d 195	7 🔆	"Traitorous 8" leave Shockley Labs, found Fairchild Semiconductor
	Mar	*	TI– Jack Kilby tests the world's first integrated circuit (SgI-transistor oscillator on germanium)
	1959	***	Jean Hoerni of Fairchild demos his "planar process" (world first)
			Bob Noyce documents a method for building ICs using that planar process
May	1960	*	Fairchild group makes first <u>IC</u>
		.*.	[Courts and the tech community decided to give equal rights to the invention of the IC to both Kilby and Noyce] MOS (linear) invented: first MOSEET amplifier domonstrated
MOS	1963	•••	standard logic families are introduced using DTL and TTL structures
		•••	CMOS process was invented by Eairchild Semiconductor in a 1963 paper and patent
	1064	**	MOS (digital) – 1 st products released by General Microelectronics for a calculator chinset
	1904	**	Linear IC's – 1 st analog ICs introduced by Fairchild Semiconductor
Apr	1965	**	Moore's Law born - Gordon Moore publishes his first version
		**	CMOS– 1 st parts by RCA
		*	ROM– 1 st Semiconductor
		**	DIP packages
	1966	*	RAM Bipolar RAMs (SRAM) introduced
RAM		**	DRAM– IBM conceives DRAM cell (1T, 1C)
	1968	**	CMOS SRAM– 1 st parts by RCA



2020 ARM intro's "backside power" process Mfg -- process

Silicon Valley





Portion of Silicon Valley map. drawn by Maryanne Regal Hoburg (1982). Courtesy: The David Rumsey Map Center, Stanford University Library

Silicon Valley









Founding Fathers









https://computerhistory.org/blog/beckman-shockleyand-the-60th-anniversary-of-the-birth-of-silicon-valley

CHM BLOG CURATORIAL INSIGHTS , REMARKABLE PEOPLE

BECKMAN, Shockley and the 60th anniversary of the birth of Silicon Valley

By David Laws | February 10, 2016



Museum CHM on Shockley



None would have the same lasting impact on the fortunes of the future Silicon Valley and beyond as Dr. Arnold Beckman's disclosure of an agreement signed the previous day for "the establishment in the Stanford community of the Shockley Semiconductor Laboratory to develop and produce transistors and other semiconductor devices."

ABOUT THE AUTHOR

David A. Laws [AMD 1975-1986, V.P. Business Development] is a high-technology business consultant with a focus on marketing and strategic planning. He earned a B.Sc. (Physics) in the UK and after moving to California in 1968 worked for Silicon Valley companies, including Fairchild Semiconductor, Advanced Micro Devices (AMD), and Altera Corporation, in roles from product marketing engineer to CEO.



My Genesis Article



Genesis: A Silicon Valley Tale



TECH HISTORY ARTICLE

BY DR JEFF DROBMAN

Highlights

- Fairchild founding
- Intel founding
- ✤ AMD history
- ✤ AMD Intel rivalry
- Search for CMOS
- * RISC CPU Architecture
- Legendary Parties & Conferences
- Anecdotes
- Valley Significant Others
- ✤ Genesis org-chart
- Process Technology Evolution
- * Anniversaries of Technologies







My Genesis Article



The Legend

It has long been legendary that companies in Silicon Valley got started in garages and beach houses, and I am setting the record straight: *It is true*. **Apple** was started in Steve Wozniak's garage, when friend Steve Jobs came by and saw his hobby computer. Advanced Micro Devices (AMD) got its start in founding president Jerry Sanders' rented Malibu beach house, on a chilly December evening in 1968 - though the house was heated considerably by those entrepreneurial fires. AMD was incorporated 5 months later (May 1969).







https://www.eetimes.com/the-new-silicon-frontier-chapter-4-startup-

fever-and-venture-capital/

DESIGNLINES | EE LIFE

The New Silicon Frontier Chapter 4: Startup Fever and Venture Capital

MELTING POT FOR THE FAIRCHILDREN

Sheldon Roberts, Eugene Kleiner, and Jean Hoerni's collective decision to leave and compete against Fairchild, just over three years after the company was founded, was the first of what would be many subsequent defections and spinouts, eventually known as "Fairchildren," directly or indirectly creating dozens of corporations, including Intel and AMD. In doing so, Fairchild sowed the seeds of innovation across multiple companies in the region that would eventually become known as Silicon Valley.

While it is unclear who came up with the moniker, "Silicon Valley," Don Hoefler, a technology reporter for the industry publication *Electronic News*, is often credited with popularizing the name in a 1971 column about the region's chip industry. Hoefler also promoted the area's innovative qualities, and was one of the first writers to chronicle the Northern Californian technology industry as a community.

Don Hoefler





Local watering holes, restaurants and other hot spots provided venues for Silicon Valley's "work hard, play hard" ethos, where industry folk gathered after work to drink, gossip, brag, trade war stories, talk shop, exchange ideas, change jobs and develop new contacts. Key venues included the Wagon Wheel, Lion & Compass, and Ricky's, along with the Peppermill and the Sunnyvale Hilton.





THE FAIRCHILD LEGACY

Throughout the first half of the 1960s, Fairchild was the undisputed semiconductor leader, forging ahead across all industry segments, be it design, technology, production or sales. Early sales and marketing efforts were modest and military-oriented; that changed in 1961 when Robert Noyce and Tom Bay recruited a group of aggressive salesmen and marketing specialists, including Jerry Sanders III and Floyd Kvamme. The newcomers transformed Fairchild's sales and marketing departments into one of the industry's legends.

Among the pivotal moments was Fairchild's entry into the consumer TV market. Attracted by potential high volumes, Sanders wanted to replace the tube (valve) CRT driver with a transistor, but the target price was U.S. \$1.50. Transistors at that time were selling to the military for \$150.00. In what can only be regarded as a massive leap of faith, Noyce's instructions to Sanders were, "Go take the order, Jerry. We'll figure out how to do it later. Maybe we'll have to build it in Hong Kong and put it in plastic, but right now let's just do it."







The TTL Data Book for Design Engineers. By always ensuring any bill of materials included at least one TTL part that was only available from it, Texas Instruments was able to stay one step ahead of the competition and own the T'TL market for the best part of 30 years, until standard logic eventually fell victim to the 1980s applicationspecific IC revolution.





Charles Sporck, Noyce's operations manager often credited with running the industry's tightest ship, left in early 1968 along with Pierre Lamond to join Widlar and Talbert at National Semiconductor. That triggered Noyce and Moore's departure from the firm later that same year–a pivotal moment in the eventual demise of the firm. The collective exodus of Sporck, Noyce, and Moore, along with so many other executives, signaled the end of an era, prompting Sherman Fairchild to bring in a new management team, led by C. Lester Hogan, then vice president of Motorola Semiconductor.

Sporck \rightarrow National

HOGAN'S HEROES

Hogan's arrival, and the subsequent displacement of Fairchild managers, demoralized the firm even further, prompting a further exodus of employees who would launch a host of new companies. Leading a group dubbed "Hogan's Heroes," the ultra-conservative Motorola executives immediately clashed with Sanders, Fairchild's flamboyant sales chief.

Hogan/Wilf/Sanders




While initially slow to respond to the changing market under Sander's direction, Fairchild embarked on a strategy of leapfrogging Texas Instruments by focusing on more complex large scale, 30-plus gate parts, instead of simpler small and medium scale devices under 30 gates — a strategy that was proving popular and successful with engineers. The move forced Texas Instruments to recognize the threat and copy all of Fairchild's 9300 series parts under 74 series numbers (for example the 9300 became the 74195 and the 9341 the 74181.)

Sander's entire strategy collapsed, however, when Hogan capitulated to Ken Olsen, founder and CEO of Digital Equipment Corporation and a key Fairchild customer. Olsen wanted Fairchild to give up on its proprietary TTL technology and instead second-source Texas Instruments' 74 Series TTL. Against Sanders' wishes, Hogan agreed, signing the death warrant for Fairchild's TTL strategy. Sanders was, understandably livid. *"*You've just killed the company, Ken," Sander's fumed.

Hogan's betrayal was the last straw for Sanders. He, together with a group of Fairchild engineers, quit to start Advanced Micro Devices. With Sanders installed as president, one of his first moves was to establish the mantra: "People first, revenues and profits will follow." Sanders also gave every employee stock options in the new company, an innovation at the time.

Bell Labs

Founders HoF





Wm Shockley



AMD cofounder

founded Maxim Integrated

Jerry Sanders CEO, AMD

1969-2002

From left: W. Jerry Sanders III, President and Chairman of the Board. D. John Carey, Managing Director of Complex Digital Operations. Svent. E. Simonean, Director of Engineering, Complex Digital Operations, Frank T, Botte, Director of Development, Analog Operations, Jamer R, Giles, Director al Engineering, Analog Operations, Edinal J, Turney, Director of Sales and Administration. Jack F, Gilford, Director el Marketing and Business Development, R. Lavrence Stonger, Managing Director, Analog Operations.

Products





Fairchild founders (8)



Fairchild Chairman/CEO, LSI Logic founder

Wilf Corrigan



Bob Noyce



Gordon Moore

T. J. Rodgers



Cypress Semi founder

Intel



Intel Originals



L to R: Andy Grove Bob Noyce Gordon Moore

> Founders: Bob Noyce Gordon Moore

1968

Intel CEO Gordon Moore





Screenshot of Gordon Moore featured in Scientists You Must Know by the Science History Institute. Courtesy of the Science History Institute.

Intel CEO Andy Grove



...



Jeff Drobman 1 min · ♣ ▼

Andy Grove was Intel's feisty CEO successor to Bob Noyce, hence 2nd one. Andy reigned over Intel in the 1980's. Andy battled AMD's CEO Jerry Sanders over the rights to the i80386 chip design awarded by the legendary 2nd source contract -- culminating in a \$1B lawsuit by AMD.



Source: VentureBeat

<u>"Bad Companies Are Destroyed</u> by Crises ... Great Companies Are Improved by Them"



AMD



manager

Chuck Keough mid-america area sales manager Chicago (312) 297-4115/6

Steve Marks eastern area sales Ed Turney director of sales coordinates all New York (212) 343-2220/1 sales activities Sunnyvale

Steve Zelencik western area sales manager Los Angeles (213) 360-2102/3

05.

Dr Jeff

Jeff Drobman ©2016-23

DR JEFF

SOFTWARE

INDIE APP DEVELOPER

1969-2014 45 YEARS OF

INSPIRED TECHNOLOGY.

AMD Sunnyvale groundbreaking, May 1969.

From left: W. Jerry Sanders III, President and Chairman of the Board. D. John Carey, Managing Director of Complex Digital Operations, Sven E. Simonsen, Director of Engineering, Complex Digital Operations, Frank T. Botte, Director of Development, Analog Operations, James N. Giles, Director of Engineering, Analog Operations, Edwin J. Turney, Director of Sales and Administration. Jack F, Gifford, Director of Marketing and Business Development, R. Lawrence Stenger, Managing Director, Analog Operations.





Jerry Sanders president coordination & implementation of AMD goals/ objectives Sunnyvale

Jolene Trout customer service, delivery, scheduling Sunnyvale

Shel Schumaker digital marketing/ headquarter sales, DIC pricing, specs, new product, coordination of distributor & international activities Sunnyvale





and

Motorola was a major chip company, having pioneered in	Fate	Demerged into Motorola Mobility an Motorola Solutions in 2011
digital logic and microprocessors such as the 6800/68000 and	Successors	Motorola Mobility
PPC. but Motorola no longer exists. it was first split into 2		Motorola Solutions NXP Semiconductors
companies: Motorola Solutions and Motorola Mobility, in 2011		ON Semiconductor
(sold to Google in 2012, then to Lenovo in 2014). that was		Cambium Networks
after the chip business was split up: ON Semi in 1999, and then	Founded	September 25, 1928; 92 years ago
Freescale in 2004, which was then sold to NIXP Semi (Philins) in	Founders	Paul and Joseph Galvin
and the second the solution of the solution of the second to the second	Defunct	January 4, 2011; 10 years ago
2015.		

from Wikipedia:

Motorola, Inc. (/<u>mootə'roolə/[2]</u>) was an American <u>multinational telecommunications</u> company founded on September 25, 1928, based in <u>Schaumburg, Illinois</u>. After having lost \$4.3 billion from 2007 to 2009, the company demerged into two independent public companies, <u>Motorola</u> <u>Mobility</u> and <u>Motorola Solutions</u> on January 4, 2011.[3] Motorola Inc. was renamed Motorola Solutions and is legally the direct successor to the original company after the demerger from Motorola Mobility.[4][5] Motorola Mobility was sold to <u>Google</u> in 2012, and acquired by <u>Chinese</u> company <u>Lenovo</u> in 2014

Electronic Device Cos.

(2



1972

	Nationwide Distributor for the world's finest lines of quality electronic components!
HAMILTON,	HAMILTON ELECTRO SALES
3) 870-7171	10912 W. Washington Blvd. • Culver City, California 90230
RANGE COUNTY (714) 522-8200 ZENITH 2-1474	ADVANCED MICRO DEVICES Integrated Circuits
	BOURNSPotentiometers
	BURROUGHSNixie Tubes, Digital Displays
	FAIRCHILDSemiconductors, Integrated Circuits, Microwave & Opto Devices
	GENERAL ELECTRIC Semiconductors, Capacitors, Meters, Opto Devices, Relays, Volt Pacs
	INTELIntegrated Circuits
	KEMETSolid Tantalum, Ceramic Capacitors
	LITRONIXOpto Devices
	MEPCO/ELECTRAPotentiometers, Resistors
	MONSANTOOpto Devices
	MOTOROLA Semiconductors, Integrated Circuits, Opto Devices
STOCKING INVENTORY	NATIONAL SEMICONDUCTOR Semiconductors, Integrated Circuits, Opto Devices
	POTTER & BRUMFIELD – WOOD Relays, Switches, Circuit Breakers
IN EVERY LOCATION	RCASemiconductors, Integrated Circuits, Opto Devices
	SIGNETICS Integrated Circuits
	SILICONIXSemiconductors, Integrated Circuits
	WESTINGHOUSESemiconductors

Chips on a PCB





Dev Boards



POPULAR DEVELOPMENT BOARDS

ultimate design flexibility





For engineers designing compact smart solutions for the IoT market, this FPGA IoT Maker Board is an excellent tool to speed up the development process and enter the market with a highperformance, reliable product.



