DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Rev 7-14-23

Tech History

Part 1

Computers

by

Dr Jeff Drobman

Dr Jeff Software

Index

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

*see separate slide sets

Computer History Museum

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Computer History Museum

http://www.computerhistory.org/fellowawards/hall/

Image Gallery

CHM Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Decades of Silicon Valley Innovation*

Tube

1930s - Radio Communications (Litton)

1940s - Defense Systems (Varian)

1950s - Magnetic Storage (IBM)

Transistor

1960s – Semiconductors (Fairchild)

1970s - Personal Computing (Apple)

Microchip

1980s – Networks (Xerox)

1990s – Internet (Cisco)

2000s - Mobile (Apple)

* in hardware

Personal Communications

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1684	❖ Optical Telegraph (smoke, mirrors)	
1792	❖ Semaphores (flags)	
1837	❖ Telegraph, Electrical (w/Morse code)	
1843	❖ Fax	Wireline
1870	❖ Telephone	
	❖ Radio (shortwave)	Wireless
	❖ TV (broadcast)	
1969	❖ CATV (interactive)	Wireline
	❖ **Internet**	
	❖ **Email** (1971 private, 1991 public)	
	❖ Teletype (TTY)	Wireline
1989	❖ FAX	
	❖ **WWW** (Web 1.0 – over the Internet)	
	❖ Cell phones (analog voice + data)	
	❖ Messaging	Wireless
	▪ **Texting** (SMS – *thumbing*)	
	▪ **IM**	
2000	❖ Web 2.0 (smart apps)	
	❖ Smart phones – OS alerts, **apps**	
	❖ **Social media** websites	

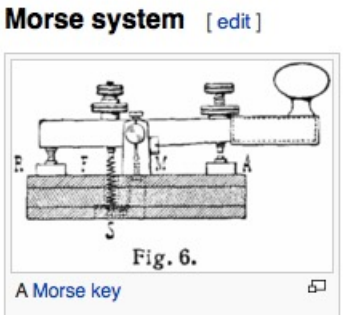

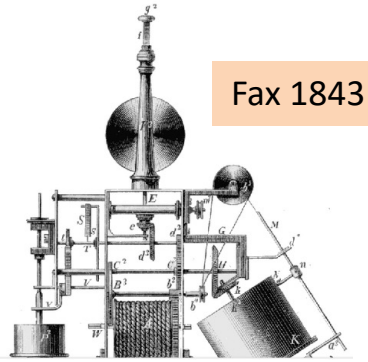

Fax 1843

The *Internet* itself is both –
Wireline (fiber & copper)
& *Wireless* (WiFi, cellular)

Latest trend:
Media streaming services

Fax Machine

Fax 1843

The first fax machine was invented in **1843** by Alexander Bain, a Scottish philosopher. Bain used two pendulums and the mechanisms from a clock to make his machine. The machine, called the Electric Printing Telegram printed messages line by line.

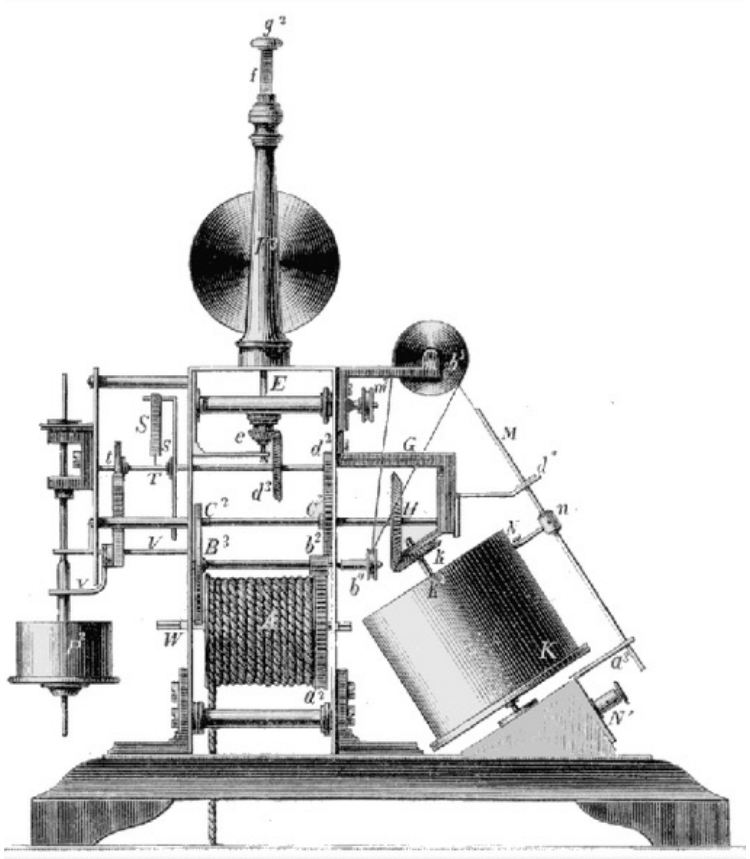

Fax Machine

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1843-1924

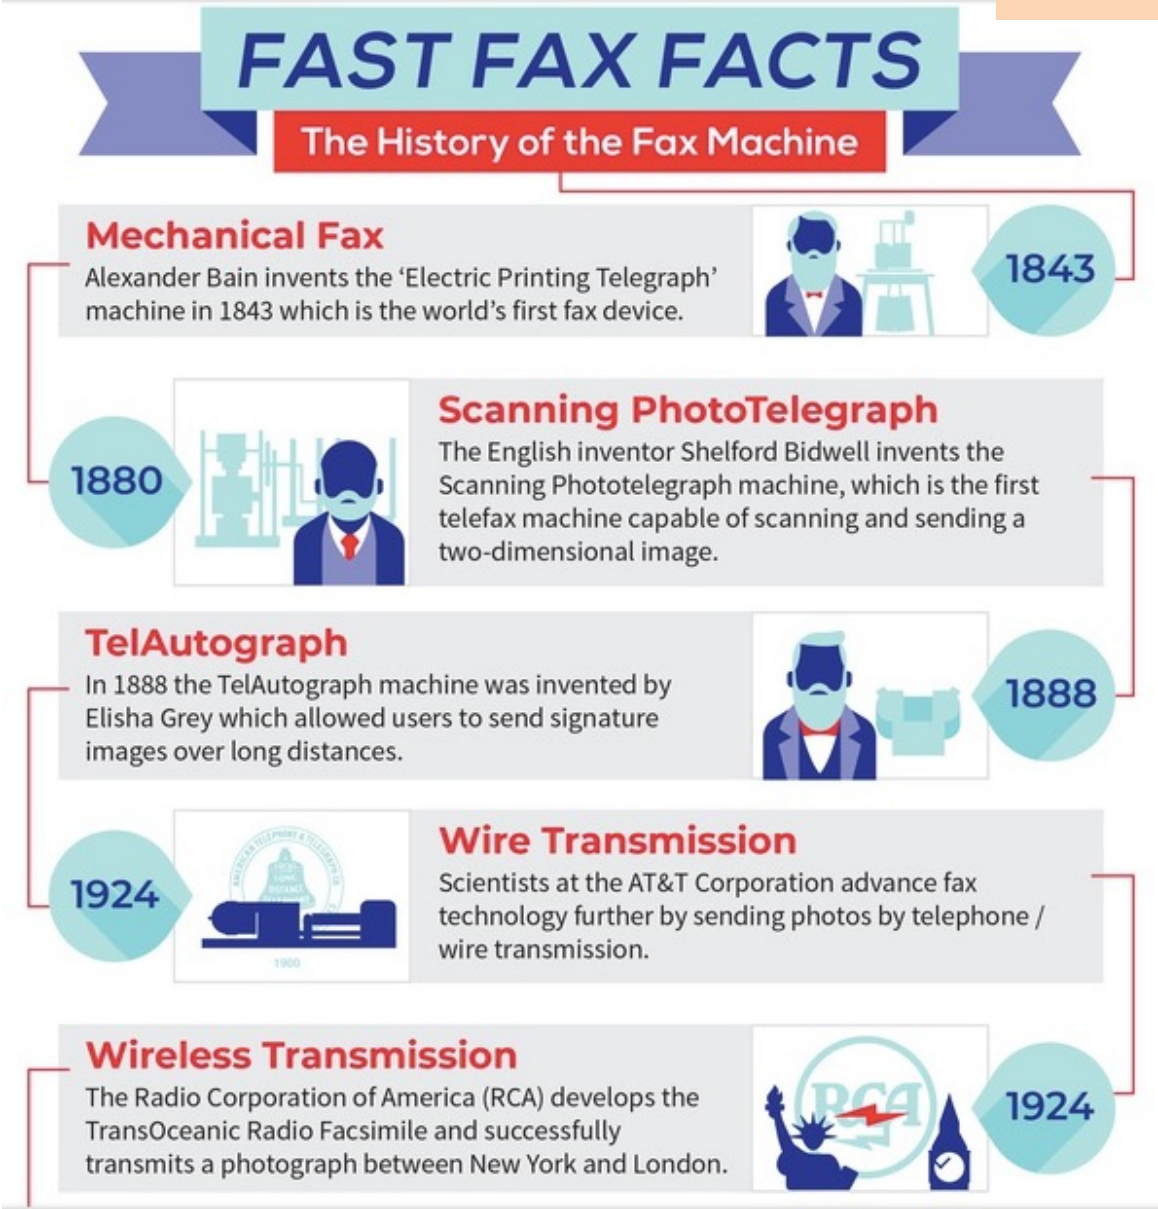

Fax Machine

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1924-1982

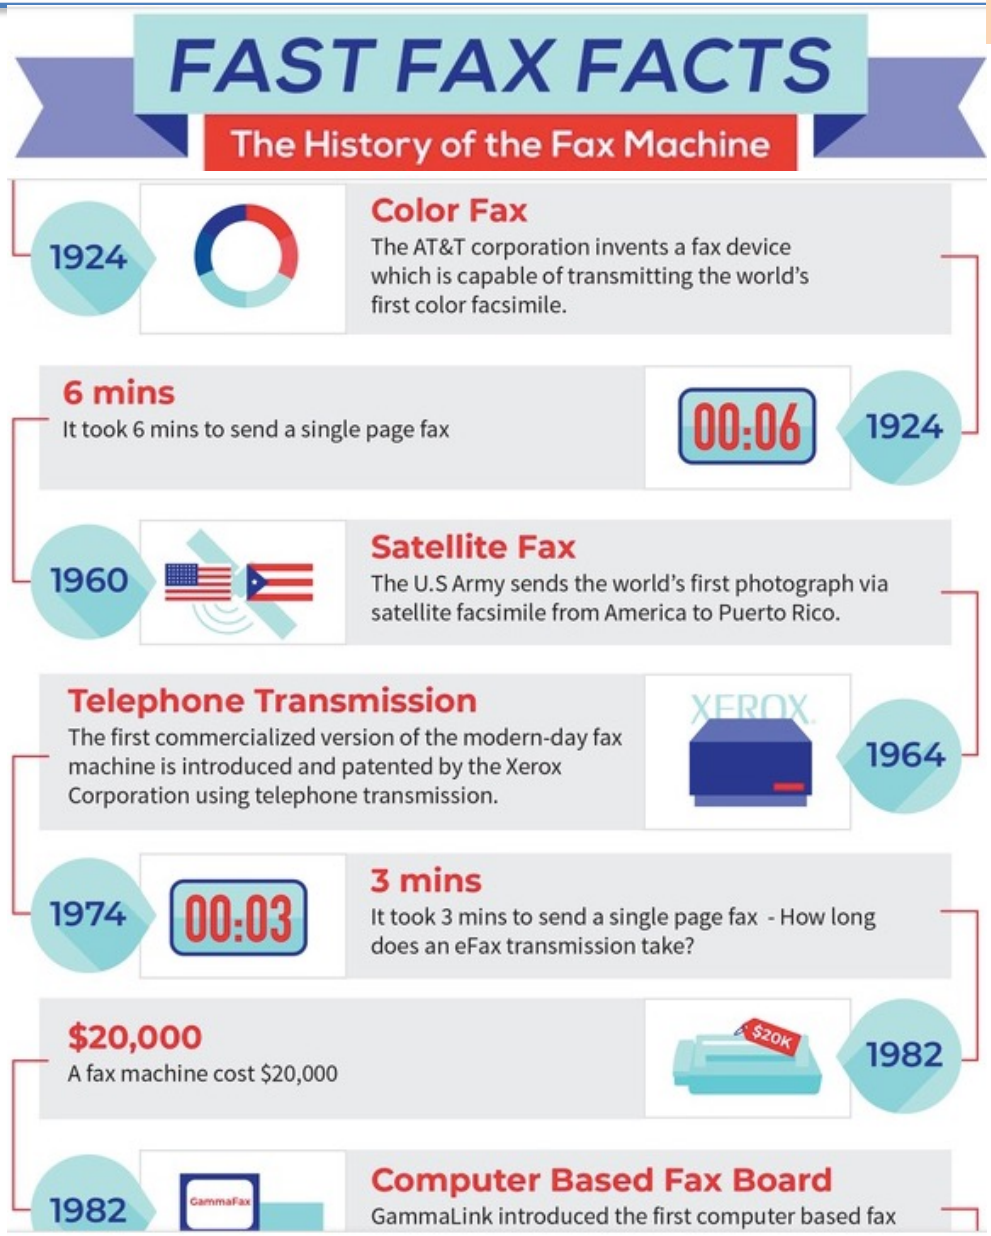

Tech Companies

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Top Tech Titans

Tech Titan Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Revenue in $B | Historical Perspective | As of 2021

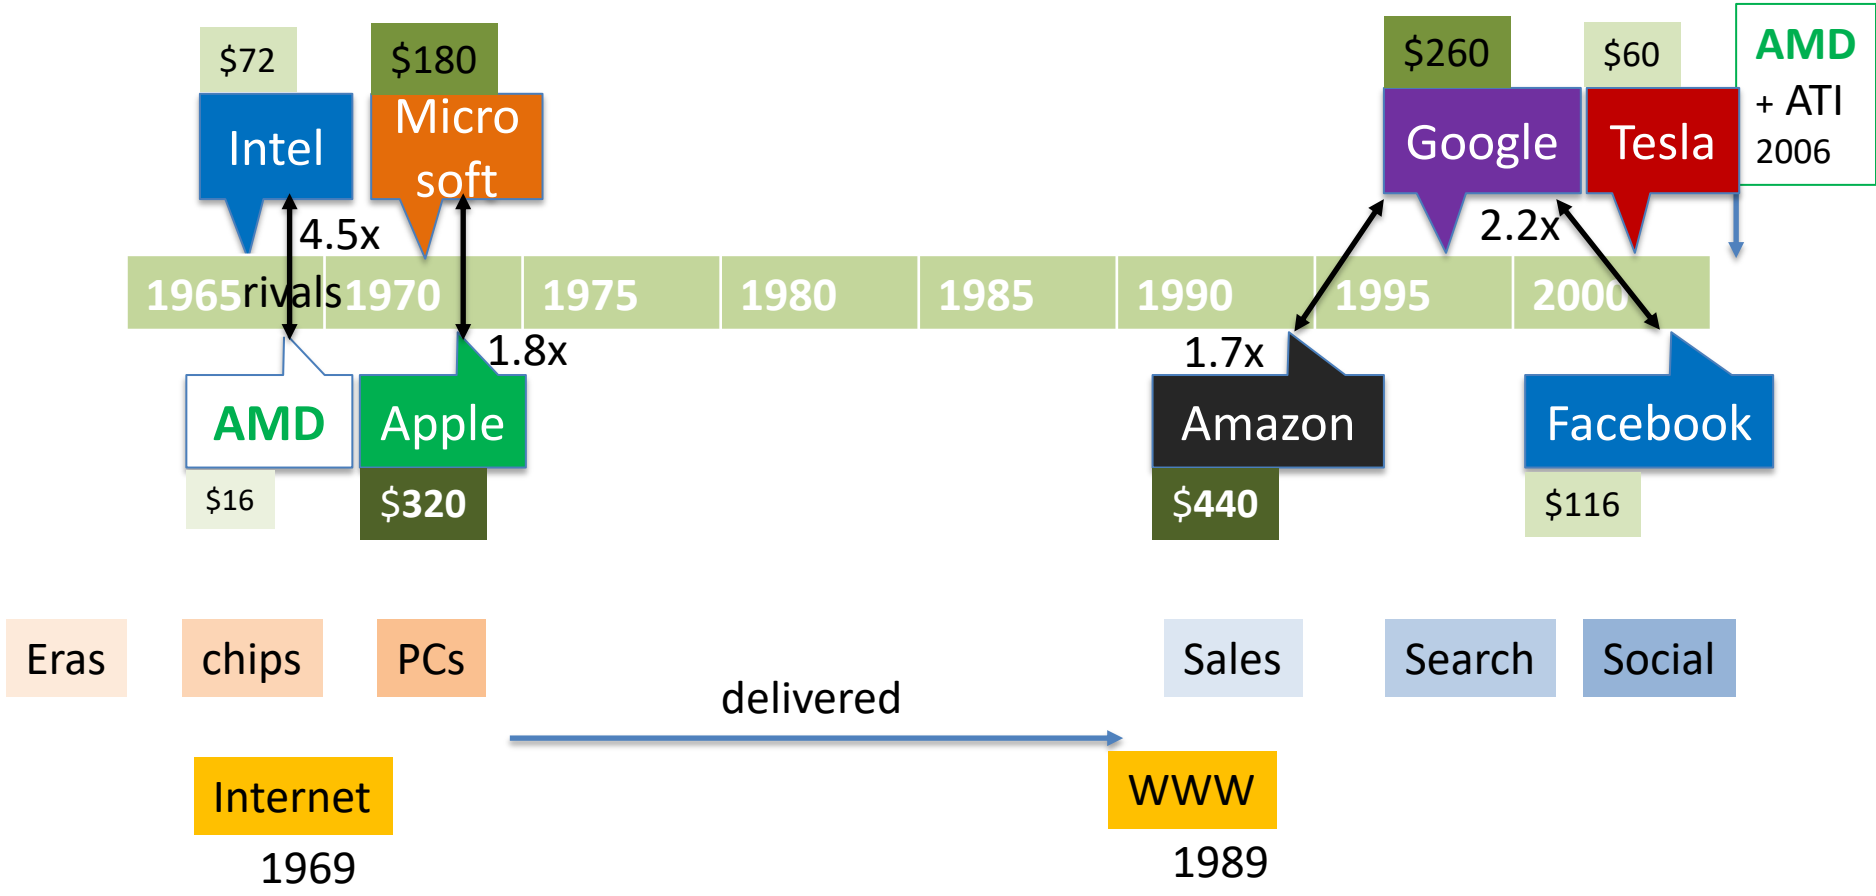

Tech Titan Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Qtr Revenue in $B | Current Perspective | **Q2** 20**23**