Google Coral Dev Board

The Coral Dev Board is now in-stock and available for free 1-day shipping at Arrow.com. Prototype, scale, and deploy with more flexibility using the Coral Dev Board and accessories with Google.

ARM-A72 Development Boar Configure



Raspberry Pi 4 Your tiny, dual-display, desktop computer

DR JEFF



Chips



Microprocessors MPU/MCU

Early MPUs







TRANSISTORS, the electronic amplifiers and switches found at the heart of everything from pocket radios to warehouse-size supercomputers, were invented in 1947. Early devices were of a type called bipolar transistors, which are still in use. By the 1960s, engineers had figured out how to combine multiple bipolar transistors into single integrated circuits. But because of the complex structure of these transistors, an integrated circuit could contain only a small number of them. So although a minicomputer built from bipolar integrated circuits was much smaller than earlier computers, it still required multiple boards with hundreds of chips. ¶ In 1960, a new type of transistor was demonstrated: the metal-oxide-semiconductor (MOS) transistor. At first this technology wasn't all that promising. These transistors were slower, less reliable, and more expensive than their bipolar counterparts. But by 1964, integrated circuits based on MOS transistors boasted higher densities and lower manufacturing costs than those of the bipolar competition. Integrated circuits continued to increase in complexity, as described by Moore's Law, but now MOS technology took the lead.



Intel i4004 and i1103A

World's 1st MPU & DRAM



The **Intel 4004** is a 4-bit central processing unit (CPU) released by Intel Corporation in 1971. It was the first commercially available microprocessor, and the first in a long line of Intel CPUs. The chip design, implemented with the MOS silicon gate technology, started in April 1970, and was created by Federi





The **1103** is a dynamic random-access memory (DRAM) integrated circuit (IC) developed and fabricated by Intel. Introduced in October 1970, the 1103 was the first commercially available DRAM IC; and due to its small physical size and low price relative to magnetic-core memory, it replaced the la

> 1st **DRAM** 1K x1 1970

Intel i8008



World's 1st 8-bit MPU



The venerable 16-pin side-brazed DIP. (Click image to view full size)





Intel i4004/i8008



World's 1st 4/8-bit MPU's



Intel MCS-4 and MCS-8 design team and CPU chips

Intel i4004/i8008



World's 1st 4/8-bit MPU's



1 mm

EVERYTHING'S BIGGER IN TEXAS: Although Texas Instruments' TMX 1795 and Intel's 8008 had a similar number of transistors, the former required a much larger silicon die. Indeed, the TMX 1795 was larger than the Intel 8008 and 4004 combined. Intel's engineers believed that its large size made the TI chip impractical to produce in commercial quantities, but TI's very successful TMS 0100 calculator chip, introduced at about the same time, had an even larger die. So the connection between die size and commercial viability must not have been straightforward.

3,078 transistors



4004 2,300 transistors



8008 3,098 transistors

8-bit i8080



Wikipedia



018928-2

Venerable Zilog Z80





- Improved i8080 (dual register bank, for one)
- Spinoff of Intel (Federico Faggin, Bernard Peuto, et al.

AMD FPU



Am9511/12 From AMD datasheet

Single & Double precision Floating-point with transcendentals

Floating Point Processor Manual Am9511A/Am9512

AMD FPU



Am9511/12 From AMD datasheet

Single & Double precision **Floating-point** with *transcendentals*

5.2 Am9511A ARITHMETIC PROCESSOR

This pioneer single-chip arithmetic processor interfaces with most popular 8-bit microprocessors such as Am9080A, Am8085, MC6800 by Motorola and Z80 by Zilog. It can also be used for 16-bit microprocessors such as AmZ8000,* but its performance with such 16-bit microprocessors is somewhat hindered by its 8-bit external data bus.

Although the external interface is only 8 bits wide, the Am9511A internally is a 16-bit microprogrammed, stack-oriented floating point machine. It includes not only floating point operations but fixed point as well. In addition to the basic add, subtract, multiply and divide operations, transcendental derived functions are also included. A data sheet of Am9511A is included in Appendix A.

AMD FPU Interface



Am9511/12 From AMD datasheet Ó 74LS04 EIN 4.7K +5V INT Am25LS2521 A1 - A7 - A7 B1 - B7 8080 MPU EOUT С Am9080A CS END 18MHz C/D Ao WR HLDA 9511 FPU DBIN \$1 Φ1 Am6224 DBIN HLDA WR 42 \$2 Am9511A SYNC SYNC DB0-7 Do-7 Do-7 DB0-7 READY READY RESET RESET Am8238 IOR RD WR STSTB IOW STSTB RDYIN 10K 1K $\phi_2 TTL$ EACK PAUSE O +12V Oŝ INTA +5V O-CLK RESET

AMD FPU





Note: Pin 1 is marked for orientation.

AMD FPU Format



Am9511/12 From AMD datasheet

FLOATING POINT FORMAT

The format for floating-point values in the Am9511A is given below. The mantissa is expressed as a 24-bit (fractional) value; the exponent is expressed as an unbiased two's complement 7-bit value having a range of -64 to +63. The most significant bit is the sign of the mantissa (0 = positive, 1 = negative), for a total of 32 bits. The binary point is assumed to be to the left of the most significant mantissa bit (bit 23). All floating-point data values must be normalized. Bit 23 must be equal to 1, except for the value zero, which is represented by all zeros.



AMD FPU Instructions



Am9511/12

From AMD datasheet

Command Mnemonics in Alphabetical Order.

ACOS	ARCCOSINE	LOG	COMMON LOGARITHM
ASIN	ARCSINE	LN	NATURAL LOGARITHM
ATAN	ARCTANGENT	NOP	NO OPERATION
CHSD	CHANGE SIGN DOUBLE	POPD	POP STACK DOUBLE
CHSF	CHANGE SIGN FLOATING	POPF	POP STACK FLOATING
CHSS	CHANGE SIGN SINGLE	POPS	POP STACK SINGLE
COS	COSINE	PTOD	PUSH STACK DOUBLE
DADD	DOUBLE ADD	PTOF	PUSH STACK FLOATING
DDIV	DOUBLE DIVIDE	PTOS	PUSH STACK SINGLE
DMUL	DOUBLE MULTIPLY LOWER	PUPI	PUSH #
DMUU	DOUBLE MULTIPLY UPPER	PWR	POWER (XY)
DSUB	DOUBLE SUBTRACT	SADD	SINGLE ADD
EXP	EXPONENTIATION (ex)	SDIV	SINGLE DIVIDE
FADD	FLOATING ADD	SIN	SINE
FDIV	FLOATING DIVIDE	SMUL	SINGLE MULTIPLY LOWER
FIXD	FIX DOUBLE	SMUU	SINGLE MULTIPLY UPPER
FIXS	FIX SINGLE	SQRT	SQUARE ROOT
FLTD	FLOAT DOUBLE	SSUB	SINGLE SUBTRACT
FLTS	FLOAT SINGLE	TAN	TANGENT
FMUL	FLOATING MULTIPLY	XCHD	EXCHANGE OPERANDS DOUBLE
FSUB	FLOATING SUBTRACT	XCHF	EXCHANGE OPERANDS FLOATING
AUSERIA/BC III		XCHS	EXCHANGE OPERANDS SINGLE



i8086 History



WikiSemi

History of the 8086

The path to the 8086 was not as direct and planned as you might expect. Its earliest ancestor was the Datapoint 2200, a desktop computer/terminal from 1970. The Datapoint 2200 was before the creation of the microprocessor, so it used an 8-bit processor built from a board full of individual TTL integrated circuits. Datapoint asked Intel and Texas Instruments if it would be possible to replace that board of chips with a single chip. Copying the Datapoint 2200's architecture, Texas Instruments created the TMX 1795 processor (1971) and Intel created the 8008 processor (1972). However, Datapoint rejected these processors, a fateful decision. Although Texas Instruments couldn't find a customer for the TMX 1795 processor and abandoned it, Intel decided to sell the 8008 as a product, creating the microprocessor market. Intel followed the 8008 with the improved 8080 (1974) and 8085 (1976) processors. (I've written more about early microprocessors here.)



Datapoint 2200 computer. Photo courtesy of Austin Roche.

i8086 History



Microcode

One of the hardest parts of computer design is creating the control logic that tells each part of the processor what to do to carry out each instruction. In 1951, Maurice Wilkes came up with the idea of microcode: instead of building the control logic from complex logic gate circuitry, the control logic could be replaced with special code called microcode. To execute an instruction, the computer internally executes several simpler micro-instructions, which are specified by the microcode. With microcode, building the processor's control logic becomes a programming task instead of a logic design task.

Microcode was common in mainframe computers of the 1960s, but early microprocessors such as the 6502 and Z-80 didn't use microcode because early chips didn't have room to store microcode. However, later chips such as the 8086 and 68000, used microcode, taking advantage of increasing chip densities. This allowed the 8086 to implement complex instructions (such as multiplication and string copying) without making the circuitry more complex. The downside was the microcode took a large fraction of the 8086's die; the microcode is visible in the lower-right corner of the die photos.3

NU U AR	M . MAN MAN . MN		
A AA	A AA A	A A	A 88 A
RM HAN			
is a saag			

A section of the microcode ROM.

WikiSemi

i8086 History



WikiSemi —

Why did the IBM PC pick the Intel 8088 processor?7 According to Dr. David Bradley, one of the original IBM PC engineers, a key factor was the team's familiarity with Intel's development systems and processors. (They had used the Intel 8085 in the earlier IBM Datamaster desktop computer.) Another engineer, Lewis Eggebrecht, said the Motorola 68000 was a worthy competitor6 but its 16bit data bus would significantly increase cost (as with the 8086). He also credited Intel's better support chips and development tools.5

In any case, the decision to use the 8088 processor cemented the success of the x86 family. The IBM PC AT (1984) upgraded to the compatible but more powerful 80286 processor. In 1985, the x86 line moved to 32 bits with the 80386, and then 64 bits in 2003 with AMD's Opteron architecture. The x86 architecture is still being extended with features such as AVX-512 vector operations (2016). But even though all these changes, the x86 architecture retains compatibility with the original 8086.

i8086 16-bit MPU

1st 16-bit MPU



(MIN MODE)

(HOLD)

(HLDA)

(WR)

(M/10)

(DT/R)

(DEN)

(ALE)

(INTA)

1978

40

39

38

37

36

35

34

33

32

31

30

29

28

27

26

25

24

23

22

21

MAX MODE

Ucc

AD 15

A16/S3

A17/S4

A18/S5

A19/S6

BHE/S7

MN/MX

RQ/GT0

RQ/GT1

LOCK

<u>52</u>

<u>S1</u>

<u>S0</u>

050

051

TEST

READY

RESET

RD





60

Intel i8086 Die WikiSemi

DR JEFF

5

Dr Jeff

S

Jeff Drobman ©2016-23



© 2022 Dr Jeff Dro

Intel i8086 Package



World's 1st 16-bit MPU



The 8086 chip, in 40-pin ceramic DIP package.



The 8086 die is visible in the middle of the integrated circuit package.

MPU/MCU Generations



MosTech 6502



Used in the Apple II



The **MOS Technology 6502** is an 8-bit microprocessor that was designed by a small team led by Chuck Peddle for MOS Technology. The design team had formerly worked at Motorola on the Motorola 6800 project; the 6502 is essentially a simplified, less expensive and faster version of that

6502 8-bit MPU



Apple II in 1977 CISC

Other than ALU, what are the basic components of a CPU?

Most of the answers are for more complicated CPUs, with caches, pipelines, DMA etc. But the basic components for a working CPU are much fewer. The 8-bit 6502 microprocessor, introduced in 1975 and used in the Apple][computer and other early personal computers, had only 3510 transistors (compared to the many billions in today's CPUs). Its basic block diagram was fairly simple and easy to understand:



Not shown is the Instruction Register (IR) and decoder logic, which holds the instruction being executed which was fetched from memory.

6502 8-bit MPU



CISC Apple II in 1977 The 6502 had one 8-bit accumulator, and two 8-bit index registers, 8-bit stack pointer, and a 16-bit program counter so it could address a maximum of 65536 bytes.

MOS 6502 registers				
FEDCBA98	76543210	(bit position, hex)		
Main registers				
	А	Accumulator		
Index registers				
	х	X index		
	Y	Y index		
00000001	S	Stack Pointer		
Program counter				
F	PC	Program Counter		
Status register				
	NV-BDIZC	P Processor flags		

The high byte of the stack address is hardwired to 1, so stack addresses ranged from 0x1FF (initial value) to 0x100.


6502 in 1974



Quora



Yowan Rajcoomar · Follow IT Engineer (2018–present) · 1y

Related How were microprocessors designed before Verilog and VHDL?

The earliest designs were hand drawn.

Example: Sheet representing the logic and buses of a 6502 from 1974:



Courtery of Donald F. Harson, Dept. of Elec. Engr., Univ. of Mississippi, University, MS 38677

MPU Generations

DR JEFF

ARE

SOF

Jeff Drobman ©2016-23

D 🔊 J

Dr Jeff



M68000 16-bit MPU



1980

Motorola introduces the 68000 microprocessor



ARM History





Legend

CPU design

Company ISA

The Acorn Archimedes is a family of personal computers designed by Acorn Computers of Cambridge, England. The systems were based on Acorn's own ARM architecture processors and the proprietary operating systems Arthur and RISC OS. The first models were introduced in 198

÷





Arm (previously officially written all caps as ARM and usually written as such today), previously Advanced RISC Machine, originally Acorn RISC Machine, is a family of reduced instruction set computing (RISC) architectures for computer processors, configured for various environments. Arm Holdings develops the architecture and licenses it to other companies, who design their own products that implement one of those architectures—including systems-on-chips (SoC) and systems-on-modules (SoM) that incorporate memory, interfaces, radios, etc. It also designs cores that implement this instruction set and licenses these designs to a number of companies that incorporate those core designs into their own products.

Processors that have a RISC architecture typically require fewer transistors than those with a complex instruction set computing (CISC) architecture (such as the x86 processors found in most personal computers), which improves cost, power consumption, and heat dissipation. These characteristics are desirable for light, portable, battery-powered devices—including smartphones, laptops and tablet computers, and other embedded systems^{[3][4][5]}—but are also useful for servers and desktops to some degree. For supercomputers, which consume large amounts of electricity, Arm is also a power-efficient solution.^[6]

CISC vs RISC:



Complex/Reduced Instruction Set Architecture

Microprocessor History

> 1971-85: CISC (8/16-bit)

- ♦ Intel i4004 (4-bit)
- ♦ Intel i8008 (8-bit) → i8080 → i8085, Z80 → i8086 (16-bit) → "x86"
- ♦ Motorola 6800 (8-bit) → 6502 → 68000 (16-bit)
- ♦ IBM PC used i8088 (8/16-bit) in 1981 \rightarrow i80n86 ("x86") \rightarrow Pentiums

(now RISC)

- ➤ 1985-2000: RISC (32/64-bit)
 - \diamond SPARC* (UC Berkeley \rightarrow Sun/Oracle)
 - ♦ MIPS* (Stanford)
 - ♦ PowerPC (Motorola/IBM)
 - \diamond AMD 29K
 - \diamond Intel i960

♦ ARM*

*still exist

CISC Instruction Cycle





MCS-8 BASIC INSTRUCTION CYCLE



RISC:



Reduced Instruction Set Architecture

- Key Architecture of RISC
 - Reduced ISA: small set of instructions (minor)
 - > Fast execution: <u>single cycle</u> only
 - Reduced impact of <u>memory</u>
 - ♦ No microprogram (key change)
 - Instructions scale to vertical microinstructions (single-cycle)
 - eliminates ~30% chip area
 - ♦ LOAD-STORE (only) memory references
 - ♦ Full general register sets
 - ♦ Cache memory
 - On-chip
 - Multi-level
 - Harvard architecture separate I and D
 - Pipelining
 - ♦ 4 or 5 stages
 - ♦ Interlocks
 - Hardware (SPARC, 29K)
 - Software (MIPS): compiler manages pipeline scheduling

RISC Pipelines





CISC/RISC Pipelines





Chips



Microcontrollers (MCU)

1st MCU



Quora

Tom Crosley

×

B.S. in Electrical Engineering, Iowa State University · June 30

Which is the first microcontroller?

The very first microcontroller was invented by Texas Instruments in 1971, and was called the TMS1802NC. It was used at TI internally in its calculator products between 1972 and 1974.



The first commercially available microcontroller was the TMS 1000, also from Texas Instruments. It was released in 1974. It combined read-only memory, read/write memory, processor and clock on one chip and was targeted at embedded systems. The TMS 1000 was used in Texas Instruments' own Speak & Spell educational toy.

Intel

In response to this, Intel came out with the 8048 family in 1976. It was mostly replaced by the 8051 in 1980, which became on of the most widely used microcontrollers.

Note this was before user programmable memory became available. These early microcontrollers all used masked read-only memory (ROM), which meant the program was developed using an emulator, and when declared ready for production, a binary file would be sent to TI or Intel, and chips would be manufactured with the program already burned into them. This resulted in long lead times, and a disaster if a bug was found after the chips were programmed.

i8051 MCU

MCS-51

1980



From Wikipedia, the free encyclopedia

Intel 8051

The **Intel MCS-51** (commonly termed **8051**) is a single chip microcontroller (MCU) series developed by Intel in 1980 for use in embedded systems. The architect of the Intel MCS-51 instruction set was John H. Wharton.^{[1][2]} Intel's original versions were popular in the 1980s and early 1990s and enhanced binary compatible derivatives remain popular today. It is an example of a complex instruction set computer, and has separate memory spaces for program instructions and data.

Intel P8051 microcontroller

Intel's original MCS-51 family was developed using N-type metal-oxide-semiconductor (NMOS) technology like its predecessor Intel MCS-48, but later versions, identified by a letter C in their name (e.g., 80C51) use complementary metal–oxide–semiconductor (CMOS) technology and consume less power than their NMOS predecessors. This made them more suitable for battery-powered devices.

The family was continued in 1996 with the enhanced 8-bit MCS-151 and the 8/16/32-bit MCS-251 family of binary compatible microcontrollers.^[3] While Intel no longer manufactures the MCS-51, MCS-151 and MCS-251 family, enhanced binary compatible derivatives made by numerous vendors remain popular today. Some derivatives integrate a digital signal processor (DSP). Beyond these physical devices, several companies also offer MCS-51 derivatives as IP cores for use in field programmable cote error (ERCA) or application encoding integrated error (ASIC) decigns.



i8051 MCU





Intel i8051 MCU 2nd Sources Jeff Drobman ©2016-23

Intel MCS-51 second sources



AMD D87C51

MHS S-80C31

OKI M80C31

DR JEFF



Philips PCB80C31

Signetics SCN8031



Temic TS80C32

MCU Block Diagram



8/16/32-bit

BASIC MODEL



MCU



Meanwhile the number of microcontrollers estimated to be shipped in 2019 was estimated at around 27 billion, twelve times as many as the total number of microprocessors. As of 2017, the split was 40% for 32-bit, 33% 8-bit, and 24% 16bit.

MCU = 12x MPU

So it can be estimated there were somewhere around nine billion 8-bit microcontrollers shipped in 2019. They are predominantly used in embedded systems that have a specific task, such as a small (air fryer, microwave oven) or large (washing machine) appliance; automobile cruise control; intelligent thermostat; etc.





Jeff Drobman Just now

as of 5 years ago (when I last checked), the i8051 was still popular along with the PIC16 and 18 (16-bit). many models sold at <\$1. Atmel's AVR is a popular microcontroller family that is customizable.

ARM Cortex M3







PIC 16F MCU



Quora

FIGURE 3-1: PIC16F8X BLOCK DIAGRAM



Embedded Control



	<i>Microprocessors</i> For COMPUTING	<i>Microcontrollers</i> For CONTROL	♦ Real-time ♦ All-in-one				
 <!--</th--><th>All 32/64-bit CPUs Large data processing applications</th><th> Small embedded control app Appliances Disk controllers Remote controllers Garage/gate openers Medium embedded control (User devices (iPods, phone Car/Airplane engine control Car/Airplane braking & saf Car transmission control Home Automation (HAN) Large embedded control (32, Car/Airplane entertainmer Car/Airplane navigation, sy </th><th><pre> viications (8-bit MCU) vieto Tiny vieto Low power vieto Low cost (16-bit MCU) es, etc.) ol fety /64-bit MCU) nt vistems management </pre></th>	All 32/64-bit CPUs Large data processing applications	 Small embedded control app Appliances Disk controllers Remote controllers Garage/gate openers Medium embedded control (User devices (iPods, phone Car/Airplane engine control Car/Airplane braking & saf Car transmission control Home Automation (HAN) Large embedded control (32, Car/Airplane entertainmer Car/Airplane navigation, sy 	<pre> viications (8-bit MCU) vieto Tiny vieto Low power vieto Low cost (16-bit MCU) es, etc.) ol fety /64-bit MCU) nt vistems management </pre>				
	Focus is <i>Memory</i> for large Data Files	 Printers (MF) Communications gear (Will 	Printers (MF) Communications gear (WiFi, cable TV boxes)				
	Large DRAM, Disk, Flash	Focus is I/O – Interre	upts				

Chips



Microprocessor Timeline (Exhaustive)



Date +	Name 🗢	Developer 🗢	Max clock (first version) +	Word size (bits) \$	Process +	Chips ^[5] ♦	Transistors +	MOSFET
1971	4004	Intel	740 kHz	4	10 µm	1	2,250	pMOS
1972	PPS-25	Fairchild	400 kHz	4		2		pMOS
1972	μPD700	NEC		4		1		
1972	8008	Intel	500 kHz	8	10 µm	1	3,500	pMOS
1972	PPS-4	Rockwell	200 kHz	4		1		pMOS
1973	μCOM-4	NEC	2 MHz	4	7.5 μm	1	2,500	NMOS
1973	TLCS-12	Toshiba	1 MHz	12	6 µm	1	2,800 silicon gates	pMOS
1973	Mini-D	Burroughs	1 MHz	8		1		pMOS
1974	IMP-8	National	715 kHz	8		3		pMOS
1974	8080	Intel	2 MHz	8	6 µm	1	6,000	NMOS
1974	μCOM-8	NEC	2 MHz	8		1		NMOS
1974	5065	Mostek	1.4 MHz	8		1		pMOS
1974	μCOM-16	NEC	2 MHz	16		2		NMOS
1974	IMP-4	National	500 kHz	4		3		pMOS
1974	4040	Intel	740 kHz	4	10 µm	1	3,000	pMOS
1974	6800	Motorola	1 MHz	8	-	1	4,100	NMOS
1974	TMS 1000	Texas Instruments	400 kHz	4	8 µm	1	8,000	
1974	PACE	National		16		1		pMOS
1974	ISP-8A/500 (SC/MP)	National	1 MHz	8		1		pMOS
1975	6100	Intersil	4 MHz	12	-	1	4,000	CMOS
1975	TLCS-12A	Toshiba	1.2 MHz	12	-	1		pMOS
1975	2650	Signetics	1.2 MHz	8		1		NMOS
1975	PPS-8	Rockwell	256 kHz	8		1		pMOS



1975	F-8	Fairchild	2 MHz	8		1		NMOS
1975	CDP 1801	RCA	2 MHz	8	5 µm	2	5,000	CMOS
1975	6502	MOS Technology	1 MHz	8	-	1	3,510	NMOS (dynamic)
1975	IMP-16	National	715 kHz	16		5		pMOS
1975	PFL-16A (MN 1610)	Panafacom	2 MHz	16	-	1		NMOS
1975	BPC	Hewlett Packard	10 MHz	16	-	1	6,000 (+ ROM)	NMOS
1975	MCP-1600	Western Digital	3.3 MHz	16	-	3		PMOS
1975	CP1600	General Instrument	3.3 MHz	16		1		NMOS
1976	CDP 1802	RCA	6.4 MHz	8		1		CMOS
1976	Z-80	Zilog	2.5 MHz	8	4 µm	1	8,500	NMOS
1976	TMS9900	Texas Instruments	3.3 MHz	16	-	1	8,000	
1976	8x300	Signetics	8 MHz	8		1		Bipolar
1976	WD16	Western Digital	3.3 MHz	16		5		PMOS
1977	Bellmac-8 (WE212)	Bell Labs	2.0 MHz	8	5 µm	1	7,000	CMOS
1977	8085	Intel	3.0 MHz	8	3 µm	1	6,500	
1977	MC14500B	Motorola	1.0 MHz	1		1		CMOS
1978	6809	Motorola	1 MHz	8	5 µm	1	9,000	
1978	8086	Intel	5 MHz	16	3 µm	1	29,000	
1978	6801	Motorola	-	8	5 µm	1	35,000	
1979	Z8000	Zilog	-	16	-	1	17,500	
1979	8088	Intel	5 MHz	8/16 ^[b]	3 µm	1	29,000	NMOS (HMOS)
1979	68000	Motorola	8 MHz	16/32 ^[c]	3.5 µm	1	68,000	NMOS (HMOS)



Date 🗢	Name 🗢	Developer 🔶	Clock ÷	Word size (bits) \$	Process +	Transistors 🗢 🗕
1980	16032	National Semiconductor	-	16/32	-	60,000
1981	6120	Harris Corporation	10 MHz	12	-	20,000 (CMOS) ^[36]
1981	ROMP	ROMP IBM		32	2 µm	45,000
1981	T-11	DEC	2.5 MHz	16	5 µm	17,000 (NMOS)
1982	RISC-I ^[37]	UC Berkeley	1 MHz	-	5 µm	44,420 (NMOS)
1982	FOCUS	Hewlett Packard	18 MHz	32	1.5 µm	450,000
1982	80186	Intel	6 MHz	16	-	55,000
1987	80C186	Intel	10 MHz	16	-	56,000 (CMOS)
1982	80188	Intel	8 MHz	8/16	-	29,000
1982	80286	Intel	6 MHz	16	1.5 µm	134,000
1983	RISC-II	UC Berkeley	3 MHz	-	3 µm	40,760 (NMOS)
1983	MIPS ^[38]	Stanford University	2 MHz	32	3 µm	25,000
1983	65816	Western Design Center	-	16	-	-
1984	68020	Motorola	16 MHz	32	2 µm	190,000
1984	NS32032	National Semiconductor	-	32	-	70,000
1984	V20	NEC	5 MHz	8/16	-	63,000
1985	80386	Intel	16–40 MHz	32	1.5 µm	275,000
1985	MicroVax II 78032	DEC	5 MHz	32	3.0 µm	125,000
1985	R2000	MIPS	8 MHz	32	2 µm	115,000
1985 ^[39]	Novix NC4016	Harris Corporation	8 MHz	16	3 μm ^[40]	16,000 ^[41]
1986	Z80000	Zilog	-	32	-	91,000
1986	SPARC MB86900	Fujitsu ^{[42][43][44]}	40 MHz	32	0.8 µm	800,000
1986	V60 ^[45]	NEC	16 MHz	16/32	1.5 µm	375,000



1986	SPARC MB86900	Fujitsu ^{[42][43][44]}	40 MHz	32	0.8 µm	800,000
1986	V60 ^[45]	NEC	16 MHz	16/32	1.5 μm	375,000
1987	CVAX 78034	DEC	12.5 MHz	32	2.0 µm	134,000
1987	ARM2	Acorn	8 MHz	32	2 µm	25,000 ^[46]
1987	Gmicro/200 ^[47]	Hitachi	-	-	1 µm	730,000
1987	68030	Motorola	16 MHz	32	1.3 µm	273,000
1987	V70 ^[45]	NEC	20 MHz	16/32	1.5 μm	385,000
1988	R3000	MIPS	25 MHz	32	1.2 μm	120,000
1988	80386SX	Intel	12–33 MHz	16/32	-	-
1988	i960	Intel	10 MHz	33/32	1.5 μm	250,000
1989	i960CA ^[48]	Intel	16–33 MHz	33/32	0.8 µm	600,000
1989	VAX DC520 "Rigel"	DEC	35 MHz	32	1.5 μm	320,000
1989	80486	Intel	25 MHz	32	1 µm	1,180,000
1989	i860	Intel	25 MHz	32	1 µm	1,000,000