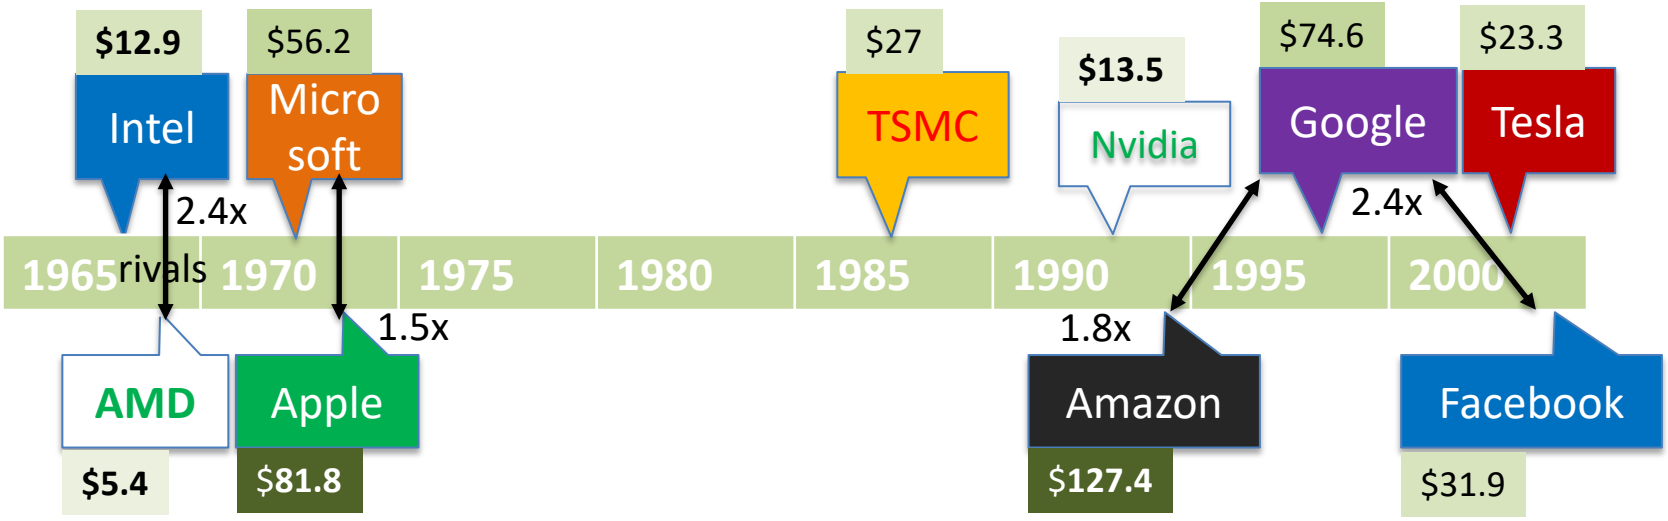

- Other *Industrials*
  - **GM** $40.0 → 2x Tesla
  - **Ford** $39.1
  - **IBM** $16.7 → ~Intel
  - QCOM $9.3
  - TI $4.2
  - NXPI $3.3

- Other *Services*
  - **Netflix** $8.2
  - Visa $8.1
  - **PayPal** $7.4

ISA/SoC Landscape

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

CPU & GPU Cores

- MIPS — Focus
- ARM
  - Mobile
    - ❑ **Apple**
    - ❑ Qualcomm
    - ❑ Samsung*
    - ❑ *Google*

- x86
  - Laptop
  - Desktop
  - Server
    - ❑ Intel* 2/3
    - ❑ AMD 1/3

- ARM
  - ❑ Apple
  - ❑ Nvidia
  - ❑ Qualcomm

*has own fab

- iOT
  - ❑ *Amazon*
  - ❑ *Google*

- Cars
  - ❑ *Tesla*

- Foundry (mfr)
  - ❑ TSMC
  - ❑ Samsung
  - ❑ GlobalFoundries (AMD)
  - ❑ Intel (new)

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD vs Intel

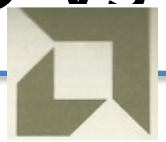

Market Segment	AMD	Intel
Desktop	Ryzen 4K/Athlon 3K	Core i7/i9 (10th gen)
Laptop	Ryzen 4000	Ice Lake
Gaming	Ryzen *Threadripper* +Radeon	Core Extreme
Server/Workstn	Epyc	Xeon

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

$70B

Intel's current logo, used since 2006

Intel's headquarters in Santa Clara, California

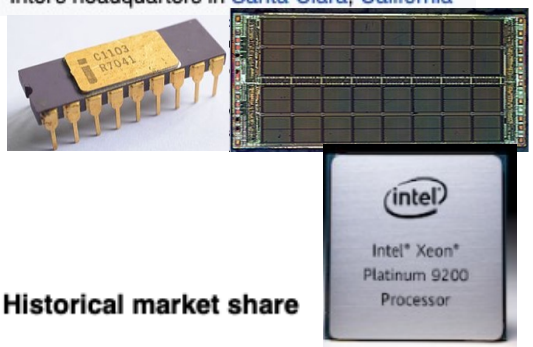

Intel logo used from 1968 to 2006

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jeff Drobman

f. 1968 | Tech Titan | 2019

Founded	July 18, 1968; 51 years ago
Founders	Gordon Moore Robert Noyce
Revenue	▲ US$70.8 billion (2018)
Operating income	▲ US$23.3 billion (2018)
Net income	▲ US$21.0 billion (2018)
Number of employees	110,200 (2019)[2]
Subsidiaries	Mobileye McAfee (49%) Here (15%)

Andy Grove, Robert Noyce and Gordon Moore in 1978

Products

- Central processing units
- Microprocessors
- Integrated graphics processing units (iGPU)
- Systems-on-chip (SoCs)
- Motherboard chipsets
- Network interface controllers
- Modems
- Mobile phones
- Solid state drives
- Wi-Fi and Bluetooth Chipsets
- Flash memory
- Vehicle automation sensors

Historical market share

In the 1980s Intel was among the top ten sellers of semiconductors (10th in 1987) in the world. In 1992,[13] Intel became the biggest chip maker by revenue and has held the position ever since. Other top semiconductor companies include TSMC, Advanced Micro Devices, Samsung, Texas Instruments, Toshiba and STMicroelectronics.

Operating segments [edit]

- **Client Computing Group** – 55% of 2016 revenues – produces hardware components used in desktop and notebook computers.[10]
- **Data Center Group** – 29% of 2016 revenues – produces hardware components used in server, network, and storage platforms.[10]
- **Internet of Things Group** – 5% of 2016 revenues – offers platforms designed for retail, transportation, industrial, buildings and home use.[10]
- **Non-Volatile Memory Solutions Group** – 4% of 2016 revenues – manufactures NAND flash memory and 3D XPoint, branded as Optane, products primarily used in solid-state drives.[10]
- **Intel Security Group** – 4% of 2016 revenues – produces software, particularly security, and antivirus software.[10]
- **Programmable Solutions Group** – 3% of 2016 revenues – manufactures programmable semiconductors (primarily FPGAs).[10]

AMD

$6.5B

f. 1969 | Tech Titan | 2019

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Advanced Micro Devices, Inc.

Old AMD Headquarters
(Sunnyvale, California)

Founder	Jerry Sanders +7
Headquarters	Santa Clara, California, U.S.
Area served	Worldwide
Key people	Lisa Su (President and CEO) John Edward Caldwell (Chairman)
Number of employees	10,100[3] (2018)

Advanced Micro Devices, Inc. (**AMD**) is an American multinational semiconductor company based in Santa Clara, California that develops computer processors and related technologies for business and consumer markets. While initially it manufactured its own processors, the company later outsourced its manufacturing, a practice known as fabless, after GlobalFoundries was spun off in 2009. AMD's main products include microprocessors, motherboard chipsets, embedded processors and graphics processors for servers, workstations, personal computers and embedded system applications.

AMD is the second-largest supplier and only significant rival to Intel in the market for x86-based microprocessors. Since acquiring ATI in 2006, AMD and its competitor Nvidia have maintained a duopoly in the discrete graphics processing unit (GPU) market.[4]

AMD Accelerated Processing Unit

Release date	2011
Architecture	AMD64
Cores	2 to 4
Transistors	32 nm 1.178b (Llano) 32 nm 1.303b (Trinity) 32 nm 1.3b (Richland) 28 nm 2.41b (Kaveri) 14 nm 4.95b (Raven Ridge)

The **AMD Accelerated Processing Unit** (**APU**), formerly known as *Fusion*, is the marketing term for a series of 64-bit microprocessors from Advanced Micro Devices (AMD), designed to act as a central processing unit (CPU) and graphics processing unit (GPU) on a single die.

AMD announced the first generation APUs, *Llano* for high-performance and *Brazos* for low-power devices in January 2011. The second generation *Trinity* for high-performance and *Brazos-2* for low-power devices were announced in June 2012. The third generation *Kaveri* for high performance devices was launched in January 2014, while *Kabini* and *Temash* for low-power devices were announced in the summer of 2013.

The Sony PlayStation 4 and Microsoft Xbox One eighth generation video game consoles both use semi-custom third generation low-power APUs.

Intel CPUs with integrated HD Graphics also have a CPU and GPU on a single die, but they do not offer HSA features (see below).

Microsoft

$126B

f. 1975 | Tech Titan | 2019

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Microsoft Corporation

Microsoft

Building 92 on the Microsoft Redmond campus in Redmond, Washington

Type	Public
Traded as	NASDAQ: MSFT

Industry	Computer software Computer hardware Consumer electronics Social networking service Cloud computing Video games Internet Corporate venture capital
Founded	April 4, 1975; 44 years ago in Albuquerque, New Mexico, U.S.
Founders	Bill Gates Paul Allen
Headquarters	One Microsoft Way, Redmond, Washington, U.S.

Products	Windows · Office · Servers · Skype · Visual Studio · Dynamics · Xbox · Surface · Mobile · List of software
Services	Azure · Bing · LinkedIn · MSDN · Office 365 · OneDrive · Outlook.com · TechNet · Pay · Microsoft Store · Windows Update · Xbox Live

Number of employees	▲ 144,106[2] (2019)

Revenue	▲ US$125.8 billion[1] (2019)
Operating income	▲ US$43.0 billion[1] (2019)
Net income	▲ US$39.2 billion[1] (2019)

Apple

$266B

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Tech Titan 2019

Apple Park in Cupertino, California, April 2018

Formerly	Apple Computer Company (1976–1977) Apple Computer, Inc. (1977–2007)
Founded	April 1, 1976; 43 years ago
Founders	Steve Jobs Steve Wozniak Ronald Wayne
Number of employees	132,000[2] (2018)
Revenue	▲ US$265.595 billion[1] (2018)
Operating income	▲ US$70.898 billion[1] (2018)
Net income	▲ US$59.531 billion[1] (2018)
Industry	Computer hardware Computer software Consumer electronics Cloud computing Digital distribution Fabless silicon design Semiconductors Financial technology Artificial intelligence
Products	Macintosh · iPod · iPhone · iPad · Apple Watch · Apple TV · HomePod · macOS · iOS · iPadOS · watchOS · tvOS · iLife · iWork · Final Cut Pro · Logic Pro · GarageBand · Shazam · Siri
Services	App Store · Apple Arcade · Apple Card · Apple Music (Beats 1) · Apple News+ · Apple Pay (Cash) · Apple Store (Genius Bar) · Apple TV+ · iBooks Store · iCloud · iMessage · iTunes Store · Mac App Store
Subsidiaries	Braeburn Capital · Beats Electronics · FileMaker Inc. · Apple Energy, LLC · Apple Sales International[3] · Apple Services[4] · Apple Worldwide Video[5] · Anobit · Beddit

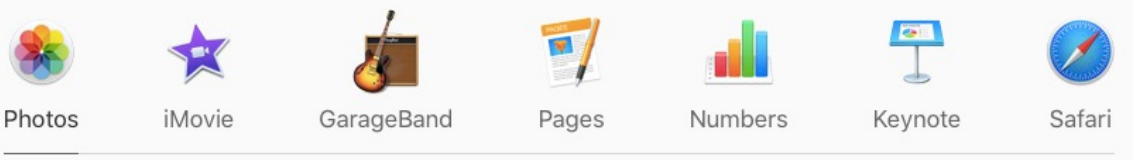

Google

$137B

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Tech Titan 2019

Google LLC

Google's logo since 2015

Google's headquarters, the Googleplex

Formerly	Google Inc. (1998–2017)
Type	Subsidiary
Industry	Internet Cloud computing Computer software Computer hardware Artificial intelligence Advertising
Founded	September 4, 1998; 20 years ago in Menlo Park, California
Founders	Larry Page Sergey Brin
Headquarters	1600 Amphitheatre Parkway, Mountain View, California, U.S.
Revenue	▲ US$136.82 billion (2018)
Operating income	▲ US$26.32 billion (2018)
Net income	▲ US$30.74 billion (2018)

Alphabet Inc. is an American multinational conglomerate headquartered in Mountain View, California. It was created through a corporate restructuring of Google on October 2, 2015, and became the parent company of Google and several former Google subsidiaries. The two founders of Googl

Number of employees	103,459 (Q1 2019)

Google Search
Google Images
YouTube
Google News
Google Finance

Google Ads
Google Marketing Platform
Google Ad Manager
AdMob
Google AdSense
Google Ad Grants

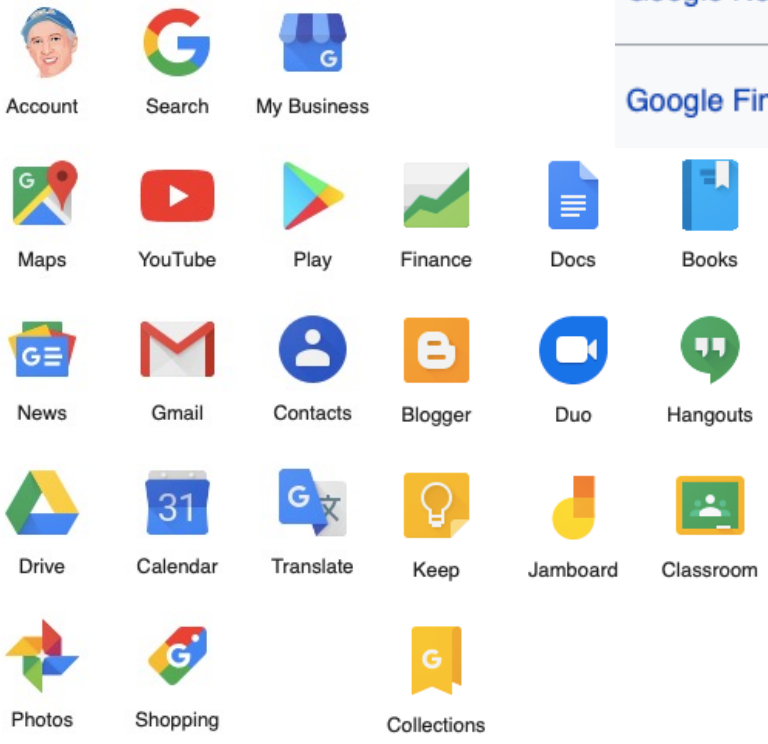

Operating systems

1. Android – Linux-base
2. Chrome OS – Linux-
3. Wear OS – version o
4. Android Auto – versi
5. Google TV – it was s
6. Android TV – versior
7. Glass OS – operatin
8. Google Fuchsia – un

Google

$137B

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

More | Tech Titan | 2019

Google's logo from 2013-2015

Then-CEO, now Chairman of Google Eric Schmidt with cofounders Sergey Brin and Larry Page (left to right) in 2008.

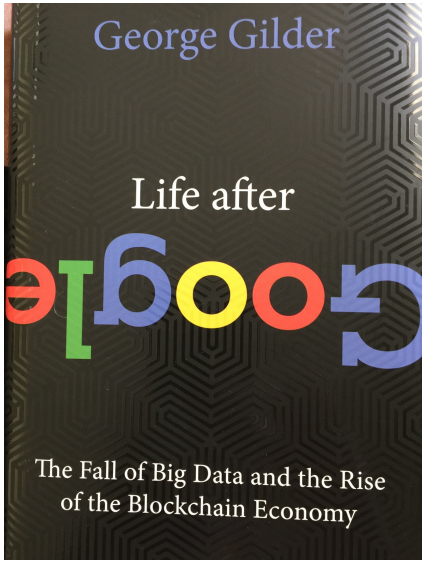

- Google started as research at Stanford by Brin and Page
  - Indexed 150 GB of Internet web pages → 15 TB (100x)
- Eric Schmidt brought "The Network" from Sun Micro to Google
- Google just built 30-acre Data Center in Oregon (at The Dalles)
  - 75,000 servers
  - 3x 10M cu. ft. warehouses
  - Connects to nearby trans-Pacific optical cable (PC-1 at 8.4 Tbps)
- Google now at "Peta/Exa scale" ($10^{15/18}$): exabytes, petaflops
- Google handles each day
  - Scores of TB of Gmail, Facebook pages, Twitter, etc.
  - 1B (10^9) YouTube videos
  - 3.5B searches → 1.5T per year
- In 2018, 1GB of hard disk costs only 2 cents (1ExB = $20M)
- Google's cloud will link to trillions of sensors around the world
  - iPhone 8 has 16 different sensors

Amazon

$233B

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Tech Titan 2019

Amazon.com, Inc.

amazon

The Amazon Spheres, part of the Amazon headquarters campus in Seattle

Founded	July 5, 1994; 25 years ago in Bellevue, Washington
Founder	Jeff Bezos
Headquarters	Seattle, Washington and Arlington, Virginia, United States

Industry	Cloud computing E-commerce Artificial intelligence Consumer electronics Digital distribution

Products	Amazon Echo · Amazon Fire · Amazon Fire TV · Amazon Fire OS · Amazon Kindle
Services	Amazon.com · Amazon Alexa · Amazon Appstore · Amazon Music · Amazon Prime · Amazon Prime Video · Amazon
Subsidiaries	A9.com · AbeBooks · Amazon Air · Alexa Internet · Amazon Books · Amazon Game Studios · Amazon Lab126 · Amazon Logistics, Inc. · Amazon Publishing · Amazon Robotics · Amazon.com Services · Amazon Studios · Audible · Body Labs · AWS · Book Depository · ComiXology · Goodreads · Graphiq · IMDb · Ring · Souq.com · Twitch Interactive · Whole Foods Market · Woot · Zappos

Revenue	▲ US$232.887 billion (2018)
Operating income	▲ US$12.421 billion (2018)
Net income	▲ US$10.073 billion (2018)

Number of employees	▲ 647,500 (2018)

Smart Speakers

Echo Dot

Add Alexa to any room

~~$49.99~~ $29.99

With Screens

NEW Echo Show 5

Compact smart display with Alexa

~~$89.99~~ $64.99

Fire Tablets

NEW Fire 7, 16 GB

Our best-selling Fire tablet now 2X the storage

$49.99

Kindle E-Readers

NEW Kindle

All-new Kindle - now with a front-light

$89.99

Doorbells

Ring Video Doorbell 2

See, hear and speak to anyone at your door

$199.00

Facebook

$56B

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Tech Titan 2019

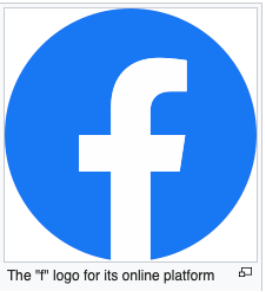

Facebook, Inc.

Founded	February 4, 2004; 15 years ago in Cambridge, Massachusetts[2]
Headquarters	1 Hacker Way (aka 1601 Willow Road) Menlo Park, California, U.S.
Employees	39,651 (June 30, 2019)[6]
Written in	C++, PHP (as HHVM),[8] D[9]
Founder(s)	Mark Zuckerberg Eduardo Saverin Andrew McCollum Dustin Moskovitz Chris Hughes
Revenue	▲ US$ 55.838 billion (2018)[4]
Operating income	▲ US$ 24.913 billion (2018)[4]
Net income	▲ US$ 22.111 billion (2018)[4]
Users	▲ 2.3 billion monthly active users (December 2018)

Products
Messenger
Watch
Portal

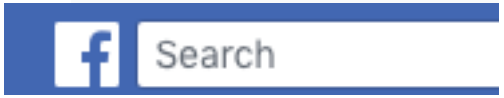

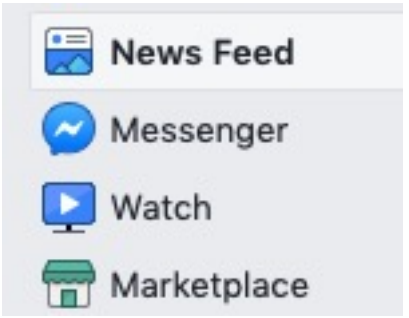

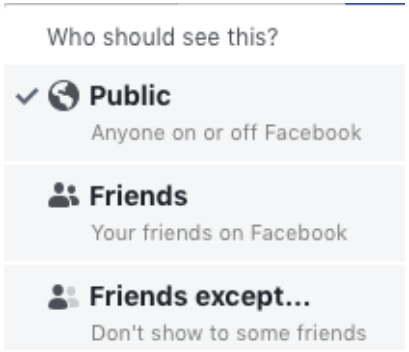

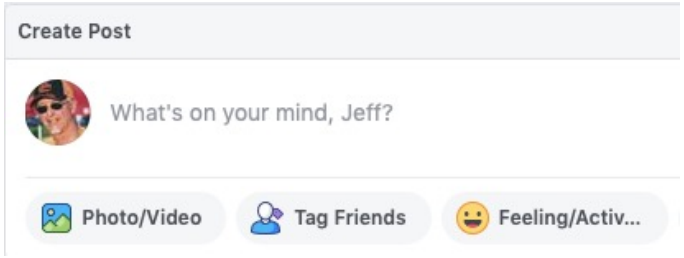

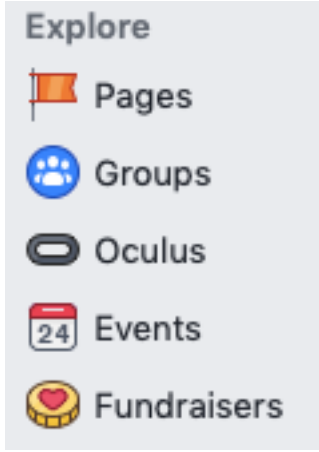

Facebook

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Server Farm

FACEBOOK says that since the company is the only entity gathering and monetizing personal data within its system, it does not engage in sales of user data to third parties. Above, a Facebook server room is seen in 2013.

JONATHAN NACKSTRAND AFP/Getty Images

Bitcoin Mine

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Server Farm

Corporate Brands

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

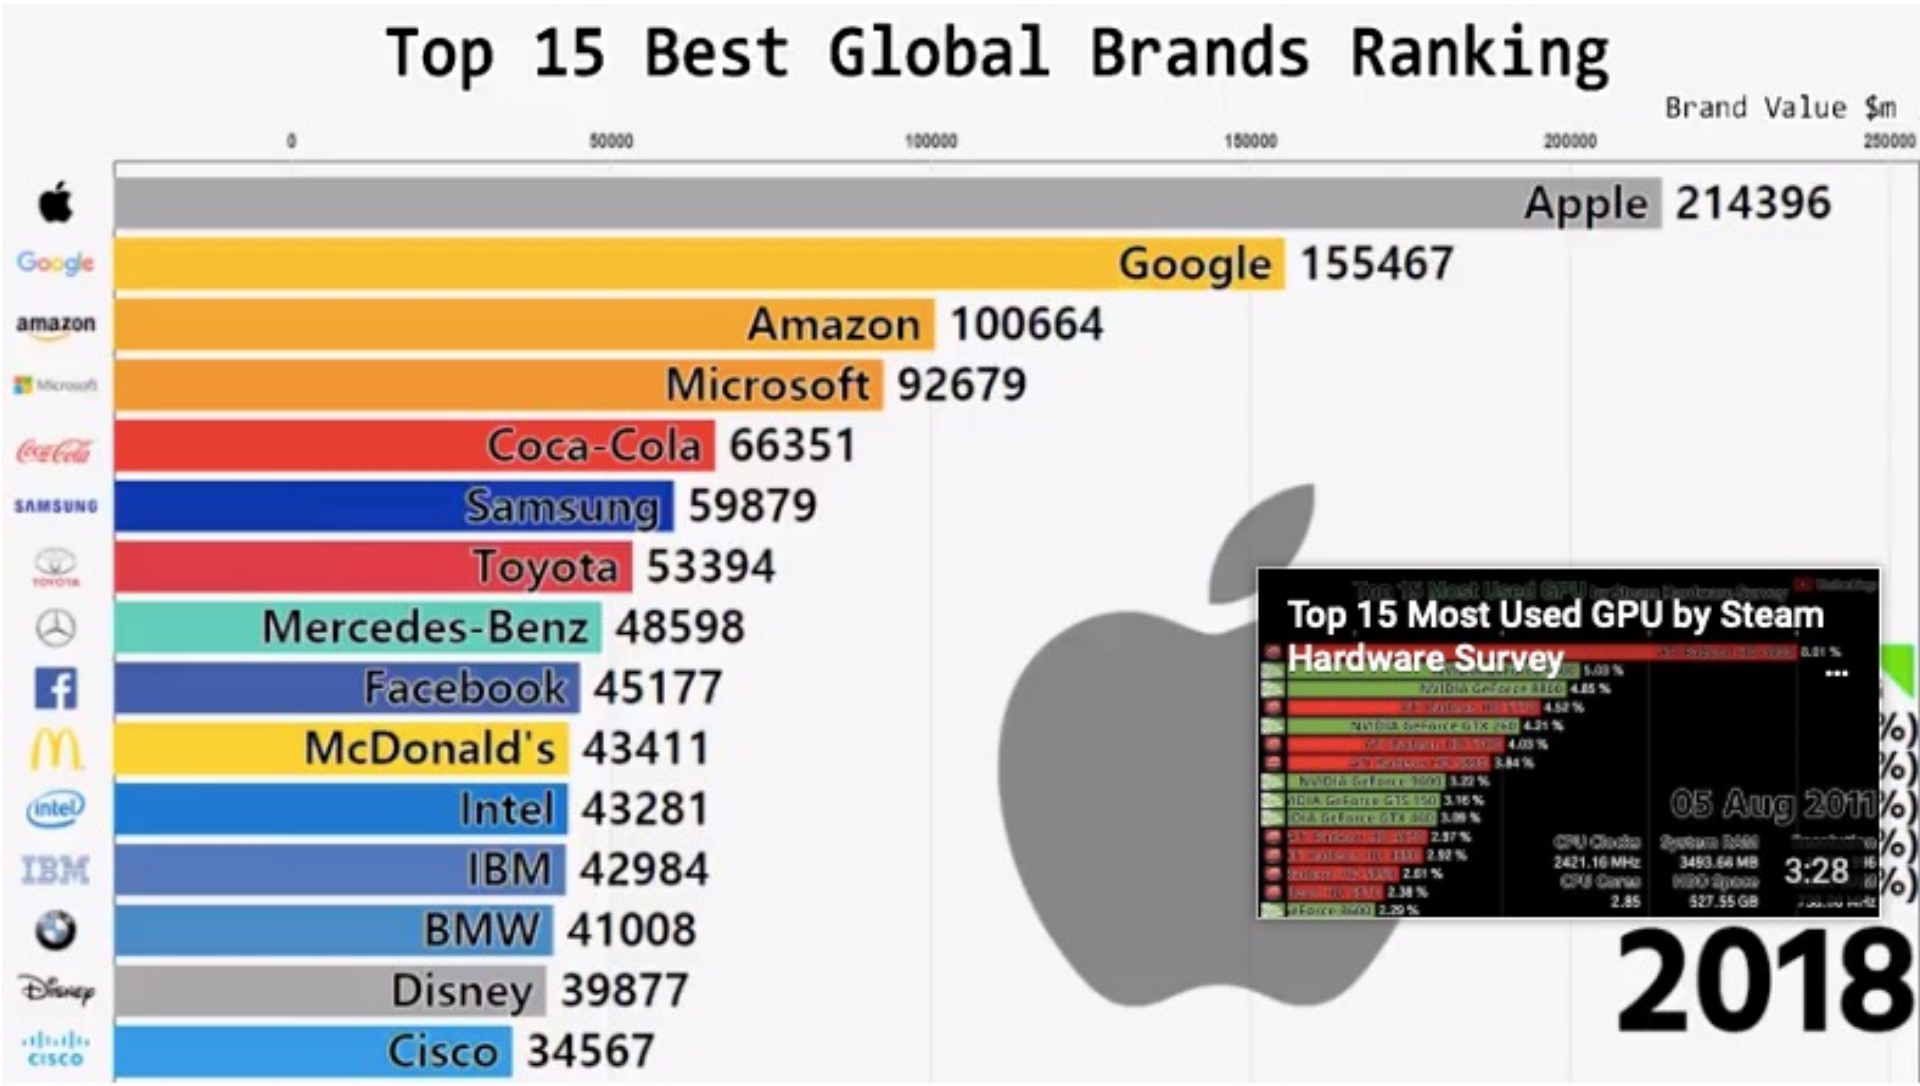

Top 15 Best Global Brands Ranking (2000-2018)