1995	UltraSPARC	Sun	143–167 MHz	64	470 nm	5.2
1995	SPARC64	HAL Computer Systems	101–118 MHz	64	400 nm	-
1995	Pentium Pro	Intel	150–200 MHz	32	350 nm	5.5
1996	Alpha 21164A	DEC	400–500 MHz	64	350 nm	9.7
1996	K5	AMD	75–100 MHz	32	500 nm	4.3
1996	R10000	MTI	150–250 MHz	64	350 nm	6.7
1996	R5000	QED	180–250 MHz	-	350 nm	3.7
1996	SPARC64 II	HAL Computer Systems	141–161 MHz	64	350 nm	-
1996	PA-8000	Hewlett-Packard	160–180 MHz	64	500 nm	3.8
1996	POWER2 Super Chip (P2SC)	IBM	150 MHz	32	290 nm	15
1997	SH-4	Hitachi	200 MHz	-	200 nm ^[54]	10 ^[55]
1997	RS64	IBM	125 MHz	64	? nm	?
1997	Pentium II	Intel	233–300 MHz	32	350 nm	7.5
1997	PowerPC 620	IBM, Motorola	120–150 MHz	64	350 nm	6.9
1997	UltraSPARC IIs	Sun	250–400 MHz	64	350 nm	5.4
1997	S/390 G4	IBM	370 MHz	32	500 nm	7.8
1997	PowerPC 750	IBM, Motorola	233–366 MHz	32	260 nm	6.35
1997	K6	AMD	166–233 MHz	32	350 nm	8.8
1998	RS64-II	IBM	262 MHz	64	350 nm	12.5
1998	Alpha 21264	DEC	450–600 MHz	64	350 nm	15.2
1998	MIPS R12000	SGI	270–400 MHz	64	250–180 nm	6.9
1998	RM7000	QED	250–300 MHz	-	250 nm	18
1998	SPARC64 III	HAL Computer Systems	250–330 MHz	64	240 nm	17.6
1998	S/390 G5	IBM	500 MHz	32	250 nm	25
1998	PA-8500	Hewlett Packard	300–440 MHz	64	250 nm	140



Date +	Name 🗢	Developer 🗢	Clock 🗢	Process +	Transistors	Cores per die /
		-			(millions)	Dies per module
2000	Athlon XP	AMD	1.33–1.73 GHz	180 nm	37.5	1/1
2000	Duron	AMD	550 MHz–1.3 GHz	180 nm	25	1 / 1
2000	RS64-IV	IBM	600–750 MHz	180 nm	44	1/2
2000	Pentium 4	Intel	1.3–2 GHz	180–130 nm	42	1/1
2000	SPARC64 IV	Fujitsu	450–810 MHz	130 nm	-	1/1
2000	z900	IBM	918 MHz	180 nm	47	1 / 12, 20
2001	MIPS R14000	SGI	500–600 MHz	130 nm	7.2	1/1
2001	POWER4	IBM	1.1–1.4 GHz	180–130 nm	174	2 / 1, 4
2001	UltraSPARC III	Sun	750–1200 MHz	130 nm	29	1/1
2001	Itanium	Intel	733–800 MHz	180 nm	25	1/1
2001	PowerPC 7450	Motorola	733–800 MHz	180–130 nm	33	1/1
2002	SPARC64 V	Fujitsu	1.1–1.35 GHz	130 nm	190	1/1
2002	Itanium 2	Intel	0.9–1 GHz	180 nm	410	1/1
2003	PowerPC 970	IBM	1.6–2.0 GHz	130–90 nm	52	1/1
2003	Pentium M	Intel	0.9–1.7 GHz	130–90 nm	77	1/1
2003	Opteron	AMD	1.4–2.4 GHz	130 nm	106	1/1
2004	POWER5	IBM	1.65–1.9 GHz	130–90 nm	276	2 / 1, 2, 4
2004	PowerPC BGL	IBM	700 MHz	130 nm	95	2/1
2005	Opteron "Athens"	AMD	1.6–3.0 GHz	90 nm	114	1/1
2005	Pentium D	Intel	2.8–3.2 GHz	90 nm	115	1/2
2005	Athlon 64 X2	AMD	2–2.4 GHz	90 nm	243	2/1
2005	PowerPC 970MP	IBM	1.2–2.5 GHz	90 nm	183	2/1
2005	UltraSPARC IV	Sun	1.05–1.35 GHz	130 nm	66	2/1



				1		
2005	UltraSPARC T1	Sun	1–1.4 GHz	90 nm	300	8 / 1
2005	Xenon	IBM	3.2 GHz	90–45 nm	165	3 / 1
2006	Core Duo	Intel	1.1–2.33 GHz	90–65 nm	151	2/1
2006	Core 2	Intel	1.06–2.67 GHz	65–45 nm	291	2 / 1, 2
2006	Cell/B.E.	IBM, Sony, Toshiba	3.2–4.6 GHz	90–45 nm	241	1+8 / 1
2006	Itanium "Montecito"	Intel	1.4–1.6 GHz	90 nm	1720	2/1
2007	POWER6	IBM	3.5–4.7 GHz	65 nm	790	2/1
2007	SPARC64 VI	Fujitsu	2.15–2.4 GHz	90 nm	543	2/1
2007	UltraSPARC T2	Sun	1–1.4 GHz	65 nm	503	8 / 1
2007	TILE64	Tilera	600–900 MHz	90–45 nm	?	64 / 1
2007	Opteron "Barcelona"	AMD	1.8–3.2 GHz	65 nm	463	4 / 1
2007	PowerPC BGP	IBM	850 MHz	90 nm	208	4 / 1
2008	Phenom	AMD	1.8–2.6 GHz	65 nm	450	2, 3, 4 / 1
2008	z10	IBM	4.4 GHz	65 nm	993	4 / 7
2008	PowerXCell 8i	IBM	2.8–4.0 GHz	65 nm	250	1+8 / 1
2008	SPARC64 VII	Fujitsu	2.4–2.88 GHz	65 nm	600	4 / 1
2008	Atom	Intel	0.8–1.6 GHz	65–45 nm	47	1/1
2008	Core i7	Intel	2.66–3.2 GHz	45–32 nm	730	2, 4, 6 / 1
2008	TILEPro64	Tilera	600–866 MHz	90–45 nm	?	64 / 1
2008	Opteron "Shanghai"	AMD	2.3–2.9 GHz	45 nm	751	4 / 1
2009	Phenom II	AMD	2.5–3.2 GHz	45 nm	758	2, 3, 4, 6 / 1
2009	Opteron "Istanbul"	AMD	2.2–2.8 GHz	45 nm	904	6 / 1



Date +	Name 🗢	Developer 🗢	Clock 🗢	Process +	Transistors (millions) ◆	Cores per die / Dies per module	threads per core
2010	POWER7	IBM	3–4.14 GHz	45 nm	1200	4, 6, 8 / 1, 4	4
2010	Itanium "Tukwila"	Intel	2 GHz	65 nm	2000	2, 4 / 1	2
2010	Opteron "Magny-cours"	AMD	1.7–2.4 GHz	45 nm	1810	4, 6 / 2	1
2010	Xeon "Nehalem-EX"	Intel	1.73–2.66 GHz	45 nm	2300	4, 6, 8 / 1	2
2010	z196	IBM	3.8–5.2 GHz	45 nm	1400	4 / 1, 6	1
2010	SPARC T3	Sun	1.6 GHz	45 nm	2000	16 / 1	8
2010	SPARC64 VII+	Fujitsu	2.66–3.0 GHz	45 nm	?	4 / 1	2
2010	Intel "Westmere"	Intel	1.86–3.33 GHz	32 nm	1170	4–6 / 1	2
2011	Intel "Sandy Bridge"	Intel	1.6–3.4 GHz	32 nm	995 ^[58]	2, 4 / 1	(1,) 2
2011	AMD Llano	AMD	1.0–1.6 GHz	40 nm	380 ^[59]	1, 2 / 1	1
2011	Xeon E7	Intel	1.73–2.67 GHz	32 nm	2600	4, 6, 8, 10 / 1	1–2
2011	Power ISA BGQ	IBM	1.6 GHz	45 nm	1470	18 / 1	4
2011	SPARC64 VIIIfx	Fujitsu	2.0 GHz	45 nm	760	8 / 1	2
2011	FX "Bulldozer" Interlagos	AMD	3.1–3.6 GHz	32 nm	1200 ^[60]	4–8 / 2	1
2011	SPARC T4	Oracle	2.8–3 GHz	40 nm	855	8 / 1	8
2012	SPARC64 IXfx	Fujitsu	1.848 GHz	40 nm	1870	16 / 1	2
2012	zEC12	IBM	5.5 GHz	32 nm	2750	6 / 6	1
2012	POWER7+	IBM	3.1–5.3 GHz	32 nm	2100	8 / 1, 2	4
2012	Itanium "Poulson"	Intel	1.73–2.53 GHz	32 nm	3100	8 / 1	2
2013	Intel "Haswell"	Intel	1.9–4.4 GHz	22 nm	1400	4 / 1	2
2013	SPARC64 X	Fujitsu	2.8–3 GHz	28 nm	2950	16 / 1	2
2013	SPARC T5	Oracle	3.6 GHz	28 nm	1500	16 / 1	8
2014	POWER8	IBM	2.5–5 GHz	22 nm	4200	6, 12 / 1, 2	8



POWER8	IBM	2.5–5 GHz	22 nm	4200	6, 12 / 1, 2	8
Intel "Broadwell"	Intel	1.8-4 GHz	14 nm	1900	2, 4, 6, 8, 12, 16 / 1, 2, 4	2
z13	IBM	5 GHz	22 nm	3990	8 / 1	2
A8-7670K	AMD	3.6 GHz	28 nm	2410	4 / 1	1
Zen	AMD	3.2–4.1 GHz	14 nm	4800	8, 16, 32 / 1, 2, 4	2
z14	IBM	5.2 GHz	14 nm	6100	10 / 1	2
POWER9	IBM	4 GHz	14 nm	8000	12, 24 / 1	4, 8
SPARC M8 ^[61]	Oracle	5 GHz	20 nm	~10,000 ^[62]	32	8
Intel "Cannon Lake"	Intel	2.2-3.2 GHz	10 nm	?	2/1	2
Zen+	AMD	2.8-3.7 GHz	12 nm	4800	2, 4, 6, 8, 12, 16, 24, 32 / 1, 2, 4	1, 2
Zen 2	AMD	2-4.7 GHz	7 nm	3900	6, 8, 12, 16, 24, 32, 64 / 1, 2, 4	2
z15	IBM	5.2 GHz	14 nm	9200	12/1	2
	POWER8 Intel "Broadwell" z13 A8-7670K Zen z14 POWER9 SPARC M8 ^[61] Intel "Cannon Lake" Zen+ Zen 2 z15	POWER8IBMIntel "Broadwell"Intelz13IBMA8-7670KAMDZenAMDz14IBMPOWER9IBMSPARC M8 ^[61] OracleIntel "Cannon Lake"IntelZen +AMDZen 2AMDz15IBM	POWER8IBM2.5–5 GHzIntel "Broadwell"Intel1.8-4 GHzz13IBM5 GHzA8-7670KAMD3.6 GHzZenAMD3.2–4.1 GHzz14IBM5.2 GHzPOWER9IBM4 GHzSPARC M8 ^[61] Oracle5 GHzIntel "Cannon Lake"Intel2.2-3.2 GHzZen 2AMD2.4.7 GHzz15IBM5.2 GHz	POWER8 IBM 2.5–5 GHz 22 nm Intel "Broadwell" Intel 1.8-4 GHz 14 nm z13 IBM 5 GHz 22 nm A8-7670K AMD 3.6 GHz 28 nm Zen AMD 3.2–4.1 GHz 14 nm z14 IBM 5.2 GHz 14 nm POWER9 IBM 4 GHz 14 nm SPARC M8 ^[61] Oracle 5 GHz 20 nm Intel "Cannon Lake" Intel 2.2-3.2 GHz 10 nm Zen + AMD 2.8-3.7 GHz 12 nm Zen 2 AMD 2-4.7 GHz 7 nm z15 IBM 5.2 GHz 14 nm	POWER8 IBM 2.5–5 GHz 22 nm 4200 Intel "Broadwell" Intel 1.8-4 GHz 14 nm 1900 z13 IBM 5 GHz 22 nm 3990 A8-7670K AMD 3.6 GHz 28 nm 2410 Zen AMD 3.2–4.1 GHz 14 nm 4800 z14 IBM 5.2 GHz 14 nm 6100 POWER9 IBM 4 GHz 14 nm 6100 SPARC M8 ^[61] Oracle 5 GHz 20 nm ~10,000 ^[62] Intel "Cannon Lake" Intel 2.2-3.2 GHz 10 nm ? Zen + AMD 2.8-3.7 GHz 12 nm 4800 Zen 2 AMD 2-4.7 GHz 7 nm 3900	POWER8IBM2.5–5 GHz22 nm42006, 12 / 1, 2Intel "Broadwell"Intel1.8-4 GHz14 nm19002, 4, 6, 8, 12, 16 / 1, 2, 4213IBM5 GHz22 nm39908 / 1A8-7670KAMD3.6 GHz28 nm24104 / 1ZenAMD3.2–4.1 GHz14 nm48008, 16, 32 / 1, 2, 4214IBM5.2 GHz14 nm610010 / 1POWER9IBM4 GHz14 nm800012, 24 / 1SPARC M8 ^[61] Oracle5 GHz20 nm~10,000 ^[62] 32Intel "Cannon Lake"Intel2.2-3.2 GHz10 nm?2 / 1Zen+AMD2.8-3.7 GHz12 nm48002, 4, 6, 8, 12, 16, 24, 32 / 1, 2, 4Zen 2AMD2-4.7 GHz7 nm39006, 8, 12, 16, 24, 32, 64 / 1, 2, 4Z15IBM5.2 GHz14 nm920012 / 1

2020S [edit]

Date +	Name \$	Developer +	Clock \$	Process \$	Transistors (millions) \$	Cores per die / Dies per module 🕈	threads per core +
2020	Zen 3	AMD	3.4–4.9 GHz	7 nm	?	6, 8, 12, 16 /	2
2020	M1	Apple	3.2 GHz	5 nm	16000	8	1
2021	M1 Max	Apple	3.2 GHz	5 nm	57000	10	1

Chips



Early RISC Microprocessors * AMD 29K * Intel i960, XScale * MIPS R2/3/4000



DUDENE



Bit-slice 1975-85 -



Am29116 16-bit MPU



Bit-slice 1985-88 –



Am29K



AMD Am29000

Berkeley **RISC** (Patterson)

From Wikipedia, the free encyclopedia

The AMD Am29000, commonly shortened to 29k, is a family of 32-bit RISC microprocessors and microcontrollers developed and fabricated by Advanced Micro Devices (AMD). Based on the seminal Berkeley RISC, the 29k added a number of significant improvements. They were, for a time, the most popular RISC chips on the market, widely used in laser printers from a variety of manufacturers.