World Economy

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

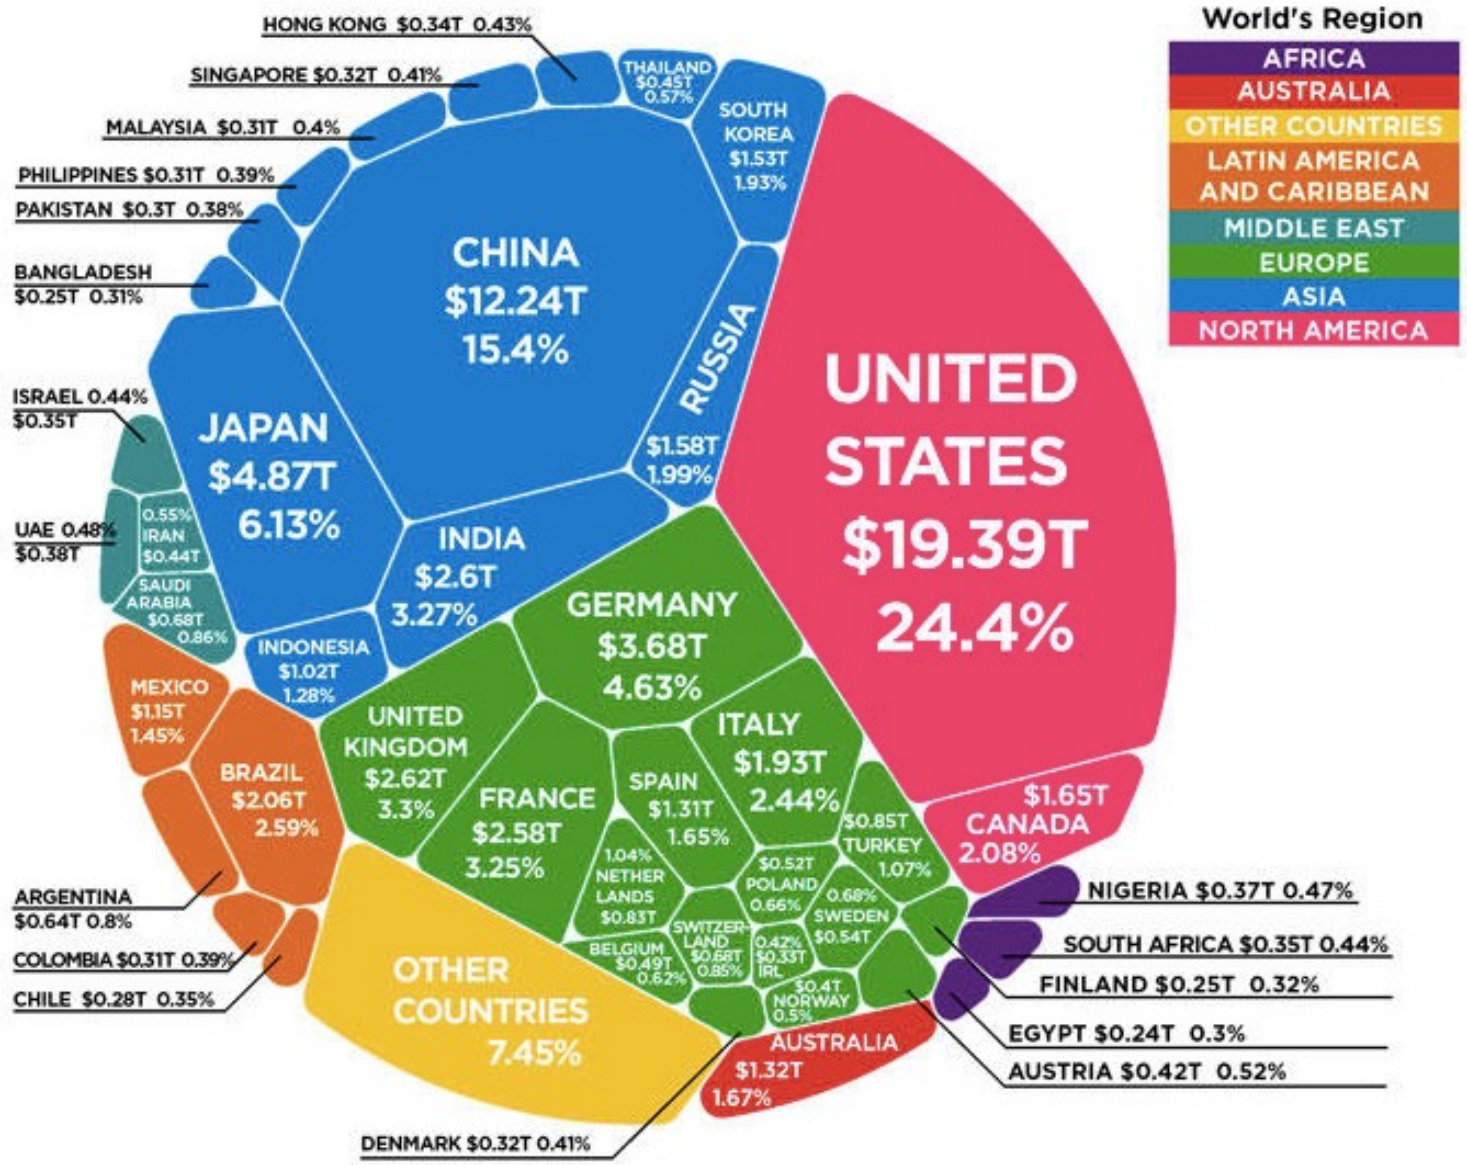

WW R&D

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

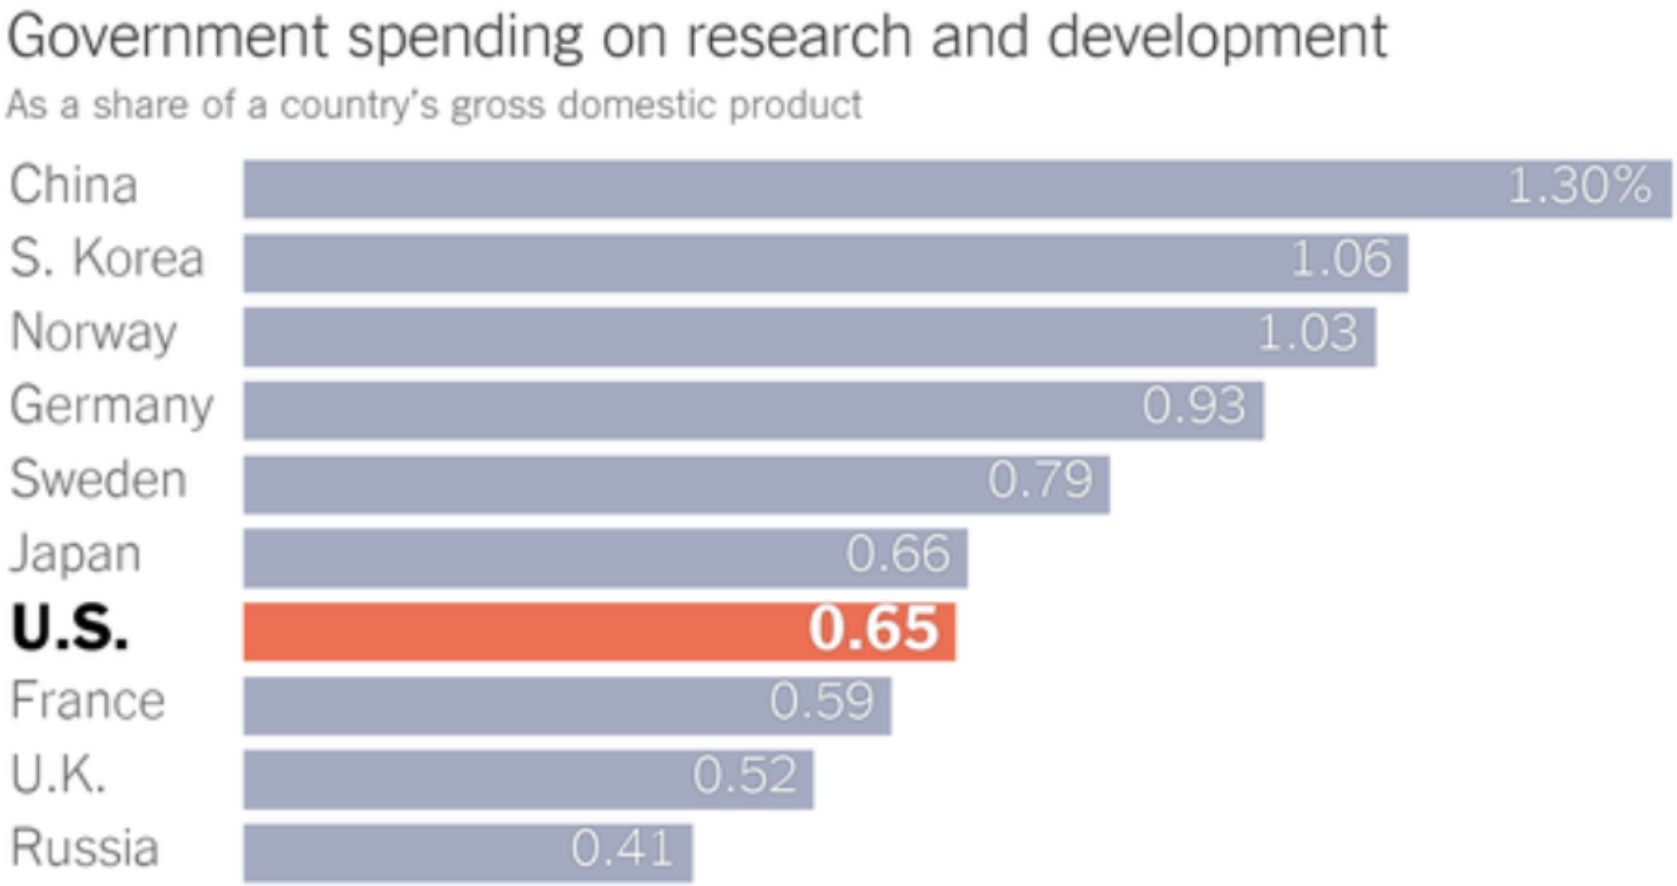

Tech Titan Founders/CEOs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

2020

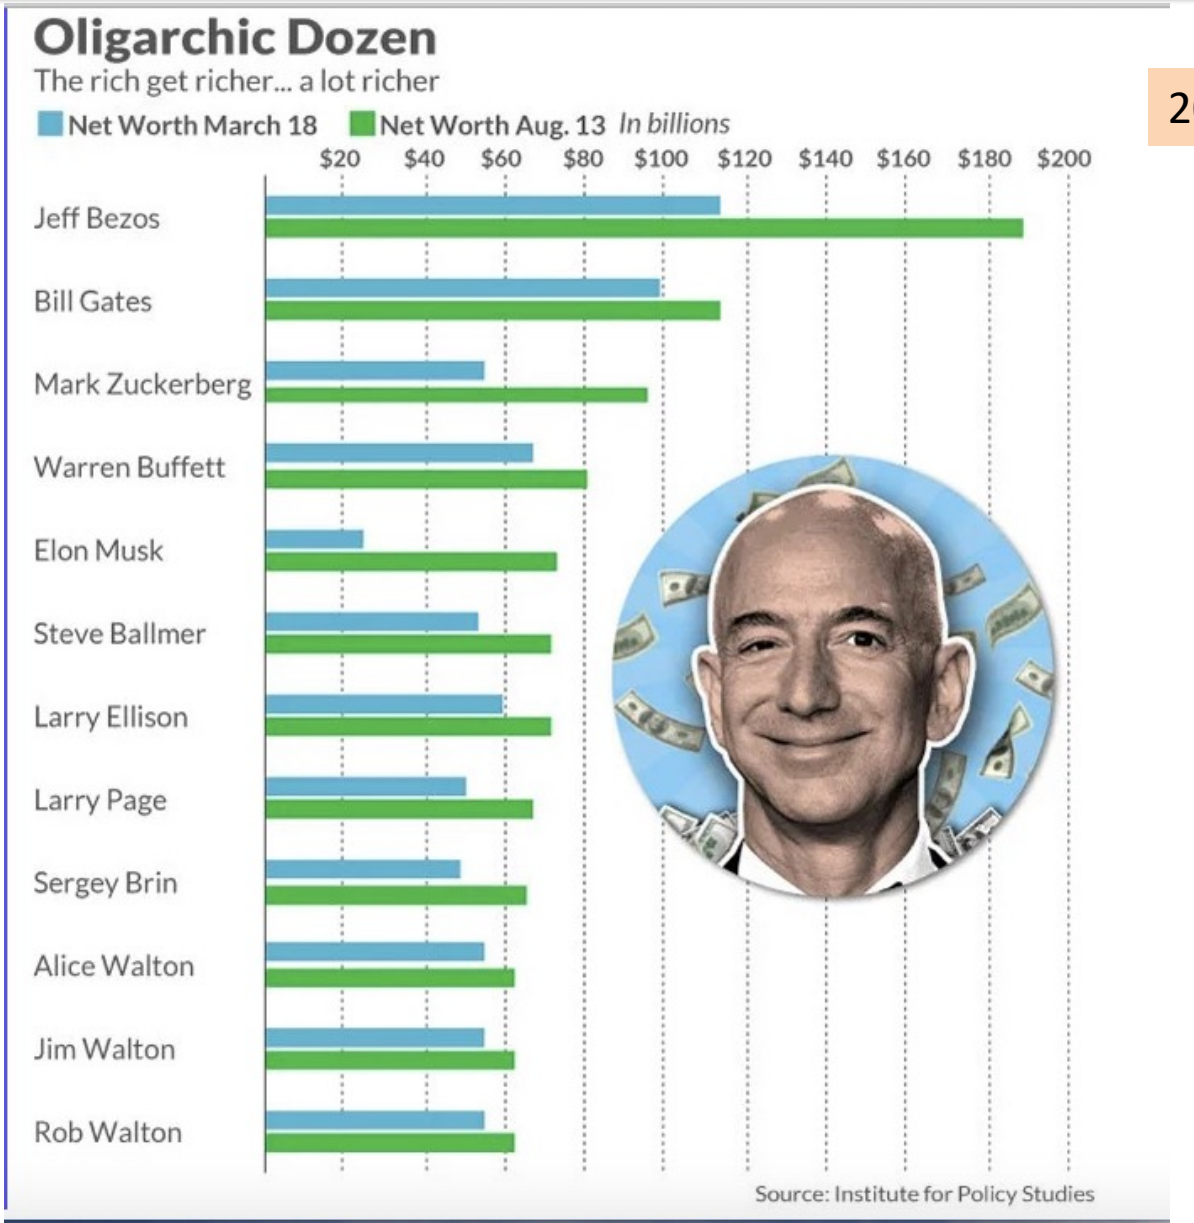

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Mainframes

1st Computer?

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quora

ABC

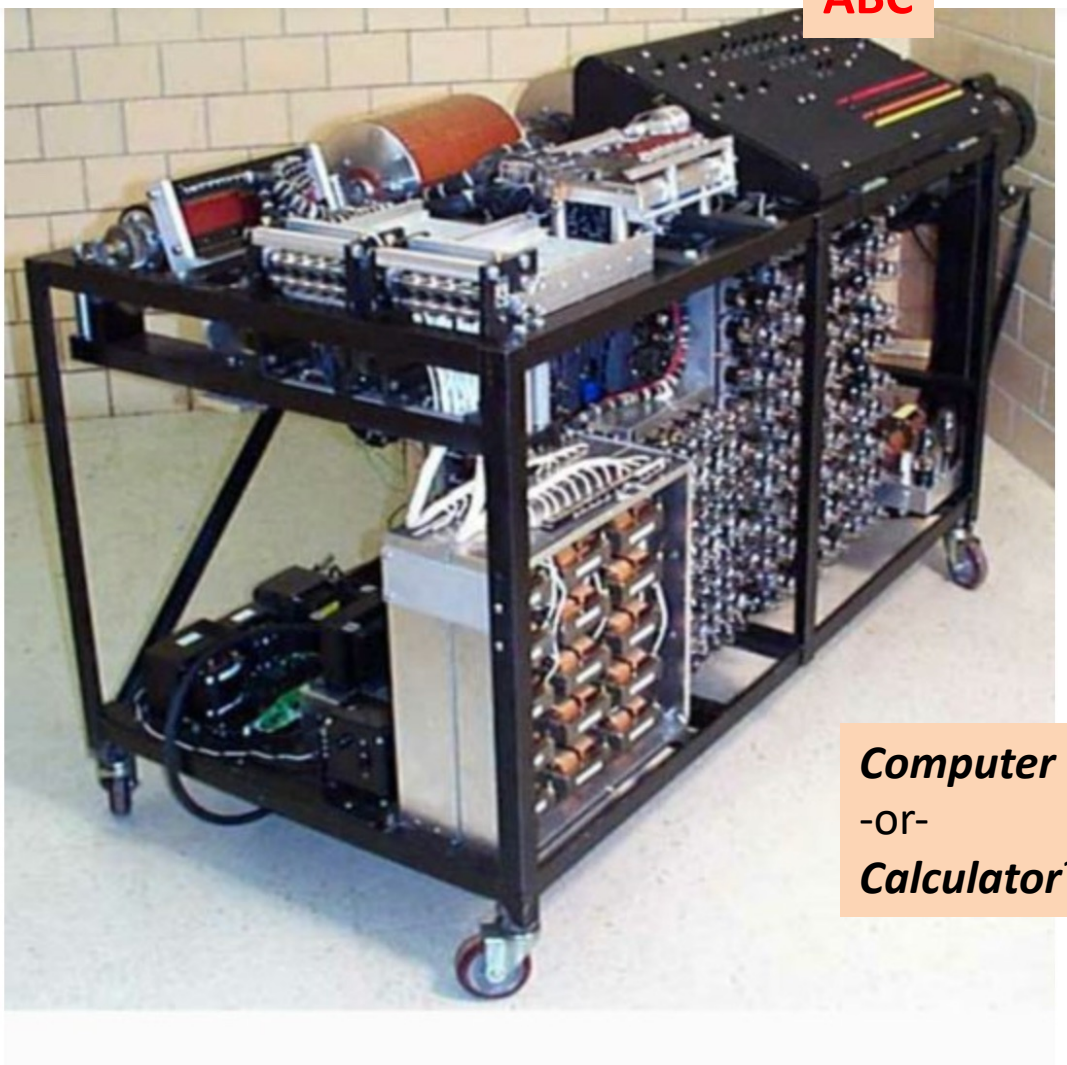

Computer -or- ***Calculator***?

The ABC weighed over 700 pounds and used vacuum tubes. It had a rotating drum, a little bigger than a paint can, that had small capacitors on it. A capacitor is device that can store an electric charge, like a battery.

Atanasoff–Berry computer (ABC)

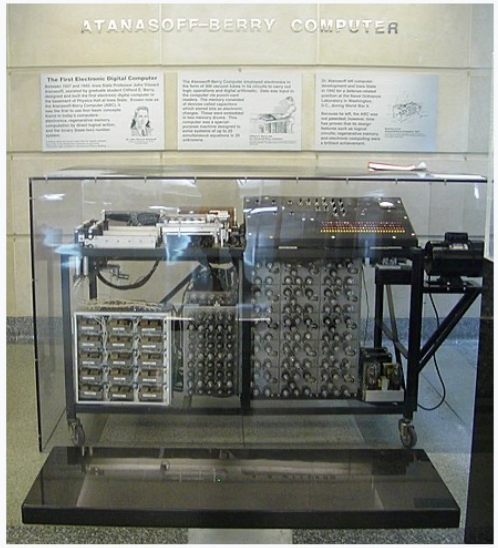

Atanasoff–Berry computer replica at Durham Center, Iowa State University

Developer	John Vincent Atanasoff with help of graduate student Clifford Berry
Release date	1942; 81 years ago
Units sold	1
CPU	More than 300 vacuum tubes @ 60 Hz
Memory	3000 bits
Display	Decimal, via a front panel display
Input	Decimal, via standard IBM 80-column punched cards
Mass	700 pounds (320 kg)

1st Computer?

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Wikipedia

ABC

The **Atanasoff–Berry computer** (**ABC**) was the first automatic electronic digital computer.[1] Limited by the technology of the day, and execution, the device has remained somewhat obscure. The ABC's priority is debated among historians of computer technology, because it was neither programmable, nor Turing-complete.[2] Conventionally, the ABC would be considered the first electronic ALU (arithmetic logic unit) – which is integrated into every modern processor's design.
Its unique contribution was to make computing faster by being the first to use vacuum tubes to do the arithmetic calculations. Prior to this, slower electro-mechanical methods were used by Konrad Zuse's Z1 computer, and the simultaneously developed Harvard Mark I. The first electronic, programmable, digital machine,[3] the Colossus computer from 1943 to 1945, used similar tube-based technology as ABC.

First Computer?

Quora

ABC or ENIAC?

Tom Crosley, M.S. Computer Science, Illinois Institute of Technology Chicago - Illinois Tech

Answered Sep 18

When was the first electronic computer ever produced?

Originally Answered: What was the first electronic computer?

The Atanasoff–Berry computer (ABC) was the first automatic electronic digital computer. (It was not the first *general-purpose programmable* electronic computer; that honor goes to ENIAC as described in another answer.)

Conceived in 1937, the machine was built by Iowa State College (now Iowa State University, my undergraduate alma mater) mathematics and physics professor John Vincent Atanasoff with the help of graduate student Clifford Berry. It was first successfully tested in 1942.

The ABC computer was built for a specific purpose; the solution of systems of simultaneous linear equations. It could handle systems with up to twenty-nine equations.

First Computer?

Quora

ENIAC?

The machine contained about 300 vacuum tubes. The memory of the Atanasoff–Berry Computer was a system called regenerative capacitor memory, a forerunner of DRAM but which used a drum.

Eckert & Mauchly

On June 26, 1947, J. Presper Eckert and John Mauchly were the first to file for patent on a digital computing device (ENIAC), much to the surprise of Atanasoff. The ABC had been examined by John Mauchly in June 1941, and Isaac Auerbach, a former student of Mauchly's.

The ENIAC patent did not issue until 1964. In 1967 Honeywell sued Sperry Rand in an attempt to break the ENIAC patents, arguing the ABC constituted prior art. The case was legally resolved on October 19, 1973, when U.S. District Judge Earl R. Larson held the ENIAC patent invalid, ruling that the ENIAC derived many basic ideas from the Atanasoff–Berry Computer.

ENIAC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1944

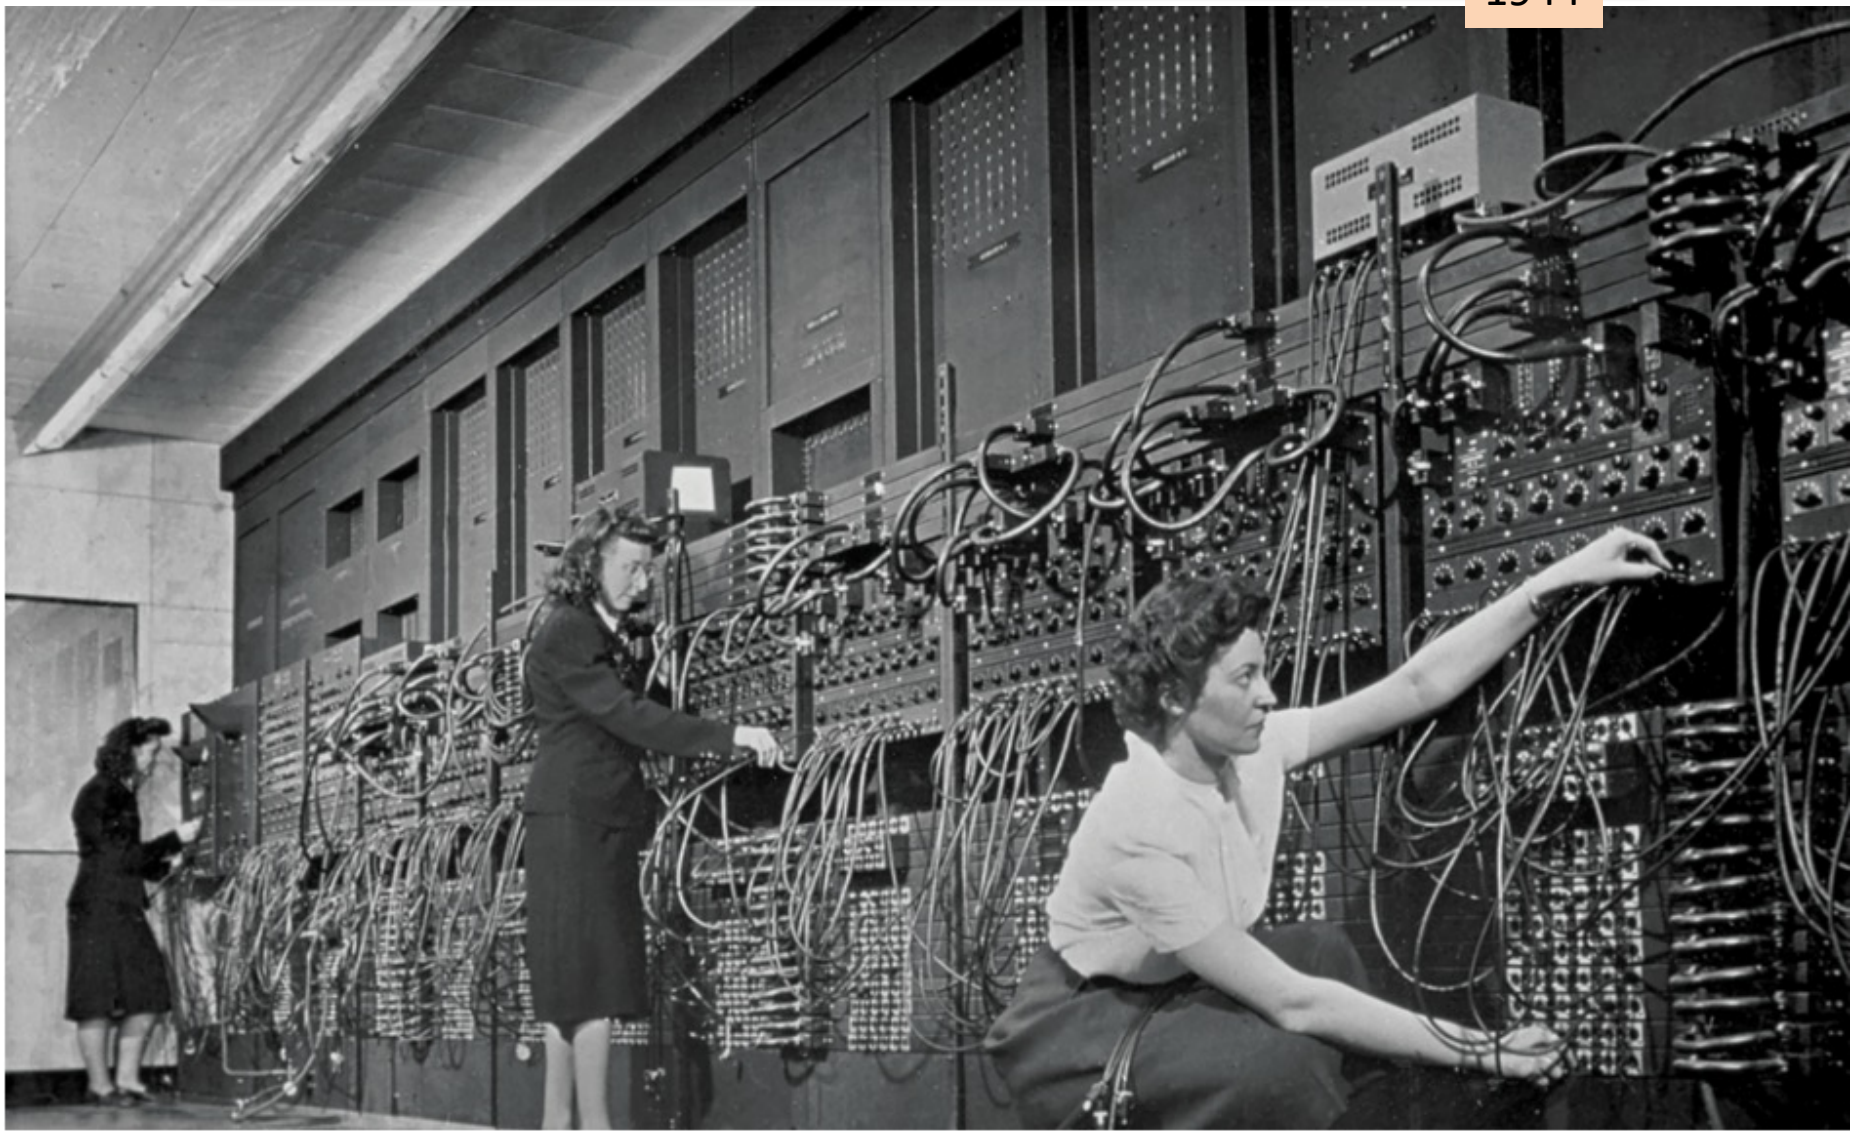

Computer operators with an Eniac — the world's first programmable general-purpose computer.
Corbis/Getty Images

ENIAC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quora

1944

Andrew Silverman, Long-time Microsoft employee, tinkerer, maker, and computer history buff.

Updated Aug 12, 2021

How much memory did the ENIAC computer have?

ENIAC's architecture was rather different from what we see in modern computers but the shortcomings of this machine were precisely what steered Von Neumann, Eckert, and Mauchly towards the combined instruction/data memory architectures that followed.

Strictly speaking, ENIAC had 1248 decimal digits and 208 signs (pos/neg) of "function tables" (essentially read-only numerical data switches that were usable as either program instructions or data in later configurations), and a mere 200 decimal digits of writable data, spread across about twenty individual accumulators. This lack of sufficient memory and the enormous number of tubes required implementing each digit eventually led Eckert to design the mercury delay-line memory circuit. Alongside cathode ray-tube memory, which was invented around the same time, these newer options drastically increased storage capacity and lowered costs in the designs that followed ENIAC.

ENIAC & Colossus

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

ENIAC was the first Turing-complete electronic device, and performed ballistics trajectory calculations for the United States Army.[93]

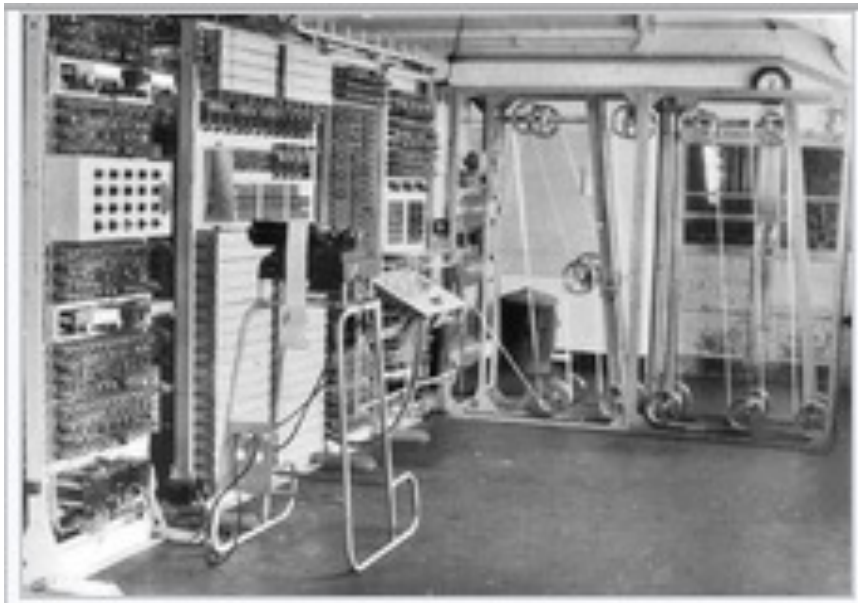

Wartime photo of Colossus No. 10

1st Gen Mainframes

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1944-50

ENIAC

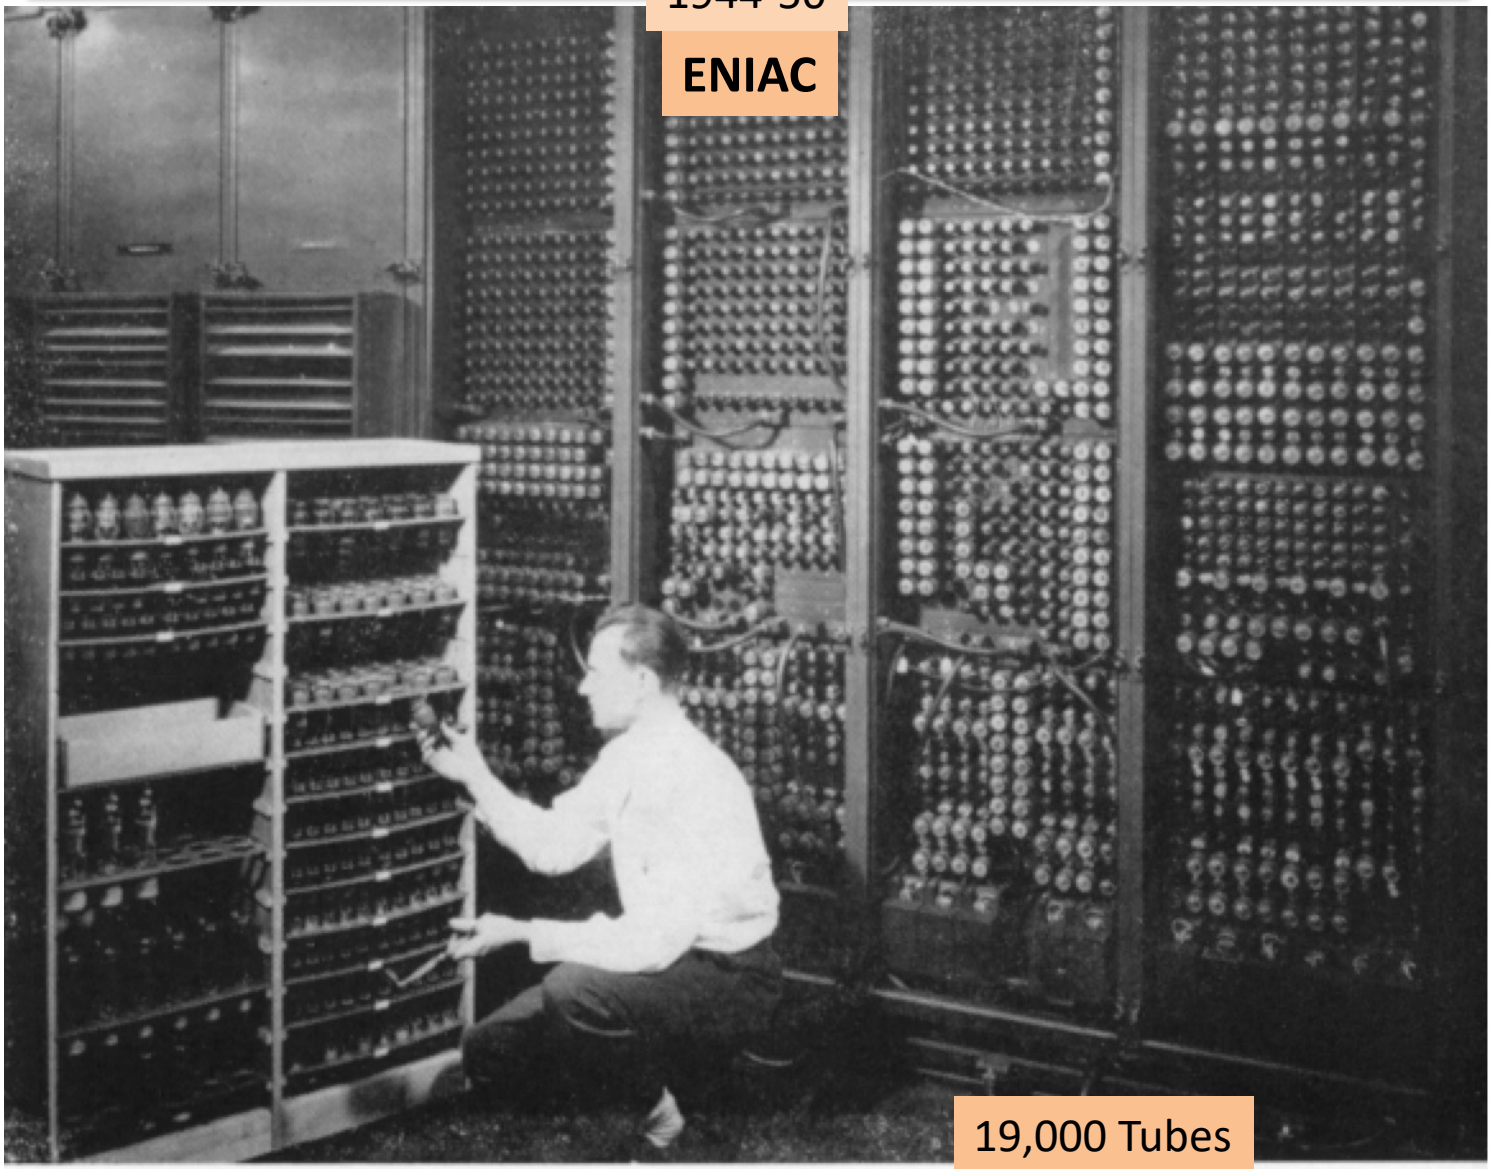

Replacing a bad tube meant checking among ENIAC's 19,000 possibilities.