Several versions were introduced during the period from 1988 to 1995, beginning with the 29000. The final model, the **29050**, was the first superscalar version, retiring up to four instructions per cycle, and also including a greatly improved floating point unit (FPU).

In late 1995 AMD dropped development of the 29k because the design team was transferred to support the PC side of the business. What remained of AMD's embedded business was realigned towards the embedded 186 family of 80186 derivatives. The majority of AMD's resources were then concentrated on their high-performance, desktop x86 clones, using many of the ideas and individual parts of the latest 29k to produce the AMD K5.

The 29000 evolved from the same Berkeley RISC design that also led to the Sun SPARC and Intel i960.

One design element used in all of the Berkeley-derived designs is the concept of register windows, a technique used to speed up procedure calls significantly. The idea is to use a large set of registers as a stack, loading local data into a set of registers


Register Windows

The 29000 evolved from the same Berkeley RISC design that also led to the Sun SPARC and Intel i960.

Design [edit]

One design element used in all of the Berkeley-derived designs is the concept of register windows, a technique used to speed up procedure calls significantly. The idea is to use a large set of registers as a stack, loading local data into a set of registers during a call, and marking them "dead" when the procedure returns. Values being returned from the routines would be placed

during a call, and marking them "dead" when the procedure returns. Values being returned from the routines would be placed in the "global page", the top eight registers in the SPARC (for instance). The competing early RISC design from Stanford University, the Stanford MIPS, also looked at this concept but decided that improved compilers could make more efficient use of general purpose registers than a hard-wired window.

In the original Berkeley design, SPARC, and i960, the windows were fixed in size. A routine using only one local variable would still use up eight registers on the SPARC, wasting this expensive resource. It was here that the 29000 differed from these earlier designs, using a variable window size. In this example only two registers would be used, one for the local variable, another for the return address. It also added more registers, including the same 128 registers for the procedure stack, but adding another 64 for global access. In comparison, the SPARC had 128 registers in total, and the global set was a standard window of eight. This change resulted in much better register use in the 29000 under a wide variety of workloads.

The 29000 also extended the register window stack with an in-memory (and in theory, in-cache) stack. When the window filled the calls would be pushed off the end of the register stack into memory, restored as required when the routine returned. Generally, the 29000's register usage was considerably more advanced than competing designs based on the Berkeley concepts.

Another difference with the Berkeley design is that the 29000 included no special-purpose condition code register. Any register could be used for this purpose, allowing the conditions to be easily saved at the expense of complicating some code. An instruction prefetch buffer was used that stored up to 16 instructions, used to improve performance during branches—the 29000 did not include any branch prediction system so there was a delay if a branch was taken (nor was it originally superscalar, so it could not "do both sides" as is common in some designs). The buffer mitigated this by storing four instructions from the other side of the branch, which could be run instantly while the buffer was re-filled with new instructions from memory.

Condition Codes = Flags

a.23992	30
AMD 29030.	





The first 29000 was released in 1988, including a built-in MMU but floating point support was offloaded to the 29027 FPU. Units with failed MMU's or BTC's were sold as the 29005.



Versions [edit]

The first 29000 was released in 1988, including a built-in MMU but floating point support was offloaded to the 29027 FPU. Units with failed MMU's or BTC's were sold as the 29005.

The last general-purpose version was the **29050**. This was a significant upgrade to the original design, the first superscalar version which could execute instructions out-of-order and speculatively. Up to six instructions could be worked on at the same time through various pipeline stages, and four could be retired at any cycle. The 29050 also included a significantly improved FPU. The 29050 was initially available with clock rates of 25, 50, and 75 MHz. A 100 MHz version was introduced later.^[1]

Several portions of the 29050 design were used as the basis for the K5 series of x86-compatible processors. The FPU adder and multiplier were carried over with some layout changes, a nanocode engine was added to the FPU to accommodate the complex instructions found in x86 but not on the 29050, while the rest of the core design was used along with complex microcode to translate x86 instructions to 29k-like 'uops' on the fly.

The Honeywell 29KII is a cpu based on the AMD 29050, and it was extensively used in real-time avionics.

29050 → K5 (x86 Pentium)





Am29000

Am29030



Am29040

Am29050

iAPX 432





The **iAPX 432** is a discontinued computer architecture introduced in 1981. It was Intel's first 32-bit processor design. The main processor of the architecture, the *general data processor*, is implemented as a set of two separate integrated circuits, due to technical limitations at the time. Although some equi

Forerunner of **i960**





Intel i960

From Wikipedia, the free encyclopedia

Intel's **i960** (or **80960**) was a RISC-based microprocessor design that became popular during the early 1990s as an embedded microcontroller. It became a best-selling CPU in that segment, along with the competing AMD 29000.^[2] In spite of its success, Intel stopped marketing the i960 in the late 1990s, as a result of a settlement with DEC whereby Intel received the rights to produce the StrongARM CPU. The processor continues to be used for a few military applications.



i960



Die photos













Intel 80960MX

Intel 80960KA

Intel 80960SA

Intel 80960CA

Intel 80960CF

Intel 80960JA











Intel

i960



Origin [edit] The i960 design was begun in response to the failure of Intel's iAPX 432 design of the early 1980s. The iAPX 432 was intended to directly support high-level languages that supported tagged, protected, garbage-collected memory—such as Ada and Lisp—in hardware. Because of its instruction-set complexity, its multi-chip implementation, and design flaws, the iAPX 432 was very slow in comparison to other processors of its time.

In 1984, Intel and Siemens started a joint project, ultimately called BiiN, to create a high-end, fault-tolerant, object-oriented computer system programmed entirely in Ada. Many of the original i432 team members joined this project, although a new lead architect, Glenford Myers, was brought in from IBM. The intended market for the BiiN systems was high-reliability-computer users such as banks, industrial systems, and nuclear power plants.

Architecture [edit]

To avoid the performance issues that plagued the i432, the central i960 instruction-set architecture was a RISC design, which was only implemented in full in the i960MX. The memory subsystem was 33-bits wide—to accommodate a 32-bit word and a "tag" bit to implement memory protection in hardware. In many ways, the i960 followed the original Berkeley RISC design, notably in its use or register windows, an implementation-specific number of caches for the per-subroutine registers that allowed for fast subroutine calls. The competing Stanford University design MIPS, did not use this system, instead relying on the compiler to generate optimal subroutine call and return code. In common with most 32-bit designs, the i960 has a flat 32-bit memory space, with no memory segmentation, except for the i960MX, which could support up to 2²⁶ "objects", each up to 2³² bytes in size.^[4] The i960 architecture also anticipated a superscalar implementation, with instructions being simultaneously dispatched to more than one unit within the processor.

Intel ARMv5 XScale

WIKIPEDIA

The Free Encyclopedia



PXA255A0C300

439

4390316

KOREA 74391083A

INTEL @ ©'01

XScale is a <u>microarchitecture</u> for <u>central processing units</u> initially designed by <u>Intel</u> implementing the <u>ARM architecture</u> (version 5) <u>instruction set</u>. XScale comprises several distinct families: IXP, IXC, IOP, PXA and CE (see more below), with some later models designed as <u>system-on-a-chip</u> (SoC). Intel sold the PXA family to <u>Marvell Technology Group</u> in June 2006.^[1] Marvell then extended the brand to include processors with other <u>microarchitectures</u>, like <u>ARM's Cortex</u>.

The XScale architecture is based on the **ARMv5TE** <u>ISA</u> without the <u>floating-point</u> instructions. XScale uses a <u>seven-stage</u> integer and an <u>eight-stage</u> memory <u>super-pipelined</u> <u>microarchitecture</u>. It is the successor to the Intel <u>StrongARM</u> line of <u>microprocessors</u> and <u>microcontrollers</u>, which Intel acquired from <u>DEC</u>'s Digital Semiconductor division as part of a settlement of a lawsuit between the two companies. Intel used the **StrongARM** to replace its ailing line of outdated <u>RISC</u> processors, the <u>i860</u> and <u>i960</u>

All the generations of XScale are 32-bit **ARMv5TE** processors manufactured with a 0.18 μ m or 0.13 μ m (as in IXP43x parts) proce have a 32 <u>KB</u> data <u>cache</u> and a 32 KB instruction cache.

ISA





MIPS I (32-bit) [R2000/3000]



*MIPS III (64-bit) [R4000]

✤MIPS64 (64-bit)

- Superset of 32-bit ISA
- Adds 64-bit ops ("Double")



MIPS microprocessors [edit]

The first MIPS microprocessor, the **R2000**, was announced in 1985. It added multiple-cycle multiply and divide instructions in a somewhat independent on-chip unit. New instructions were added to retrieve the results from this unit back to the register file; these result-retrieving instructions were interlocked.

The R2000 could be booted either big-endian or little-endian. It had thirty-one 32-bit general purpose registers, but no condition code register (the designers considered it a potential bottleneck), a feature it shares with the AMD 29000 and the Alpha. Unlike other registers, the program counter is not directly accessible.

The R2000 also had support for up to four co-processors, one of which was built into the main CPU and handled exceptions, traps and memory management, while the other three were left for other uses. One of these could be filled by the optional **R2010** FPU, which had thirty-two 32-bit registers that could be used as sixteen 64-bit registers for double-precision.



MIPS Technologies

From Wikipedia, the free encyclopedia (Redirected from MIPS Computer)

MIPS Technologies, Inc., formerly **MIPS Computer Systems, Inc.**, was an American fabless semiconductor design company that is most widely known for developing the MIPS architecture and a series of RISC CPU chips based on it.^{[1][2]} MIPS provides processor architectures and cores for digital home, networking, embedded, Internet of things and mobile applications.^{[3][4]}

MIPS Technologies, Inc. is owned^[5] by Wave Computing, who acquired it from Tallwood MIPS Inc., a company indirectly owned by Tallwood Venture Capital. Tallwood bought it on 2017-10-25 from Imagination Technologies, a UK-based company best known for their PowerVR graphics processor family.^[6] Imagination Technologies had previously bought MIPS after CEVA, Inc. pulled out of a bidding on 2013-02-08.

MIPS Technologies, Inc.



The former MIPS Technologies building in Santa Clara Type Subsidiary Industry RISC microprocessors Acquired in 2018 by Wave Fate Computing 1984; 36 years ago Founded John L. Hennessy / Founder 2013 / Defunct Headquarters Sunnyvale, California, U.S. Key people Sandeep Vij Products Semiconductor intellectual property Number of up to 50 (according to LinkedIn employees

in May 2018), previously 146 (September 2010) Wave Computing

Parent



History [edit]	Defunct	2013 🖍
MIRS Computer Systems Inc. was founded in 1094 ^{[7][8]} by a group of respectively from Stanford University that	Headquarters	Sunnyvale, California, U.S.
included John L. Honnessy and Chris Davier. These researchers had worked on a preject called MIDS (for	Key people	Sandeep Vij
Included John L. Hennessy and Chris Rowen. These researchers had worked on a project called MIPS (for	Products	Semiconductor intellectual
Microprocessor without Interlocked Pipeline Stages), one of the projects that pioneered the RISC concept. Other	property	
ncipal founders were Skip Stritter, formerly a Motorola technologist, and John Moussouris, formerly of IBM. ^[9] Number of up to 50 (accord		up to 50 (according to LinkedIn
The initial CEO was Vaemond Crane, formerly President and CEO of Computer Consoles Inc., who arrived in	employees in May 2018), previously 146 (September 2010)	
February 1985 and departed in June 1989. He was replaced by Bob Miller, a former senior IBM and Data General	Parent	Wave Computing 🖍
executive. Miller ran the company through its IPO and subsequent sale to Silicon Graphics.	Website	www.mips.com &

In 1988, MIPS Computer Systems designs were noticed by Silicon Graphics (SGI) and the company adopted the

MIPS architecture for its computers.^[10] A year later, in December 1989, MIPS held its first IPO. That year, Digital Equipment Corporation (DEC) released a Unix workstation based on the MIPS design.

After developing the R2000 and R3000 microprocessors, a management change brought along the larger dreams of being a computer vendor. The company found itself unable to compete in the computer market against much larger companies and was struggling to support the costs of developing both the chips and the systems (MIPS Magnum). To secure the supply of future generations of MIPS microprocessors (the 64-bit R4000), SGI acquired the company in 1992^[11] for \$333 million^{[12][13]} and renamed it as MIPS Technologies Inc., a wholly owned subsidiary of SGI.^[14]

During SGI's ownership of MIPS, the company introduced the R8000 in 1994 and the R10000^[15] in 1996 and a follow up the R12000 in 1997.^[16] During this time, two future microprocessors code-named *The Beast* and *Capitan* were in development; these were cancelled after SGI decided to migrate to the Itanium architecture^[17] in 1998.^{[12][18]} As a result, MIPS was spun out as an intellectual property licensing company, offering licences to the MIPS architecture as well as microprocessor core designs.



Company timeline [edit]

Year 🕈	\$
1981	Dr. John Hennessy at Stanford University founds and leads Stanford MIPS, a research program aimed at building a microprocessor using RISC principles.
1984	MIPS Computer Systems, Inc. co-founded by Dr. John Hennessy, Skip Stritter, and Dr. John Moussouris ^[43]
1986	First product ships: R2000 microprocessor, Unix workstation, and optimizing compilers
1988	R3000 microprocessor
1989	First IPO in November as MIPS Computer Systems with Bob Miller as CEO
1991	R4000 microprocessor
1992	SGI acquires MIPS Computer Systems. Transforms it into internal MIPS Group, and then incorporates and renames it to MIPS Technologies, Inc. (a wholly owned subsidiary of SGI)
1994	R8000 microprocessor
1994	Sony PlayStation released, using an R3000 CPU with custom GTE coprocessor
1996	R10000 microprocessor; Nintendo 64 released, incorporating a cut down R4300 processor.
1998	Re-IPO as MIPS Technologies, Inc
1999	Sony PlayStation 2 released, using an R5900 cpu with custom vector coprocessors
2002	Acquires Algorithmics Ltd, a UK-based MIPS development hardware/software and consultancy company.
September 6, 2005	Acquires First Silicon Solutions (FS2), a Lake Oswego, Oregon company as a wholly owned subsidiary. FS2 specializes in silicon IP, design services and OCI (On-Chip Instrumentation) development tools for programming, testing, debug and trace of embedded systems in SoC, SOPC, FPGA, ASSP and ASIC devices.
2007	MIPS Technologies acquires Portugal-based mixed-signal intellectual property company Chipidea
February 2009	MIPS Joins Linux Foundation ^[44]
May 8, 2009	Chipidea is sold to Synopsys.
June 2009	Android is ported to MIPS ^[45]



Wiki

R3000: 32-bit CPU → *pipelined* (5 stages)

The R3000 succeeded the R2000 in 1988, adding 32 KB (soon increased to 64 KB) caches for instructions

and data, along with support for shared-memory multiprocessing in the form of a cache coherence protocol. While there were flaws in the R3000s multiprocessing support, it was successfully used in several successful multiprocessor computers. The R3000 also included a built-in MMU, a common feature on CPUs of the era. The R3000, like the R2000, could be paired with a **R3010** FPU. The R3000 was the first successful MIPS design in the marketplace, and eventually over one million were made. A speed-bumped version of the R3000 running up to 40 MHz, the **R3000A** delivered a performance of 32 VUPs (VAX Unit of Performance). The MIPS R3000A-compatible **R3051** running at 33.8688 MHz was the processor used in the Sony PlayStation though it didn't have FPU or MMU. Third-party designs include Performance Semiconductor's **R3400** and IDT's **R3500**, both of them were R3000As with an integrated R3010 FPU. Toshiba's **R3900** was a virtually first SoC for the early handheld PCs that ran Windows CE. A radiation-hardened variant for space applications, the Mongoose-V, is a R3000 with an integrated R3010 FPU.

The **R4000** series, released in 1991, extended MIPS to a full 64-bit architecture, moved the FPU onto the main die to create a single-chip microprocessor, and had a high clock frequency of 100 MHz at introduction. However, in order to achieve the clock frequency, the caches were reduced to 8 KB each and they took three cycles to access. The high operating frequencies were achieved through the technique of deep pipelining (called super-pipelining at the time). The improved **R4400** followed in 1993. It had larger 16 KB primary caches, largely bug-free 64-bit operation, and support for a larger L2 cache.

MIPS, now a division of SGI called MTI, designed the low-cost **R4200**, the basis for the even cheaper **R4300i**. A derivative of this microprocessor, the NEC VR4300, was used in the Nintendo 64 game console.^[1]

R4000: 1st 64-bit CPU \rightarrow super-pipelined (8 stages)

Quantum Effect Devices (QED), a separate company started by former MIPS employees, designed the R4600 Orion, the R4700 Orion, the R4650 and the R5000. Where the R4000 had pushed clock frequency and sacrificed cache capacity, the QED designs emphasized large caches which could be accessed in just two cycles and efficient use of silicon area.

MIPS I– Base (R2000) Org

DR JEFF



IDT's MIPS R3000 Die





First MIPS RISC CPUs



32-bit







Wiki

R4700



R4700 Orion with the exposed silicon chip, fabricated by IDT, designed by Quantum Effect Devices



Top-side view of package for R4700 ₽ Orion

Chips



Advanced RISC Microprocessors Apple A14/M1 Intel Core i3/5/7/9 AMD Zen 3 Mobile SoC's

Apple M1



5-nanometer process

The first personal computer chip built with this cutting-edge technology.

16 billion transistors

The most we've ever put into a single chip.



Apple Event





AMD vs Intel: CPU Families Jeff Drobman @2016-23

		(intel)
Market Segment	AMD	Intel
Desktop	Ryzen 4K/Athlon 3K	Core i7/i9 (10 th gen)
Laptop	Ryzen 4000	Ice Lake
Gaming	Ryzen Threadripper +Radeon	Core Extreme
Server/Workstn	Ерус	Xeon

According to the company, the AMD Ryzen 4700 G series desktop processor offers up to 2.5x multi-threaded performance compared to the previous generation, up to 5% greater single-thread performance than the Intel Core i7-9700, up to 31% greater multithreaded performance than the Intel Core i7-9700, and up to 202% better graphics performance than the Intel Core i7-9700.

Intel New Chips



Ø

Why doesn't Intel have as strong of integrated graphics on their CPUs, such as their Intel UHD 630 graphics, compared to AMD's Vega 11 integrated graphics?



Brett Bergan, Building PC's for 25 years

Answered 48m ago

Unfortunately for AMD fans, Vega 11 was a great product that found its way to ONE processor (actually two if you consider the 2400G and 3400G two different CPUs)

But that detail aside, Intel has been recycling the same 14nm "Skylake" HD 530/630 GPU for five years already. I have a sneaky suspicion that 10th gen Comet Lake CPU models consist of a lot of recycled Coffee Lake silicon that didn't get sold in 2018. The i3–10100 hyperthreaded quad is essentially a i7–7700 that has a locked multiplier set at 4.3GHz

Same CPU. Same GPU. Just three generations later.

10nm Ice Lake with its somewhat improve ... (more)

AMD Ryzen



8

8



Norman Latifov, knows Turkish

Answered 2h ago

It is the latest mobile CPU from AMD. Ryzen 4000 CPUs are only available on laptops and they are the fastest mobile CPUs available right now. I am using a Lenovo Yoga slim 7 with r7 4800u and before I bought it I did a lot of research. Based on reviews, they are even faster than 11th gen Intel CPUs that are yet to come. Although their integrated GPU is not as good as iris graphics (Intel 11th gen CPUs' integrated GPU), vega series GPUs are still a good option. And in my opinion, ryzen 4000 CPUs have a great multicore performance.



Zachary Hawkshaw, AMD Hardware Connoisseur Answered 10h ago

I don't know where you got your information, but that's not true. The 2700X only has 4.8 billion transistors while the 3700X has 19.2 billion.

Transistor count doesn't necessarily mean more performance by itself. The 3700X is better because it has a newer architecture (Zen2 vs Zen+) with improvements to the Infinity Fabric, as well as a higher turbo frequency.

AMD Ryzen



8



John B. Anderson III, IT Consultant, PC Integrator, 20+ Years in IT and Gaming.

Answered Mon

Well, it's only on the Laptop side... Since they skipped it to make them both match desktop/laptop for Zen 3 architecture.

Laptops with the 4XXX name are actually Zen 2 processors.

See some guy in marketing thought it'd be a good idea to call Zen processors on laptops 2XXX series, and so when the 2XXX series came out on desktop the laptop was already at 2XXX so they called them 3XXX.

Example Desktop CPU Ryzen 5 2600x the laptop version would be a Ryzen 5 3550h

AMD Ryzen 5 2600 vs 3550H ☑ with the typical lower performance on laptop vs desktop CPU. Same Zen+ architecture in both.

On the Zen 2 Archetuxture

AMD Ryzen 5 3600 vs 4600HS 🖄 we see the Ryzen 5 4600 HS ~5% of the speed.

There are no Zen 3 laptop processors as of the time I'm writing 11/2020

Typically we call it a "Gen / Generation" when they name a model with a change in the first digit.

Example Core i5 10XXX would be a 10th gen i5. Ryzen 7 1800x would be a 1st Gen Ryzen, However, Since the laptops, 4XXX were the only "Gen 4s" but they were technically Gen 3's well the answer is somewhat tricky.

AMD vs Intel



Sockets

The current Threadripper sTRX4 socket is an LGA design with 4094 pins.. they're kind of mindboggling to look at. Modern EPYC processors use a mechanically identical but electrically tweaked Socket SP3r3 socket. Older chips use TR4 and SP3r2 sockets, respectively. The EPYC series and Threadripper Pro actually support up to 2TiB DDR4 DRAM on eight channels and 128 PCI Express 4.0 links. Today's standard Threadrippers support 128 GiB or 256GiB DDR4 DRAM on four channels, with up to 88 PCI Express 4.0 links, and of course, up to 64 CPU cores on all three platforms.



So you're looking for an Intel processor for consumers/workstations that's more or less similar to Threadripper. Intel's answer to EPYC is Xeon, so are there any mainstream i-series processors that correspond closely? The latest high-end

AMD vs Intel



Sockets

AMD Threadrippers are essentially consumer versions of the EPYC line of server processors. There are differences, but the basic idea is the same: four DDR4 memory channels, high core count, etc. My aging Threadripper system has "only" sixteen processor cores and the usual four 64-bit DDR4 memory channels.

Like all Ryzen family processors, Threadrippers are made of multiple "chiplets" connected by ultra high speed Infinity Fabric links. Each chiplet so far contains up to eight processor cores. The central chip in generation 2 and later Threadrippers is a I/O chip, supporting PCI Express links, that sort of thing.



The current Threadripper sTRX4 socket is an LGA design with 4094 pins.. they're kind of mindboggling to look at. Modern EPYC processors use a mechanically identical but electrically tweaked Socket SP3r3 socket. Older chips use TR4 and SP3r2 sockets, respectively. The EPYC series and Threadripper Pro actually support up to 2TiB DDR4 DRAM on eight channels and 128 PCI Express 4.0 links. Today's standard Threadrippers support 128 GiB or 256GiB DDR4 DRAM on four channels, with up to 88 PCI Express 4.0 links, and of course, up to 64 CPU cores

AMD vs Intel i9



So you're looking for an Intel processor for consumers/workstations that's more or less similar to Threadripper. Intel's answer to EPYC is Xeon, so are there any mainstream i-series processors that correspond closely? The latest high-end Xeons and Phi processors use Intel's LGA 3647 socket. This socket supports six channels of DDR4 memory, but there is no consumer version of an LGA3647 processor.



So the Intel answer to complete with Threadripper is the LGA2066 socket, also called Socket R4. There are lower-end Xeons that also use this socket. This supports DDR4 up to 256GiB on four channels, 48 PCI Express 3.0 lanes (with an additional 24 PCI Express 3.0 links in the X299 chipset). Current LGA2066 chips offer up to 18 CPU cores.

RISC-V





No, RISC-V is 1980s done correctly, 30 years later.

It still concentrates on fixing those problems that we had in 1980s (making instruction set that is easy to pipeline with a simple pipeline), but we mostly don't have anymore, because we have managed to find other, more practical solutions to those problems.

And it's "done correctly" because it abandons the most stupid RISC features such as delay slots. But it ignores most of the things we have learned after that.

ARMv8 is much more advanced and better instruction set which makes much more sense from a technical point of view. Many common things require much more RISC-V instruction than ARMv8 instructions. The only good reason to use RISC-V instead of ARM is to avoid paying licence fees to ARM.

MediaTek vs. Snapdragon



MediaTek Helio P60

Built on the 12nm fabrication process, MediaTek Helio P60 is the upper-mainstream processor of the MediaTek introduced in 2008 mainly for android. The processor is equipped with 4x big ARM Cortex-A73 cores and 4x small and power-efficient ARM Cortex-A53 cores in two clusters. The cores' clusters have the ability to clock the speed up to 2 GHz. The processor also integrates an ARM Mali-G72MP3 GPU and a dedicated AI processing unit.



Snapdragon 636

Built on 14nm Fabrication process, Snapdragon 636 was launched at the same time with eight cores based on Kryo 260 cores ticking at up to 1.8 GHz. It used Adreno 590 as the GPU. The cores of the processor are customizable and hence needed to be

X-Large Dice



🌑 Vladislav Zorov 🔦 · 10h ago

RTX 3080 and 3090 have 28.3 billion transistors, on a 628 mm² die.

RX 6800 XT has 26.8 billion transistors, on a 520 mm² die.

In any case, M1 is definitely not the leader. The leader is this thing, at 1.2 *trillion* transistors:



(note that even the old GPU they're comparing it to in the picture had 21 bn transistors - and the new ones have more - so no way 16 bn is the leader)

wasn't aware of the 21B transistor chip, but it is way larger than a 600 sq mm die (about 1 inch square). that large of a die is likely to be way expensive due to defect densities and silicon wafer costs. Apple still leads in density by using 5nm instead of AMD using 7nm. and that monster "chip" is 8.5 inches square — more a board size than a "chip".

1 sq in 645 sq mm

> 1 inch 25.4 mm





Moore's Law

(see separate slide set Chips & Fabs)

Gordon Moore

Jan 3, 1929 -- March 24, 2023

≪News+

Los Angeles Times Gordon E. Moore, Intel founder and creator of Moore's Law, dies at 94







Gordon Moore



Moore in 1978

Born	Gordon Earle Moore January 3, 1929 Pescadero, California, U.S. ^[1]
Died	March 24, 2023 (aged 94) Waimea, Hawaii, U.S.
Education	University of California, Berkeley (BS) California Institute of Technology (PhD)



transistorsdoublesevery 2 years



Plaque to Moore\'s Law at the technology plaza in Mountain View, beneath a model of the Silicon crystal

Dave Laws (2018)

Number of transistors

Chips-Moore's Law







Moore's Law



Looking Back

- Original in 1965: # Transistors will double every year (12 months)
- Moore revised his prediction in 1975: double every 2 years (24 months)
 - → THIS IS MOORE'S LAW
- Intel's exec David House added CPU complexity would double every 18 months
- History shows # Transistors has doubled every
 - 2 years in *logic*18 months in DRAM/SRAM




Origin of Moore's Law





DR JEFF

SOF

Jeff Drobman ©2016-23

DS.

Dr Jeff

Figure 3. Gordon Moore notes on IC device types. Collection of the Computer History Museum, 102783359.

Year	Device	Function	Transistors	Resistors	Components	LOG ₂
1959	2N697	Transistor	1	0	1	0
1962	Type G	RTL 3 - I/P gate	3	4	7	2.8
1963 (late 62)	Type R	RTL D Flip Flop	15	18	33	5.0
1964	945	DTL R-S Flip Flop	13	21	34	5.1
1965 (late 64)	958	RTL Counter	33	25	58	5.9
1966	9300	TTL Shift Register	85	40	125	7.0
1967	4500	DTL 32- Gate Array	200	64	264*	8.0

Figure 4. Table of component count for devices in photograph.