1st Gen Mainframes

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1950-55

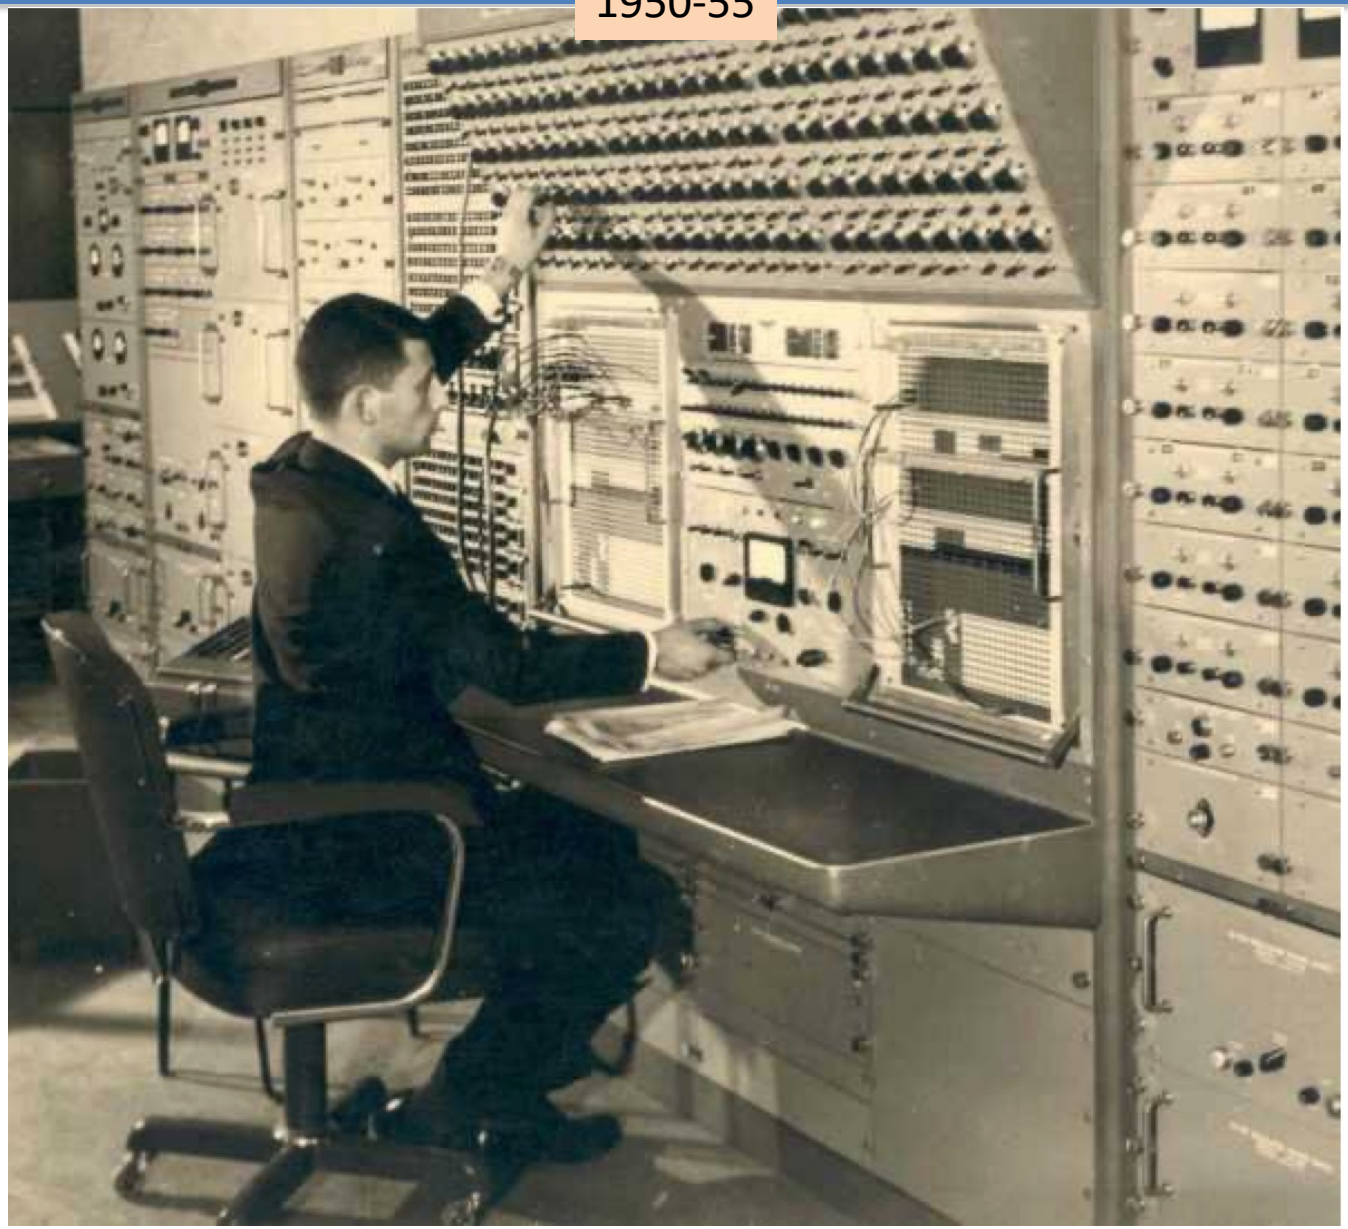

1st Gen Mainframes

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1950-55

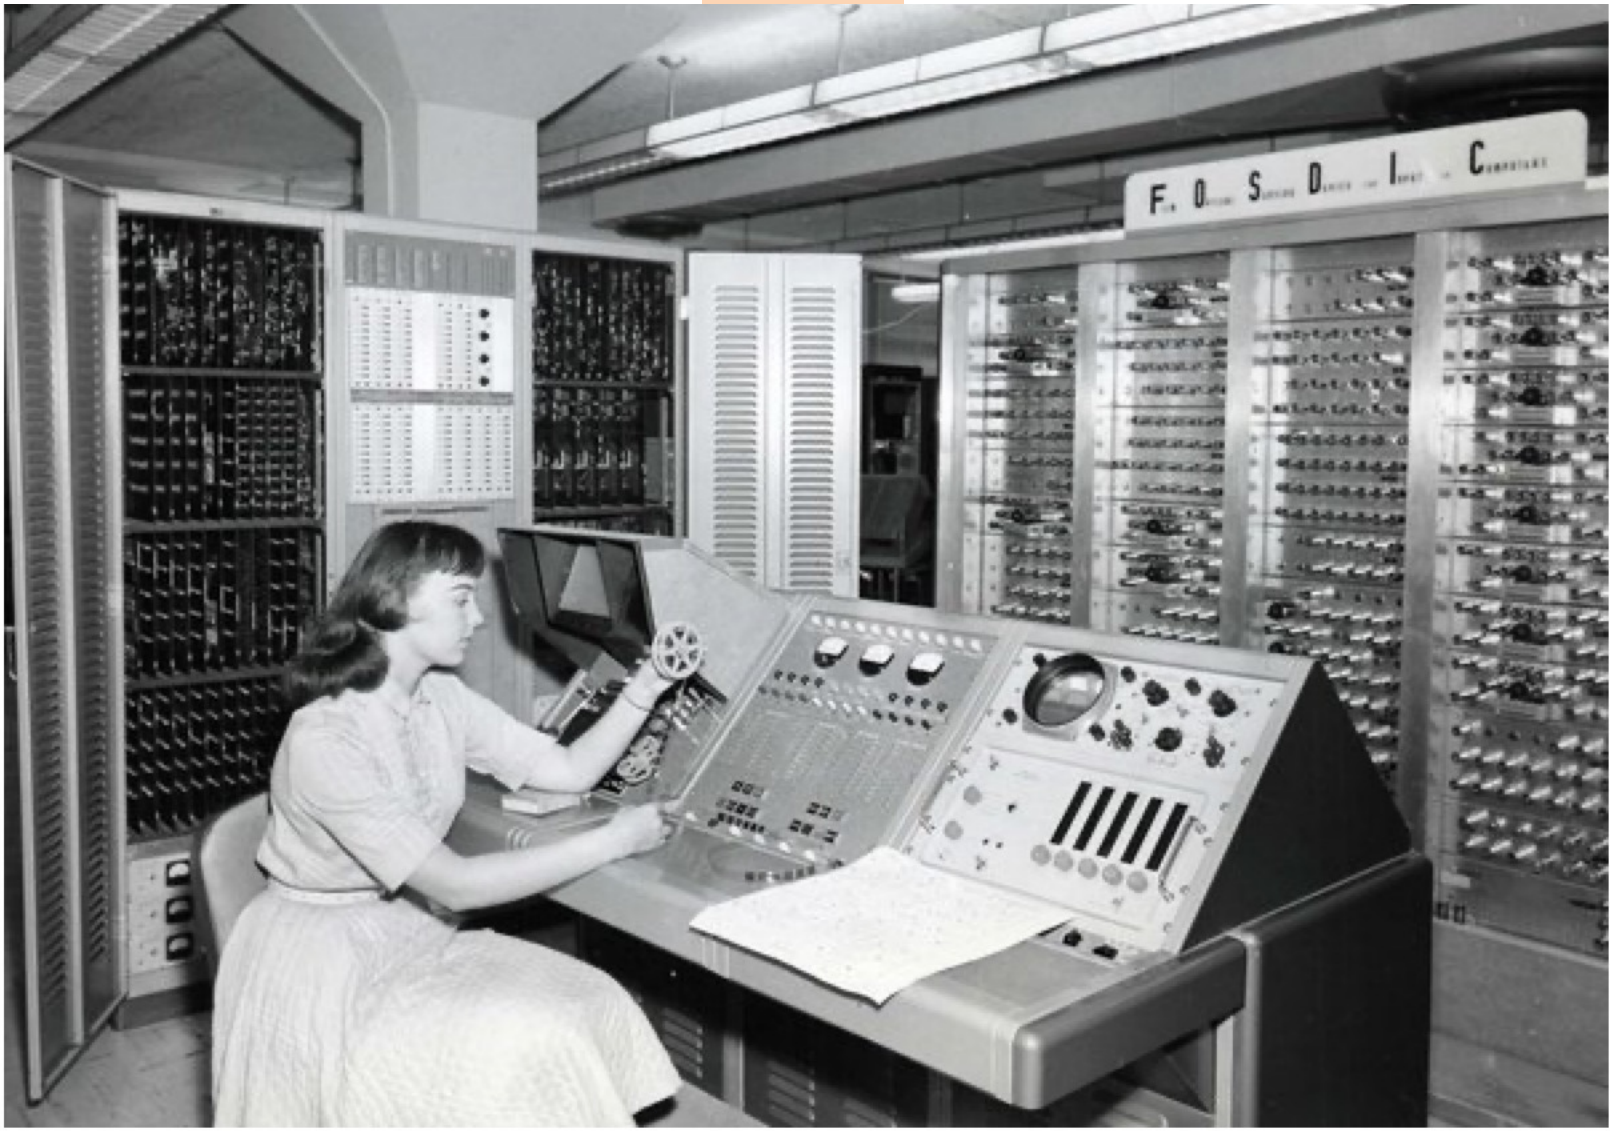

Univac 1103

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1952

An E.R.A./Univac 1103 computer in the 1950s.
Hum Images/Alamy

IBM 701

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1952

IBM 701 processor frame, showing 1071 of the vacuum tubes

Vacuum tube logic module from a 700 series IBM computer.

Williams tube from an IBM 701 at the Computer History Museum

IBM 704

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1st with FP

1954

An IBM 704 computer, with IBM 727 tape drives and IBM 780 CRT display

IBM 1401

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1959

The IBM 1401 "mainframe".

IBM 1401

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1959

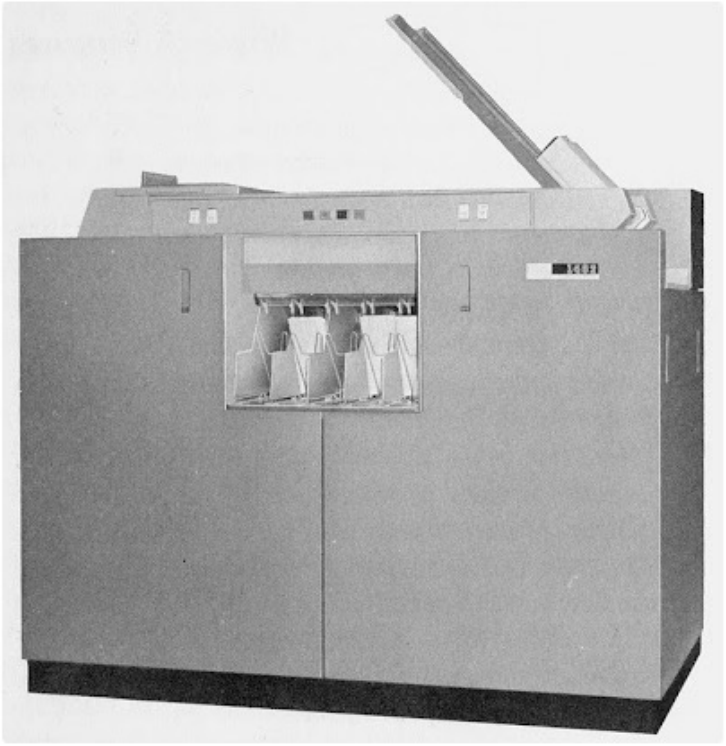

1402 Card reader/punch

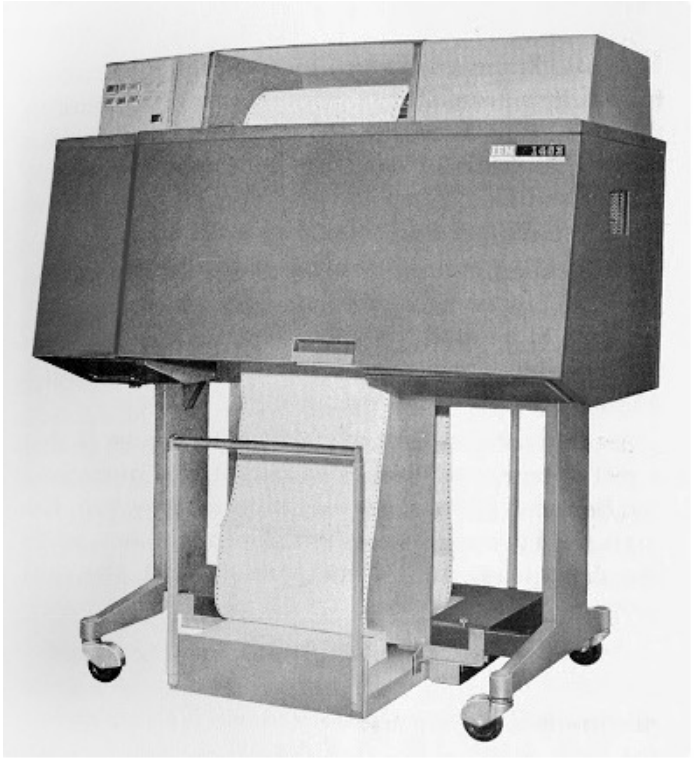

1403 Printer

IBM 1401

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1959

IBM 729 Tape Drives — with vacuum tape loop buffering

IBM 7090

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1962

Dual 7090s at NASA during Project Mercury.

IBM 7074

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1964

IBM 7094

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1965

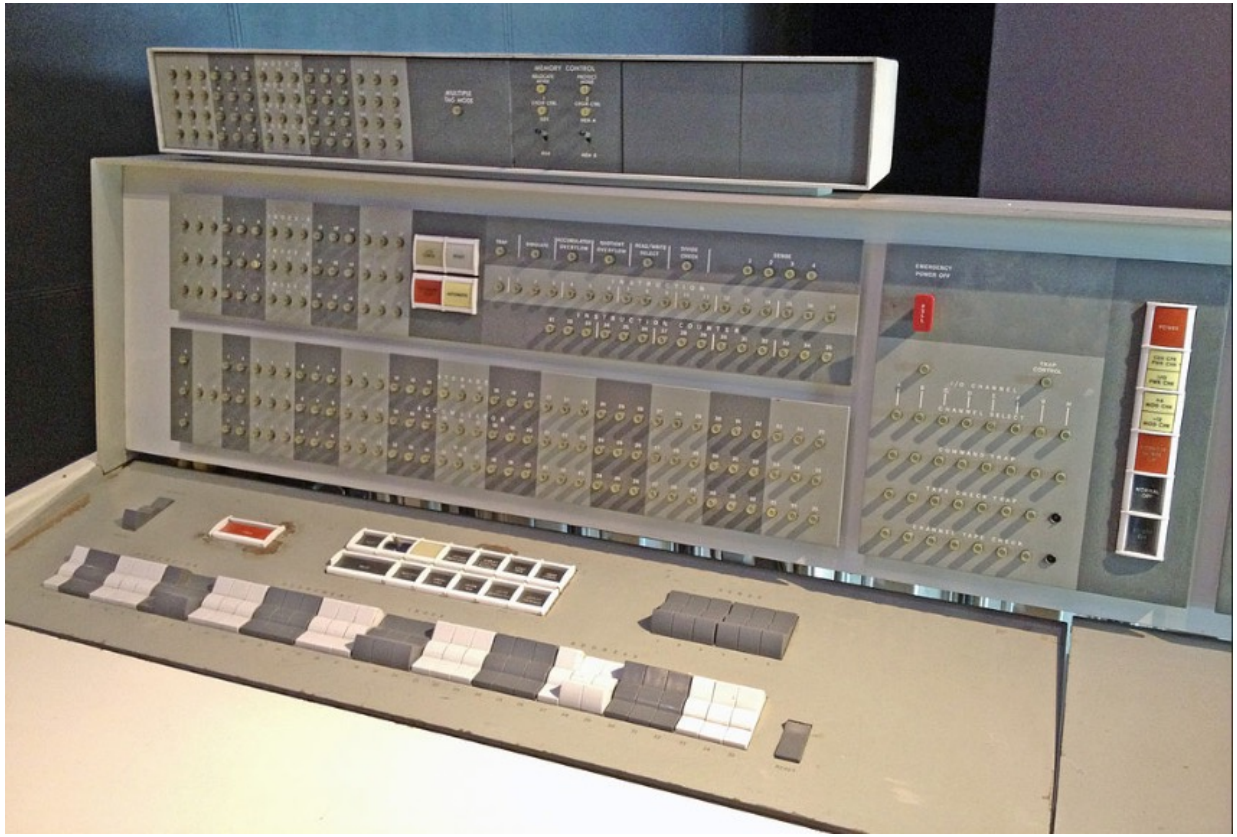

IBM 7094 operator's console showing additional index register displays in a distinctive extra box on top. Note "Multiple Tag Mode" light in the top center.

IBM System 360

1965

System/360 Model 65 operator's console, with register value lamps and toggle switches (middle of picture) and "emergency pull" switch (upper right)

IBM System/360 Model 50 CPU, computer operator's console, and peripherals at Volkswagen

IBM System 360

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1965

IBM System 360 Model 91 operator's console at NASA, sometime in the late 1960s.

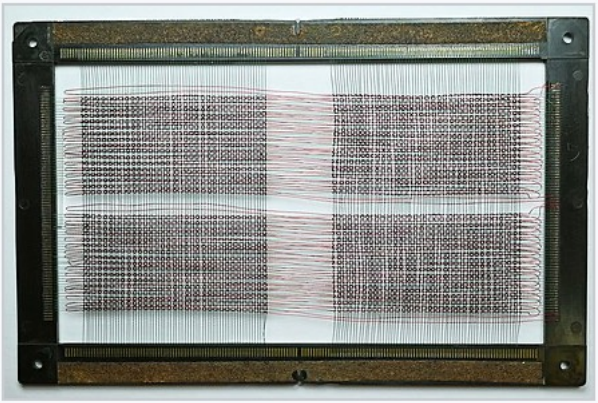

Magnetic core memory, probably from a 360

UCLA had 2 of these Model 91's (out of only 3 built) 1968-80

Wikipedia

Quora

Drum Storage

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

But did you know that the concept of regenerative capacitive memory was invented in the early 1940s? The very first electronic digital computer, the ABC (Atanasoff–Berry computer), had a pair of drums, each containing 1600 capacitors that rotated on a common shaft once per second. The capacitors on each drum were organized into 32 "bands" of 50, so it could store 32 50-bit floating point numbers.