Microprocessor Timeline



Moore's Law – The number of transistors on integrated circuit chips (1971-2016) Our World

Moore's law describes the empirical regularity that the number of transistors on integrated circuits doubles approximately every two years. This advancement is important as other aspects of technological progress – such as processing speed or the price of electronic products – are strongly linked to Moore's law.



Data source: Wikipedia (https://en.wikipedia.org/wiki/Transistor_count)

The data visualization is available at OurWorldinData.org. There you find more visualizations and research on this topic.

Memory Timeline





Hard disk storage has become denser at an exponential rate over the last 50 years, just like main memory. The dramatic increase in capacity and speed of both has fueled the increasing power of computers.

Future of Moore's Law

Chip design

Transistors

- SiGe
- FinFFT
- INT
- Chip stacks (3D hybrid)
 - Intel/Micron 3D Xpoint
- 🗖 3D
 - NAND Flash (EEPROM)
- ✤ Architecture
 - Specialized hardware (GPU, APU, et²)
 - Reconfigurable hardware (FPGA)
- Thermal/Cooling
 - Microfluidics (liquid cooling)
- Something completely different
 - Molecular computing
 - Quantum computing

"As Moore's Law slows, we are being forced to make tough choices between Power, Performance and Cost." (ARM)



A transistor is a switch. Ordinarily, current cannot flow. When a voltage is applied to the gate, the channel becomes conductive, current flows from the source to the drain, and the transistor switches on.

A finFET transistor raises the channel above the block of silicon upon which the device sits. That allows the gate to wrap around three sides of the channel, improving its electrical properties.

Gate

Drain

Dr Jeff

Jeff Drobman ©2016-23

Source: The Economist

New sorts of transistors can eke out a few more iterations of Moore's

law, but they will get increasingly expensive

Faith no Moore

Selected predictions for the end of Moore's law



Sources: Intel; press reports; The Economist

Chips





(see separate slide set Chips & Fabs)

Wafer Fabs Today



- 1968 🛠 Intel
- 1978 🛠 Micron**
- 1980 🛠 Samsung
- 1987 * TSMC* (1st foundry)
- 2009 ❖ AMD → Global Foundries*
- 2010 Chartered -> Global Foundries*
- ²⁰¹⁴ ❖ IBM → Global Foundries*
 - SMIC* (China)
 - *Pure Foundry
 - **Internal use only

Fabs – AMD, Intel





AMD Sunnyvale Fab 1 1970



Cost x10,000 in 40 years averages to 250x per year



Intel's latest Fab in Hillsboro



An aerial view of Ronler Acres, Intel's largest silicon research and development hub.

Wafers





12" Wafer (2002) (Current Standard)



Die on Wafer



Intel Pentium



Intel Pentium microprocessor die and wafer



Wafers





Quora





After taking into consideration the wasted dies intersecting the "Exclusion Edge" there are a possible 577 dies produced. In this sample, 392 have defects of some sort, leaving only 185 that can be used as intended.

Assuming these are CPU dies with integrated GPU, any defect in a GPU region can result in a CPU with the GPU disabled and be sold at a lower price. If this is a die with six CPU cores, one pair of cores can be disabled to make it a quad core. In that way, maybe half of the defective chips can be salvaged.

With say 14nm production there are fewer transistors per die, and thus a lower probability of defects as compared to 7nm production. The defect rate goes up because there simply are a lot more transistors and a lot more opportunities for a defect to occur.

The problem is greatly compounded when the die is quite large, because then the possible number of good dies per wafer goes down and the probability of getting defects goes up. GPU dies like those made by Nvidia are easier to bin because multiple compute units can be disabled and still result in a perfectly functional product.

Just as a case in point. A RTX 2080 has 46 compute units, while a RTX 2060 KO has only 30. But they both use the same TU104 silicon. The 2060 KO edition has a whopping 16 of its compute units disabled. Fully 1024 of its GPU cores are dead. You simply can't do that with CPU silicon.

Intel Process Nodes



Slower Node Transitions Versus Foundries

IC KNOWLEDGE LLC

	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023
Intel	14nm					10nm				7nm
Samsung	14nm		10nm		7nm	5nm			3nm	
TSMC		16nm	10nm	7nm		5nm		3nm		2nm?

- Intel takes bigger density jumps but less often.
- TSMC and Samsung take smaller jumps more frequently, 5 nodes versus Intel's 3.



3/24/2021

Copyright IC Knowledge LLC, all rights reserved

4

Figure 4. Node Introductions.

Chips



Memories

Memory IC Timeline



ROM- 1st Semiconductor
 DIP packages
 RAM- Bipolar RAMs (SRAM) introduced
 DRAM- IBM conceives DRAM cell (1T, 1C)
 1968 CMOS SRAM- 1st parts by RCA
 1971 Microprocessor & RAM in MOS invented by Intel

1987 * Toshiba intro's Flash EEPROM,

Memory Types



Computer memory types

Wiki

Volatile

RAM DRAM (SDRAM · DDR · GDDR · HBM) · SRAM

Historical

Williams–Kilburn tube (1946–47) • Delay line memory (1947) • Mellon optical memory (1951) • Selectron tube (1952) • Dekatron • T-RAM (2009) • Z-RAM (2002–2010)

Non-volatile

ROM Mask ROM · PROM · EPROM · EEPROM · Flash memory

NVRAM

ReRAM

Early stage NVRAM FeRAM · MRAM · PCM (3D XPoint) · FeFET memory Magnetic

Magnetic tape data storage (Linear Tape-Open) · Hard disk drive

Optical Optical disc • 5D optical data storage

In development CBRAM · Racetrack memory · NRAM · Millipede memory · ECRAM

Historical

Paper data storage (1725) · Drum memory (1932) · Magnetic-core memory (1949) · Plated wire memory (1957) · Core rope memory (1960s) · Thin-film memory (1962) · Disk pack (1962) · Twistor memory (~1968) · Bubble memory (~1970) · Floppy disk (1971)

Computer Memory Org



1.6.1: Some computer components.



Memory Models





System Org: Multilevel Memory

DR JEFF







Figure 9. MCS-8 Memory System

AMD 64-Bit Bipolar SRAM







Characteristics 3101 Typical Delay Access Time 35 ns Typical Power Dissipation 400 mW

 $V_{CC} = Pin 16$ GND = Pin 8

i1101A 256x1 SRAM



10

g

A.

A2



A₀

VDD

8

Am1101A 256x1 SRAM





Am1103 1Kx1 DRAM





Memory Chips



DRAM 1T



Dynamic random-access memory (DRAM) is a type of random access semiconductor memory that stores each bit of data in a memory cell consisting of a tiny capacitor and a transistor, typically a MOSFET. The capacitor can either be charged or discharged; these two states are taken to

SRAM 4T



Static random-access memory is a type of semiconductor random-access memory (RAM) that uses bistable latching circuitry (flip-flop) to store each bit. SRAM exhibits data remanence, but it is still *volatile* in the conventional sense that data is eventually lost when the memory is not powered.

DRAM Timeline



Patterson & Hennessy =

Figure 1.5.1: Growth of capacity per DRAM chip over time (COD Figure 1.11).

The y-axis is measured in kibibits (2¹⁰ bits). The DRAM industry quadrupled capacity almost every three years, a 60% increase per year, for 20 years. In recent years, the rate has slowed down and is somewhat closer to doubling every two years to three years.



DRAM Schematic





Memory Chips



ROM



Read-only memory (ROM) is a type of nonvolatile memory used in computers and other electronic devices. Data stored in ROM cannot be electronically modified after the manufacture of the memory device. Read-only memory is useful for storing software that is rarely changed during the life of the syste ROM (masked)
PROM
EPROM
EEPROM
Flash E²

WIKIPEDIA The Free Encyclopedia



Flash memory is an <u>electronic non-volatile</u> <u>computer memory storage</u> <u>medium</u> that can be electrically erased and reprogrammed. The two main types of flash memory, **NOR flash** and **NAND flash**, are named for the <u>NOR</u> and <u>NAND logic gates</u>. Both use the same cell design, consisting of <u>floating gate</u> <u>MOSFETs</u>. They differ at the circuit level depending on whether the state of the bit line or word lines is pulled high or low: in NAND flash, the relationship between the bit line and the word lines resembles a NAND gate; in NOR flash, it resembles a NOR gate.

Flash memory, a type of <u>floating-gate</u> memory, was invented at <u>Toshiba</u> in 1980 and is based on <u>EEPROM</u> technology. Toshiba began marketing flash memory in 1987.^[1] <u>EPROMs</u> had to be erased completely before they could be rewritten. NAND flash memory, however, may be erased, written, and read in blocks (or pages), which generally are much smaller than the entire device. NOR flash memory allows a single <u>machine word</u> to be written – to an erased location – or read independently. A flash memory device typically consists of one or more flash <u>memory chips</u> (each holding many flash memory cells), along with a separate <u>flash memory controller</u> chip.







A disassembled USB flash drive. The chip on the left is flash memory. The controller is on the right.



WIKIPEDIA The Free Encyclopedia



Flash memory stores information in an array of memory cells made from <u>floating-gate transistors</u>

- In <u>single-level cell</u> (SLC) devices, each cell stores only one bit of information.
- Multi-level cell (MLC) devices, including triplelevel cell (TLC) devices, can store more than one bit per cell.

4 levels \rightarrow 2 bits

2x DRAM density

NAND



NAND memories

WIKIPEDIA The Free Encyclopedia

NAND flash architecture was introduced by Toshiba in 1989.^[97] These memories are accessed much like <u>block devices</u>, such as hard disks. Each block consists of a number of pages. The pages are typically 512,^[98] 2,048 or 4,096 bytes in size. Associated with each page are a few bytes (typically 1/32 of the data size) that can be used for storage of an <u>error correcting code</u> (ECC) <u>checksum</u>.

Typical <u>block sizes</u> include:

•32 pages of 512+16 bytes each for a block size (effective) of 16 KiB
•64 pages of 2,048+64 bytes each for a block size of 128 KiB^[99]
•64 pages of 4,096+128 bytes each for a block size of 256 KiB^[100]
•128 pages of 4,096+128 bytes each for a block size of 512 KiB.
While reading and programming is performed on a page basis, erasure can only be performed on a block basis

WOM!



25120

April 1, 1980 **SUITETHS** FULLY ENCODED, 9046×N, RANDOM ACCESS WRITE-ONLY-MEMORY

Do Not Copy

DESCRIPTION

The Signetic² 25000 Series 9946XN Random Access Write-Only-Memory employs both enhancement and depletion mode P-Channel, N-Channel, and neu(1) channel MOS devices. Although a static device, a single TTL level clock phase is required to drive the on-board multi-port clock generator. Data refresh is accomplished during CB and LH periods⁽¹¹⁾. Quadri-state outputs (when applicable) allow expansion in many directions, depending on organization.

The static memory cells are operated dynamically to yield extremely low power dissipation. All inputs and outputs are directly TTL compatible when proper interfacing circuitry is employed.

Device construction is more or less S.O.S.(2).

FEATURES

- FULLY ENCODED MULTI-PORT ADDRESSING
- WRITE CYCLE TIME 80nS (MAX. TYPICAL)
- WRITE ACCESS TIME(3)
- POWER DISSIPATION 10uW/BIT TYPICAL
- CELL REFRESH TIME 2mS (MIN. TYPICAL)

FINAL SPECIFICATION(10)

BIPOLAR COMPATIBILITY

All data and clock inputs plus applicable outputs will interface directly or nearly directly with bipolar circuits of suitable characteristics. In any event use 1 amp fuses in all power supply and data lines.

INPUT PROTECTION

All terminals are provided with slip-on latex protectors for the prevention of Voltage Destruction. (PILL packaged devices do not require protection).

SILICON PACKAGING

Low cost silicon DIP packaging is implemented and reliability is assured by the use of a non-hermetic sealing technique which prevents the entrapment of harmful ions, but which allows the free exchange of friendly ions.

SPECIAL FEATURES

Because of the employment of the Signetics' proprietary Sanderson-Rabbet Channel the 25120 will provide 50% higher speed than you will obtain.

COOLING

The 25120 is easily cooled by employment of a six-foot

Memory Segments





Typical memory layout for a program with a 32-bit address space.

GiB/TiB (2³⁰/2⁴⁰)



	Actual										
tual			% La	Value	Abbreviation	Binary term	Value	Abbreviation	Decimal		
1024	10 <mark>24</mark>			2 ¹⁰	KiB	kibibyte	10 ³	KB	kilobyte		
0 <mark>48,</mark> 576	1,0 <mark>48,</mark> 576		5%	2 ²⁰	MiB	mebibyte	10 ⁶	MB	megabyte		
0 <mark>74</mark> ×10 ⁹	1.0 <mark>74</mark> %10 ⁹		7%	2 ³⁰	GiB	gibibyte	10 ⁹	GB	gigabyte		
95×10 ¹²	1.0 <mark>995</mark> ×10 ¹²		10%	2 ⁴⁰	TiB	tebibyte	10 ¹²	ТВ	terabyte		
		13%		2 ⁵⁰	PiB	pebibyte	10 ¹⁵	PB	petabyte		
			15%		EiB	exbibyte	10 ¹⁸	EB	exabyte		
			18%		18%		ZiB	zebibyte	10 ²¹	ZB	zettabyte
Actual	Power of 10	Power of 2	Ordin al	2 ⁸⁰	YiB	yobibyte	10 ²⁴	YB	yottabyte		
1024	10 ³	2 ¹⁰	1K								
1,048,576	10 ⁶	2 ²⁰	1M								
1.074x10	10 ⁹	2 ³⁰	1G								
1.0995x10 ¹²	10 ¹²	2 ⁴⁰	1T								
Virtual Memory





Virtual Memory





Chips







Am2505 Multiplier



Bit-slice 1971-80 -

Am2505 Four-Bit by Two-Bit 2's Complement Multiplier Advanced Micro Devices Complex Digital Integrated Circuits



Distinctive Characteristics:

5N 7437 N

#18

- Provides 2's complement multiplication at high speed without correction.
- Can be used in an iterative scheme or time sequenced mode.
- Multiplies two 12-bit signed numbers in typically 200ns.

FUNCTIONAL DESCRIPTION:

The Am2505 is a high-speed digital multiplier that can multiply numbers represented in the 2's complement notation and produce a 2's complement product without correction. The device consists of a 4x2 multiplier that can be connected to form iterative arrays able to multiply numbers either directly, or in a time sequenced arrangement. The device assumes that the most significant digit in a word carries a negative weight, and can therefore be used in arrays where the multiplicand and multiplier have different word lengths. The multiplier uses the guaternary algorithm and performs the function S = X Y + K where K is the input field used to add partial products generated in the array. At the beginning of the array the K inputs are available to add a signed constant to the least significant part of the product. Multiplication of an m bit number by an n bit number in an array results in a product having m + n bits so that all possible combinations of product are accounted for. If a conventional 2's complement product is required the most significant bit can be ignored, and overflow conditions can be detected by comparing the last two product digits. Figure 2 shows how multipliers are connected together in an array. A

Figure 2 shows how multipliers are connected together in an array. A number of connection schemes are possible. Figure 4 shows diagramatically the connection scheme that results in the fastest multiply. If higher speed is required an array can be split into several parts, and the parts added with high-speed look-ahead carry adders such as the Am9340.

added with high-speed look-ahead carry adders such as the Am9340. Provision is made in the design for multiplication in the active high (positive logic) or active low (negative logic) representations simply by reinterpreting the active level of the input operands, the product, and a polarity control P. For a more complete description and applications the user is referred to the Am2505 Application Note.

- Multiplies in active HIGH (positive logic) or active LOW (negative logic) representations.
- Easy correction for unsigned, sign-magnitude or 1's complement multiplication.
- 100% reliability assurance testing in compliance with MIL STD 883.



Am2505 Multiplier





Am2900 Family



Bit-slice 1975-85 -

AMD 2901 bit-slice processor family includes 2901 and 2903 4-bit microprocessors slices, 2909 and 2911 microprogram sequencers, 2910 microprogram controller and other support chips. The 2901 processor consists of 16 4-bit registers, 4-bit ALU and associated decoding/multiplexing circuits. The ALU accepts 9-bit microinstructions that specify source operands, ALU function and the destination register. The 2901 ALU can perform 8 different functions (they are encoded into 3 bits within the microinstruction): addition, subtraction and logic operations. Multiple 2901 bit-slice processors could be combined together to build microprocessors with any data width (in 4 bits increments).

Enhanced version of 2901, AMD 2903 has 9 new special ALU functions used for implementation of multiplication, division and normalization operations. The number of arithmetic and logic ALU functions in 2903 was increased to 15.





Am2900 Family



Members of the Am2900 family [edit]

The Am2900 Family Data Book lists:[22]

- Am2901 4-bit bit-slice ALU (1975)
- Am2902 Look-Ahead Carry Generator
- · Am2903 4-bit-slice ALU, with hardware multiply
- Am2904 Status and Shift Control Unit
- Am2905 Bus Transceiver
- Am2906 Bus Transceiver with Parity
- Am2907 Bus Transceiver with Parity
- Am2908 Bus Transceiver with Parity
- Am2909 4-bit-slice address sequencer
- Am2910 12-bit address sequencer
- Am2911 4-bit-slice address sequencer
- Am2912 Bus Transceiver
- Am2913 Priority Interrupt Expander
- Am2914 Priority Interrupt Controller
- Am2915 Quad 3-State Bus Transceiver
- Am2916 Quad 3-State Bus Transceiver
- Am2917 Quad 3-State Bus Transceiver
- Am2918 Instruction Register, Quad D Register
- Am2919 Instruction Register, Quad Register
- Am2920 Octal D-Type Flip-Flop
- Am2921 1-to-8 Decoder
- Am2922 8-Input Multiplexer (MUX)
- Am2923 8-Input MUX
- Am2924 3-Line to 8-Line Decoder

- Am2925 System Clock Generator and Driver
- Am2926 Schottky 3-State Quad Bus Driver
- Am2927/Am2928 Quad 3-State Bus Transceiver
- Am2929 Schottky 3-State Quad Bus Driver
- Am2930 Main Memory Program Control
- Am2932 Main Memory Program Control
- Am2940 Direct Memory Addressing (DMA) Generator
- Am2942 Programmable Timer/Counter/DMA Generator
- Am2946/Am2947 Octal 3-State Bidirectional Bus Transceiver

Bit-slice

1975-85

- Am2948/Am2949 Octal 3-State Bidirectional Bus Transceiver
- Am2950/Am2951 8-bit Bidirectional I/O Ports
- Am2954/Am2955 Octal Registers
- Am2956/Am2957 Octal Latches
- Am2958/Am2959 Octal Buffers/Line Drivers/Line Receivers
- Am2960 Cascadable 16-bit Error Detection and Correction Unit
- Am2961/Am2962 4-bit Error Correction Multiple Bus Buffers
- Am2964 Dynamic Memory Controller
- Am2965/Am2966 Octal Dynamic Memory Driver

2901 Block Diagram









P1 01

Cn+s

P2 G2

Am2902

Cmay

P3 G3

MPR-018

C., 1 g

Po Go

2901 Chip





2901-A-B Die







DUDENE



Bit-slice 1975-85 -



Am29116 16-bit MPU



Bit-slice 1985-88 –



Section



Debug & Test

ICE Debugger



called setting a breakpoint in the code. The debugger could show the contents of the registers and then we could tell the debugger to jump back into the **loop** function and continue execution of our program.



Figure 3: Wiring diagram for connecting an ESP32-PROG (left) to a DOIT ESP32 (right)

ICE Debugger





Figure 1: An ESP32-PROG (on the right) connected to a DOIT ESP32 device using direct connection to the JTAG pins

ICE Debugger

ARMv7



Address	Machine Code	Opcode	Operands	Description
83:	fd2591	132r	a9, 40018	Make register a9 point at the
				location where i is stored
86:	0988	132i.n	a8, a9, 0	Load the location pointed at by
				a9 into a8
88:	881b	addi.n	a8, a8, 1	Add 1 to the value in a8
8a:	0989	s32i.n	a8, a9, 0	Store a8 in the location pointed
				at by a9
8c:	fd2491	132r	a9, 4001c	Make register a9 point at the
				location where j is stored
8f:	0988	132i.n	a8, a9, 0	Load the location pointed at by
				a9 into a8
91:	880b	addi.n	a8, a8, -1	Subtract 1 from the value in a8
93:	0989	s32i.n	a8, a9, 0	Store a8 in the location pointed
				at by a9



Additional Material







Enter JTAG

As the power and complexity of microprocessors grew, it became harder to make bond-outs to expose all the internal signals that make hardware debugging possible. To address this, manufacturers formed a Joint Tag Action Group (JTAG) to define standards by which a device can expose its internal state using just a few pins.

Many circuit boards have pins labelled JTAG which are used during manufacture and testing. Sometimes these pins can also be used for hardware debugging. Not all processors support hardware debugging connections. The ATmega328P processor used in the Arduino Uno cannot be debugged in this way. However, the ESP32 does provide these connections. Some of the general-purpose input/output (GPIO) pins on an ESP32 can be used as JTAG connectors. To debug code running in hardware, you'll need some way of connecting your development computer to the JTAG signals on the target device. Espressif (the same company that makes the ESP32) produces a great device for this. It is called the ESP32-PROG.



Wikipedia

Boundary scan

From Wikipedia, the free encyclopedia (Redirected from JTAG boundary scan)

Boundary scan is a method for testing interconnects (wire lines) on printed circuit boards or sub-blocks inside an integrated circuit. Boundary scan is also widely used as a debugging method to watch integrated circuit pin states, measure voltage, or analyze sub-blocks inside an integrated circuit.

JTAG

The Joint Test Action Group (JTAG) developed a specification for boundary scan testing that was standardized in 1990 as the IEEE Std. 1149.1-1990. In 1994, a supplement that contains a description of the Boundary Scan Description Language (BSDL) was added which describes the boundary-scan logic content of IEEE Std 1149.1 compliant devices. Since then, this standard has been adopted by electronic device companies all over the world. Boundary scan is now mostly synonymous with JTAG.^{[1][2]}

Debugging [edit]

The boundary scan architecture also provides functionality which helps developers and engineers during development stages of an embedded system. A JTAG Test Access Port (TAP) can be turned into a low-speed logic analyzer.

History [edit]

James B. Angell at Stanford University proposed serial testing.^[4]

IBM developed level-sensitive scan design (LSSD).^{[5][6]}

JTAG





JTAG



Wikipedia

On-chip infrastructure [edit]

To provide the boundary scan capability, IC vendors add additional logic to each of their devices, including *scan cells* for each of the external traces. These cells are then connected together to form the external boundary scan shift register (BSR), and combined with JTAG Test Access Port (TAP) controller support comprising four (or sometimes more) additional pins plus control circuitry.

Some TAP controllers support scan chains between on-chip logical design blocks, with JTAG instructions which operate on those internal scan chains instead of the BSR. This can allow those integrated components to be tested as if they were separate chips on a board. On-chip debugging solutions are heavy users of such internal scan chains.

These designs are part of most Verilog or VHDL libraries. Overhead for this additional logic is minimal, and generally is well worth the price to enable efficient testing at the board level.

For normal operation, the added boundary scan latch cells are set so that they have no effect on the circuit, and are therefore effectively invisible. However, when the circuit is set into a test mode, the latches enable a data stream to be shifted from one latch into the next. Once a complete data word has been shifted into the circuit under test, it can be latched into place so it drives external signals. Shifting the word also generally returns the input values from the signals configured as inputs.

Test mechanism [edit]

As the cells can be used to force data into the board, they can set up test conditions. The relevant states can then be fed back into the test system by clocking the data word back so that it can be analyzed.

By adopting this technique, it is possible for a test system to gain test access to a board. As most of today's boards are very densely populated with components and tracks, it is very difficult for test systems to physically access the relevant areas of the board to enable them to test the board. Boundary scan makes access possible without always needing physical probes.

In modern chip and board design, Design For Test is a significant issue, and one common design artifact is a set of boundary scan test vectors, possibly delivered in Serial Vector Format (SVF) or a similar interchange format.

JTAG



Wikipedia -

JTAG test operations [edit]

Devices communicate to the world via a set of input and output pins. By themselves, these pins provide limited visibility into the workings of the device. However, devices that support boundary scan contain a shift-register cell for each signal pin of the device. These registers are connected in a dedicated path around the device's boundary (hence the name). The path creates a virtual access capability that circumvents the normal inputs and provides direct control of the device and detailed visibility at its outputs.^[3] The contents of the boundary scan are usually described by the manufacturer using a part-specific BSDL file.

Among other things, a BSDL file will describe each digital signal exposed through pin or ball (depending on the chip packaging) exposed in the boundary scan, as part of its definition of the Boundary Scan Register (BSR). A description for two balls might look like this:

"541	(bc_1,	*,	control,	1),"	&		
"542	(bc_1,	GPI051_ATACS1,	output3,	х,	541,	1,	Z)," &
"543	(bc_1,	GPI051_ATACS1,	input,	X),"	&		
"544	(bc_1,	*,	control,	1),"	&		
"545	(bc_1,	GPI050_ATACS0,	output3,	х,	544,	1,	Z)," &
"546	(bc_1,	GPI050_ATACS0,	input,	X),"	&		

That shows two balls on a mid-size chip (the boundary scan includes about 620 such lines, in a 361-ball BGA package), each of which has three components in the BSR: a control configuring the ball (as input, output, what drive level, pullups, pulldowns, and so on); one type of output signal; and one type of input signal.

There are JTAG instructions to SAMPLE the data in that boundary scan register, or PRELOAD it with values.

During testing, I/O signals enter and leave the chip through the boundary-scan cells. Testing involves a number of test vectors, each of which drives some signals and then verifies that the responses are as expected. The boundary-scan cells can be configured to support external testing for interconnection between chips (EXTEST instruction) or internal testing for logic within the chip (INTEST instruction).

JEDEC (EIA)





GRAPHICS DOUBLE DATA RATE 6 (GDDR6) SGRAM STANDARD

JESD250B

Published: Nov 2018

This document defines the Graphics Double Data Rate 6 (GDDR6) Synchronous Graphics Random Access Memory (SGRAM) specification, including features, functionality, package, and pin assignments. The purpose of this Specification is to define the minimum set of requirements for 8 Gb through 16 Gb x16 dual channel GDDR6 SGRAM devices. System designs based on the required aspects of this standard will be supported by all GDDR6 SGRAM vendors providing compatible devices. Some aspects of the GDDR6 standard such as AC timings and capacitance values were not standardized. Some features are optional and therefore may vary among vendors. In all cases, vendor data sheets should be consulted for specifics. This document was created based on some aspects of the GDDR5 Standard (JESD212). Item 1836.99D.

JEDEC (EIA)



JEDEC History

In 1924, the Radio Manufacturers Association (which later became the Electronic Industries Association) was established. In 1944, the Radio Manufacturers Association and the National Electronic Manufacturers Association established the Joint Electron Tube Engineering Council (JETEC), which was responsible for assigning and coordinating type numbers of electron tubes. As the radio industry expanded into the emerging field of electronics, various divisions of the EIA, including JETEC, began to function as semi-independent membership groups. The Council expanded its scope to include solid state devices, and by 1958 the organization was renamed the Joint Electron Device Engineering Council (JEDEC) – one council for tubes and one for semiconductors.



JEDEC initially functioned within the engineering department of EIA where its primary activity was to develop and assign part numbers to devices. Over the next 50 years, JEDEC's work expanded into developing test methods and product standards that proved vital to the development of the semiconductor industry. Among the landmark standards that have come from JEDEC committees are:

JEDEC (EIA)



JEDEC ® Global Standards for the Microele	ectronics Industry		
STANDARDS & DOCUMENTS	COMMITTEES	NEWS	EVENTS & MEETINGS

Why JEDEC Standards Matter

JEDEC committees develop open standards, which are the basic building blocks of the digital economy and form the bedrock on which healthy, high-volume markets are built. For example, JEDEC semiconductor memory standards - from dynamic RAM chips and memory modules to DDR synchronous DRAM and flash components – have enabled huge markets in PCs, servers, digital cameras, MP3 players, smart phones, automotive and HDTV, to name just a few.

Standards enable innovation, serving to commoditize components by lowering their prices while maintaining quality and reliability. This leads suppliers to compete more vigorously on innovative features and gives buyers more variety and a broader selection. The end result is a much larger market than proprietary products could foster, which means more potential sales and revenue.

Standards allow companies to invest more strategically in R&D rather than inventing everything from scratch. Once common form factors are set, companies can base their designs on standards and focus on innovation.