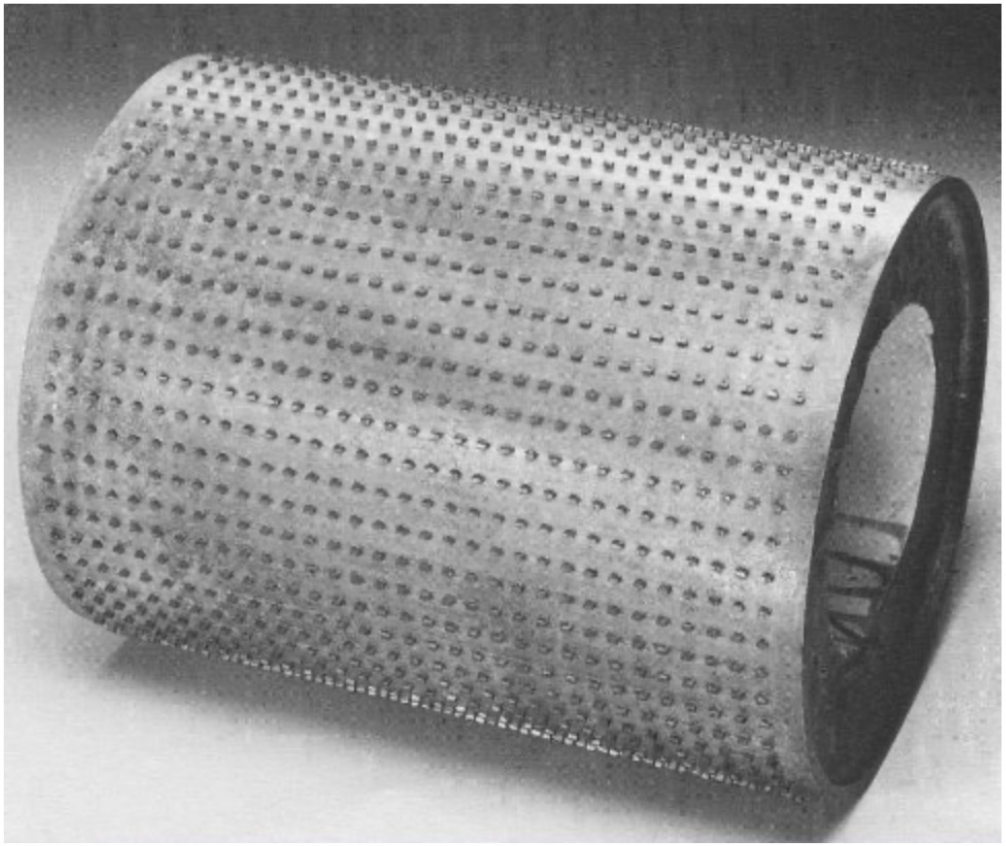

Picture from ABC Computer - Notes

Each of the little squares is connected to a 0.0015 μF paper capacitor

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

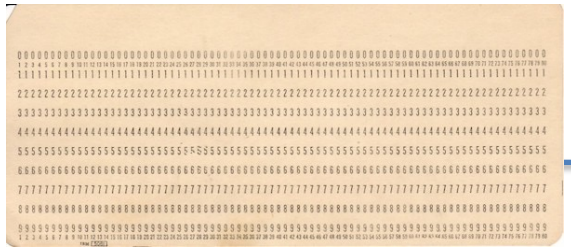

ENIAC

DEC PDP/VAX

44. ENIAC, the first computer ever built

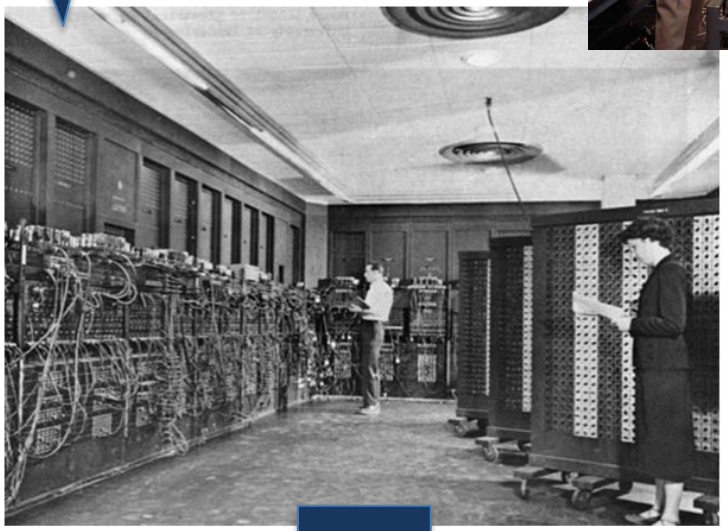

Apple II

Apple Mac

Google TPU

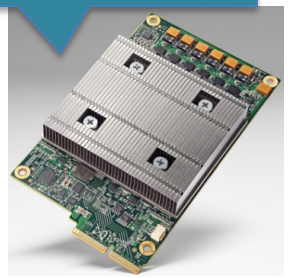

Data Center

CDC

IBM 360

IBM PC

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

TIMELINE

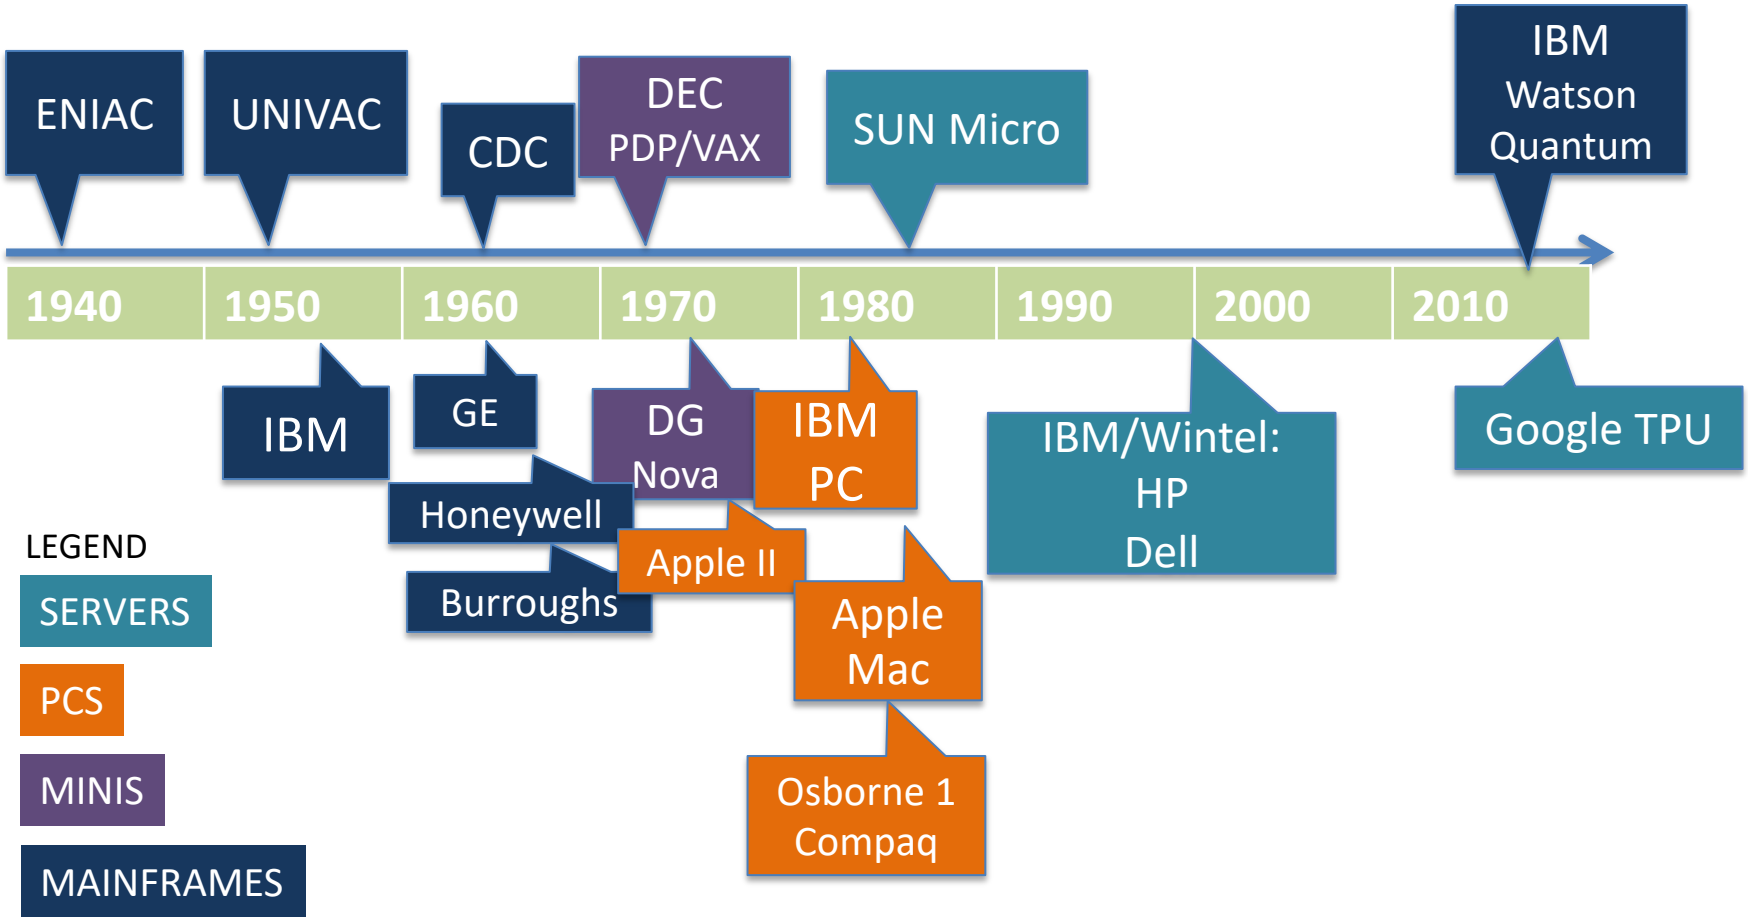

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Minicomputers of the 1970's
Used Am2900 *bit-slices*

TIMELINE

Computers of the 1990's
Used RISC CPU's

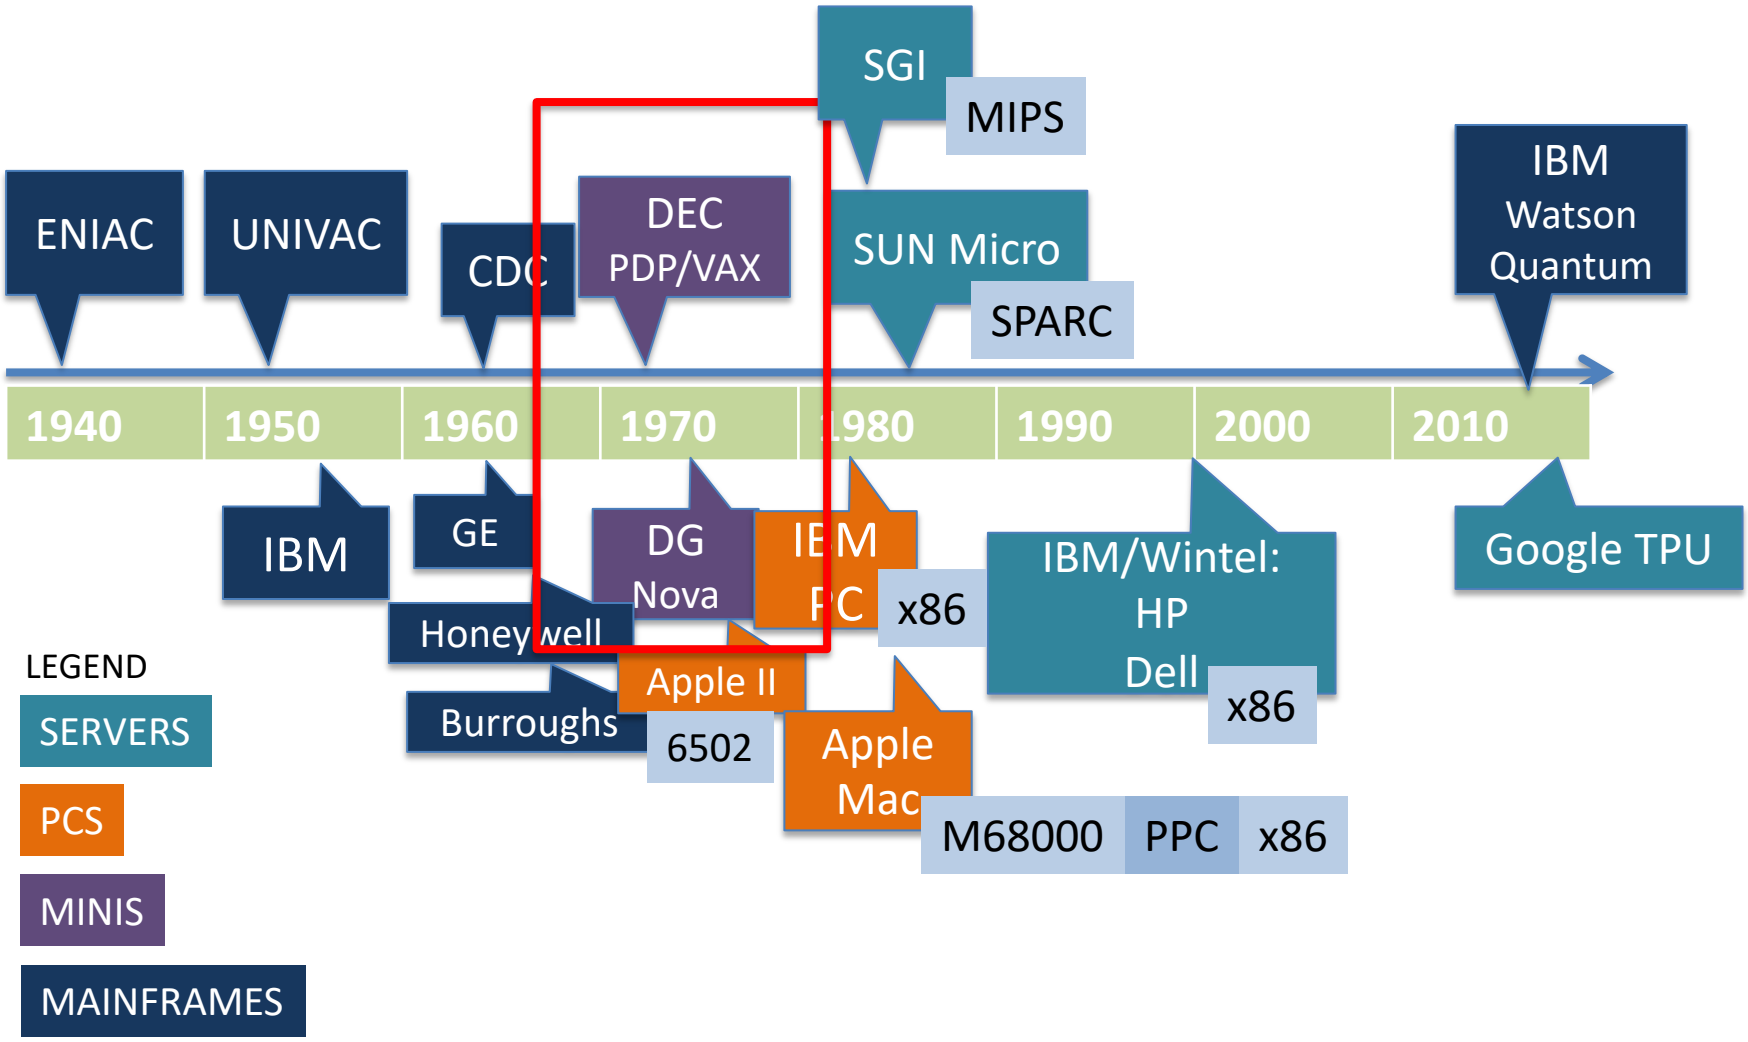

Computer Generations

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

TIMELINE

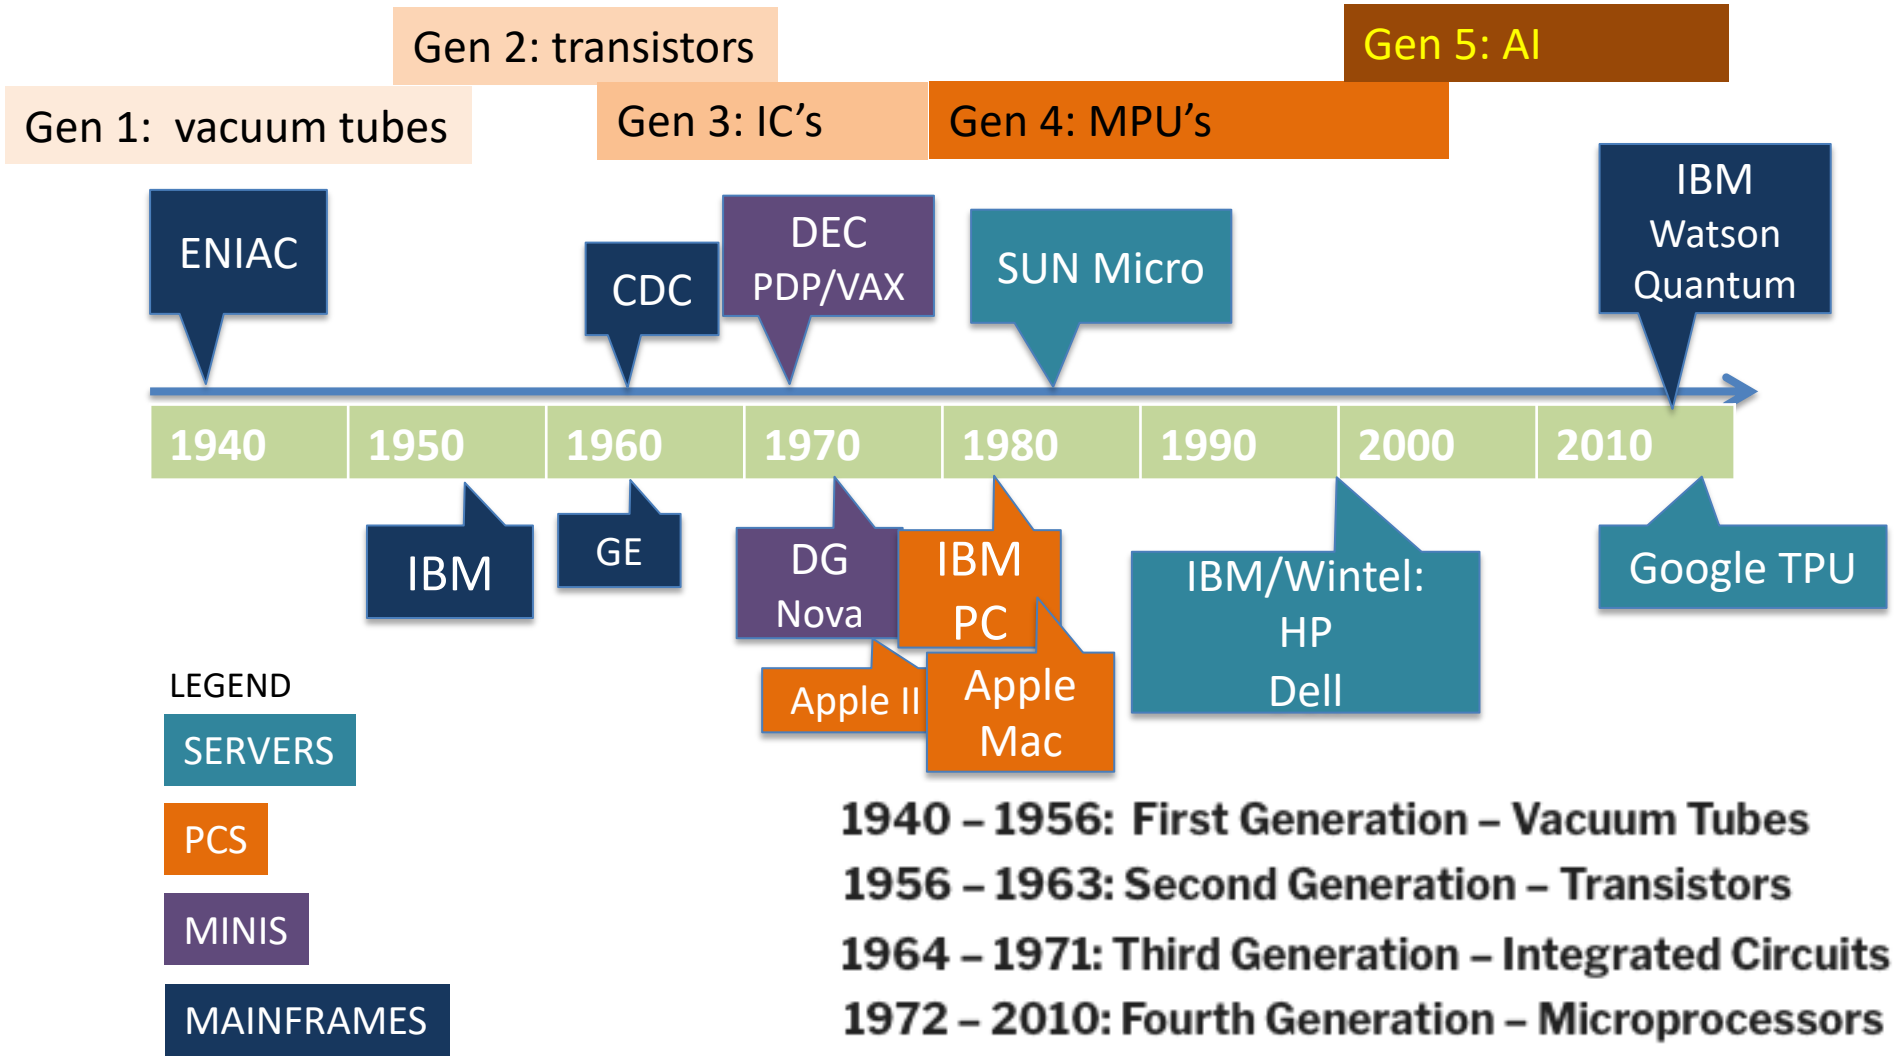

1940 – 1956: First Generation – Vacuum Tubes
1956 – 1963: Second Generation – Transistors
1964 – 1971: Third Generation – Integrated Circuits
1972 – 2010: Fourth Generation – Microprocessors
2010- : Fifth Generation – Artificial Intelligence

1st Gen Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

von Neumann

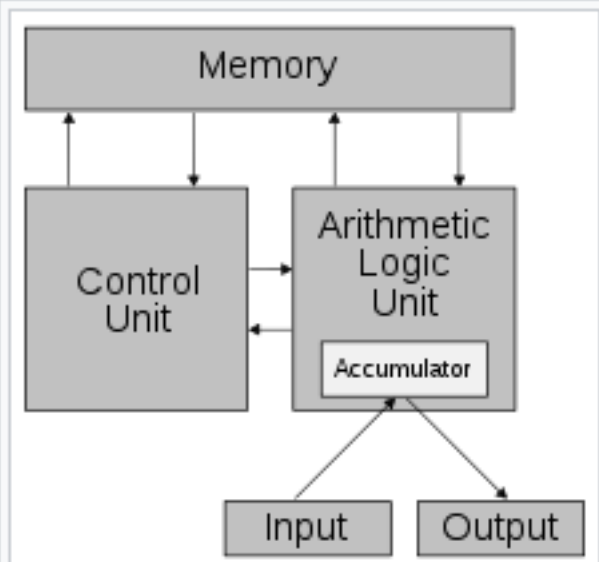

Design of the von Neumann architecture, 1947

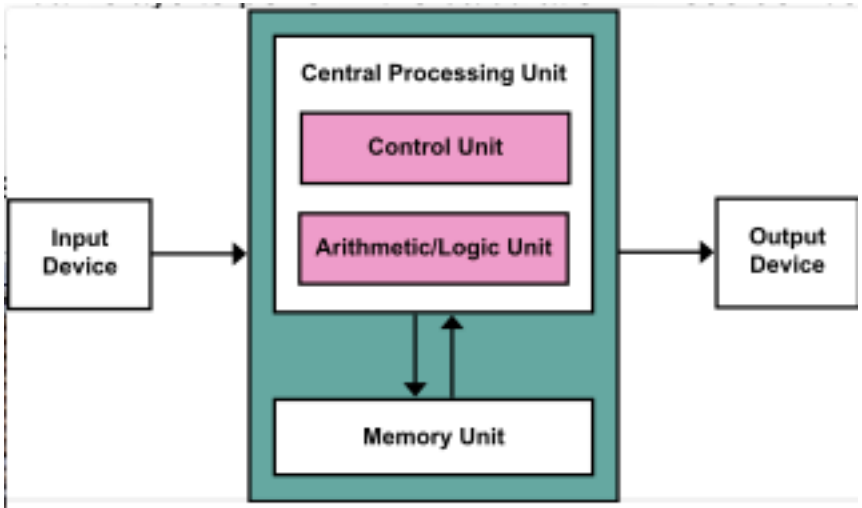

The **von Neumann architecture**—also known as the **von Neumann model** or **Princeton architecture**—is a computer architecture based on a 1945 description by Hungarian-American mathematician and physicist John von Neumann and others in the *First Draft of a Report on the EDVAC*. Th

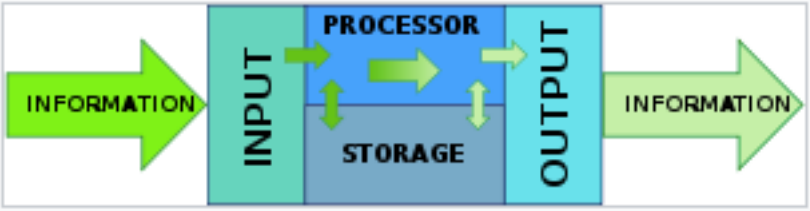

Computing hardware is a platform for information processing.

Old Computer Tech

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Vacuum Tubes

Richard Urwin, Software engineer for 30+ years

Answered 9h ago

There were no very small components in the earliest computers. They used vacuum tubes.

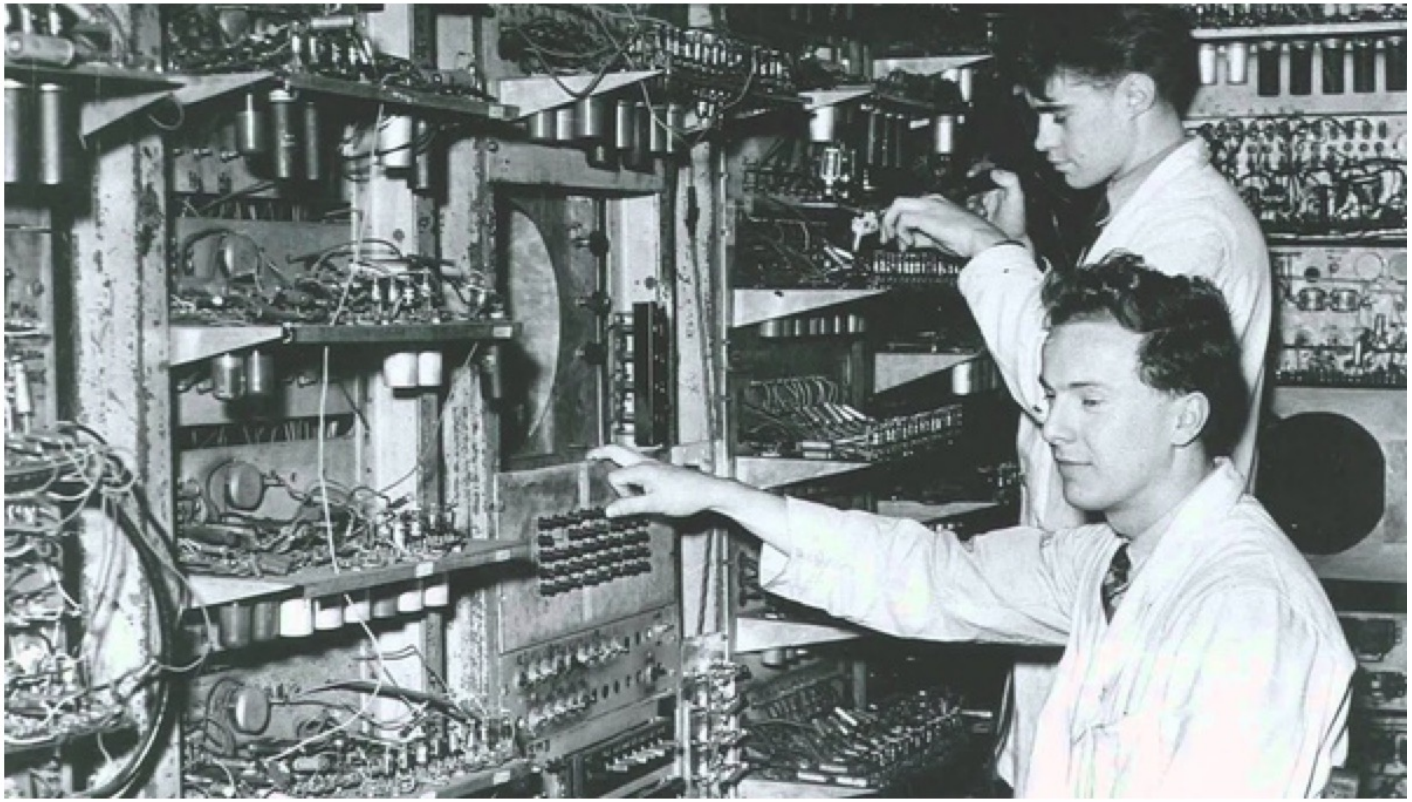

Each of the silver cylinders in the picture is a vacuum tube, approximately equivalent to a transistor. They were about an inch (2.5cm) across and two inches (5cm) high.

Old Computer Tech

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Discretes R/T

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

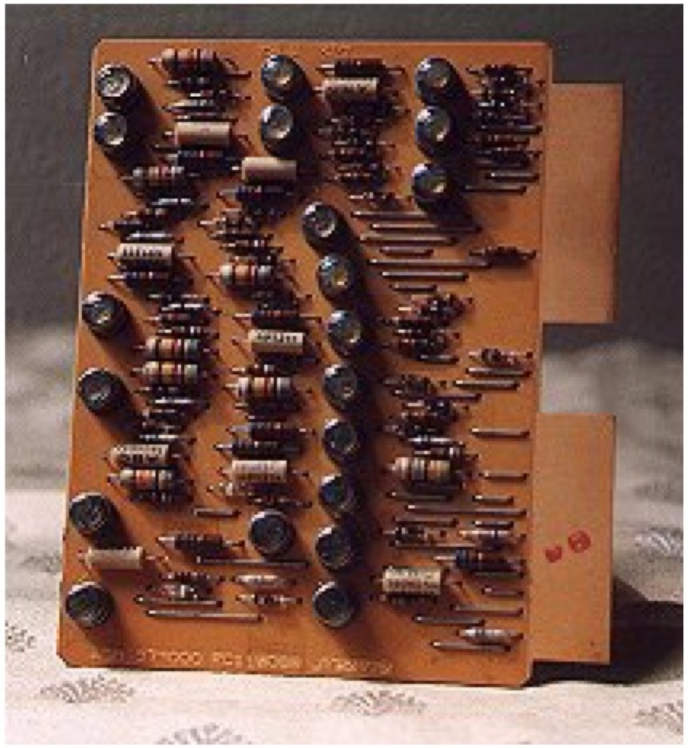

There are no small components there. Those horizontal cylinders are maybe half an inch (12mm) long.

Old Computer Tech

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

MSI/LSI IC's

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

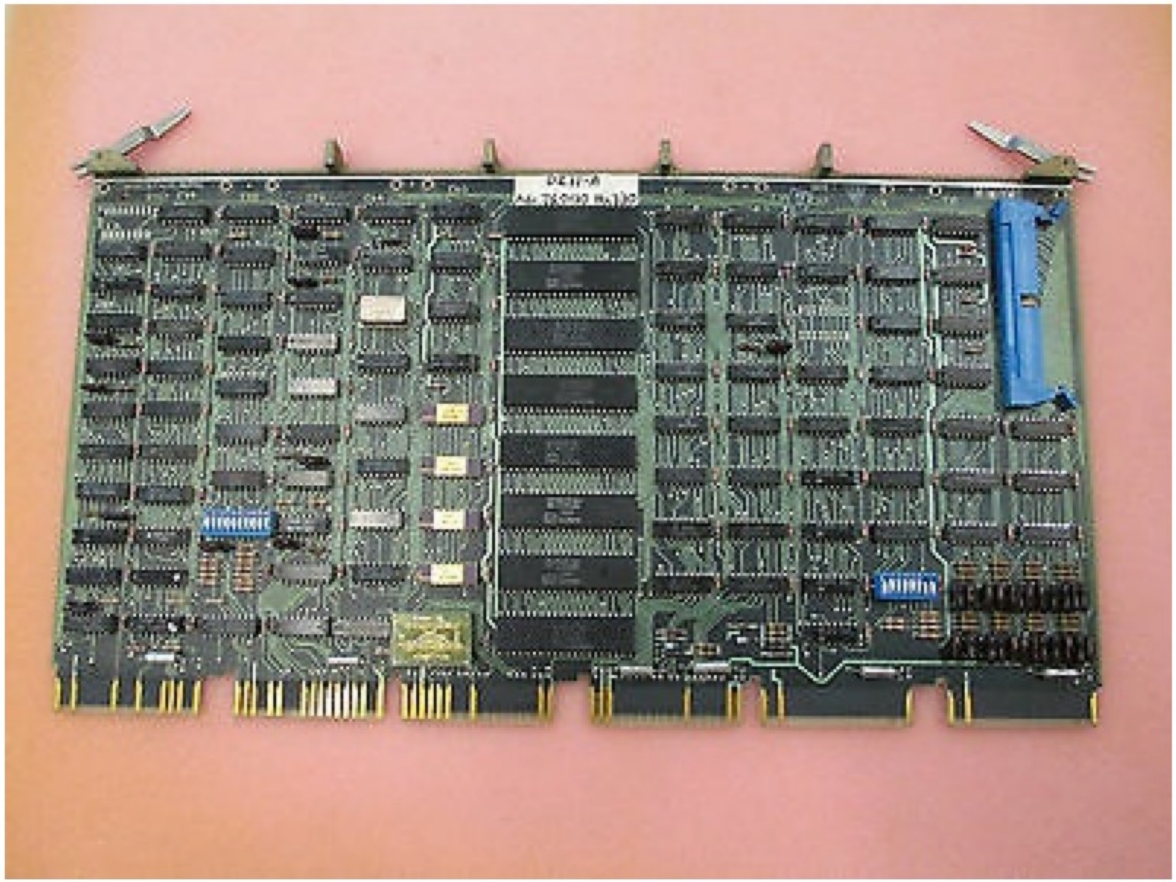

The integrated circuits were very simple with only a few dozen transistors in, so they didn't need very precise machines to make them. The boards could still be assembled by hand.

Old Computer Tech

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

LSI IC's

This is the second computer I assembled, the Acorn Atom

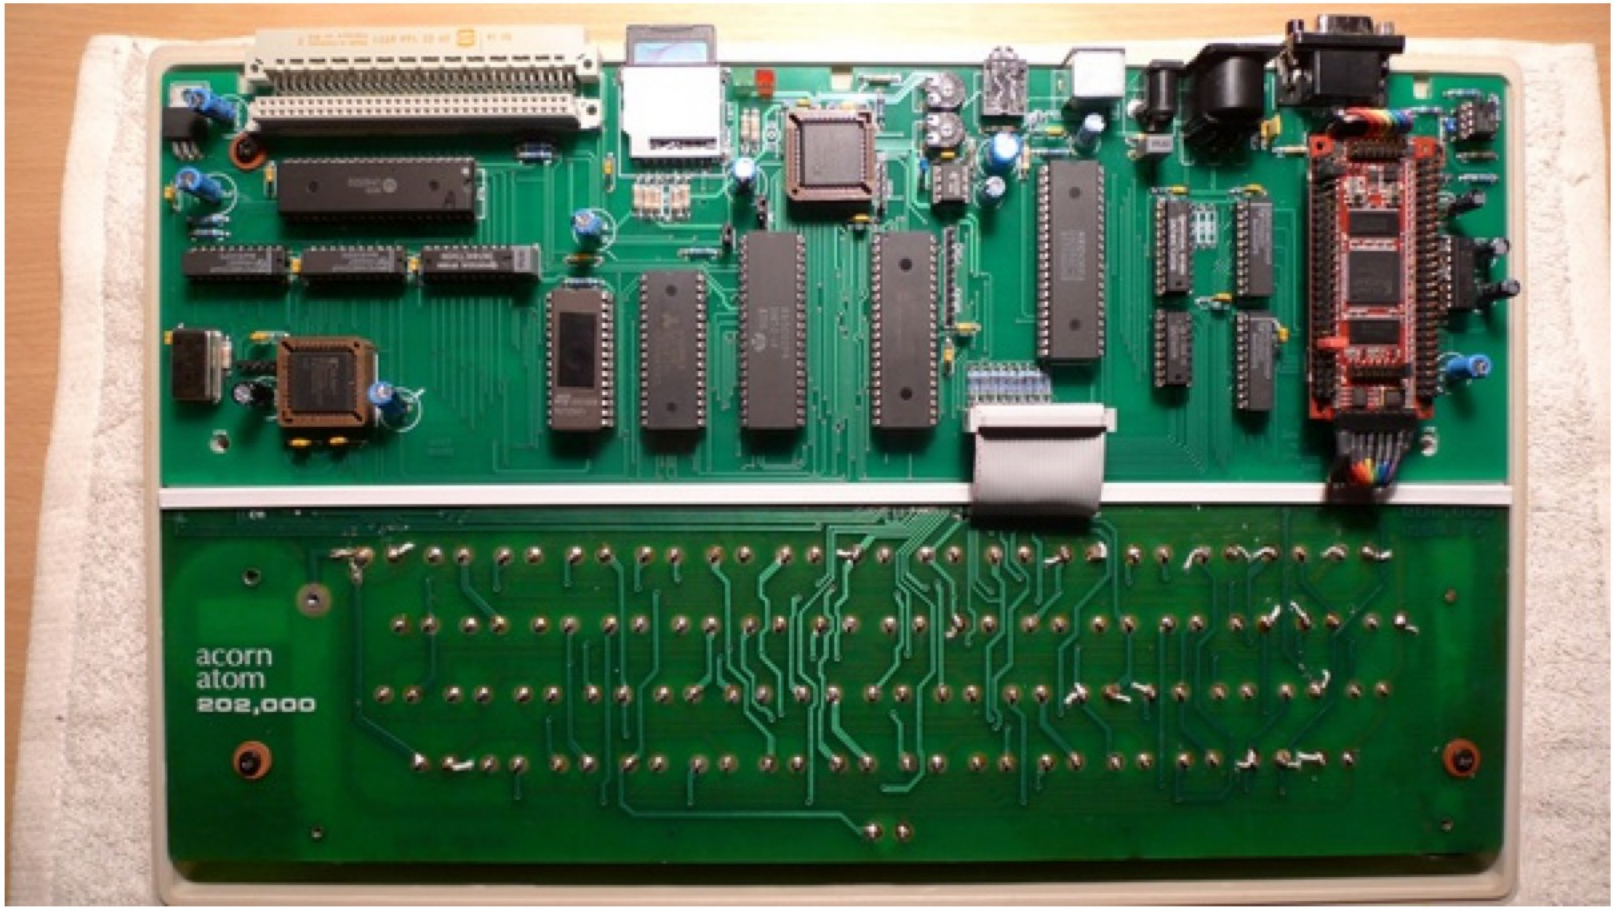

New Computer Tech

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

VLSI IC's

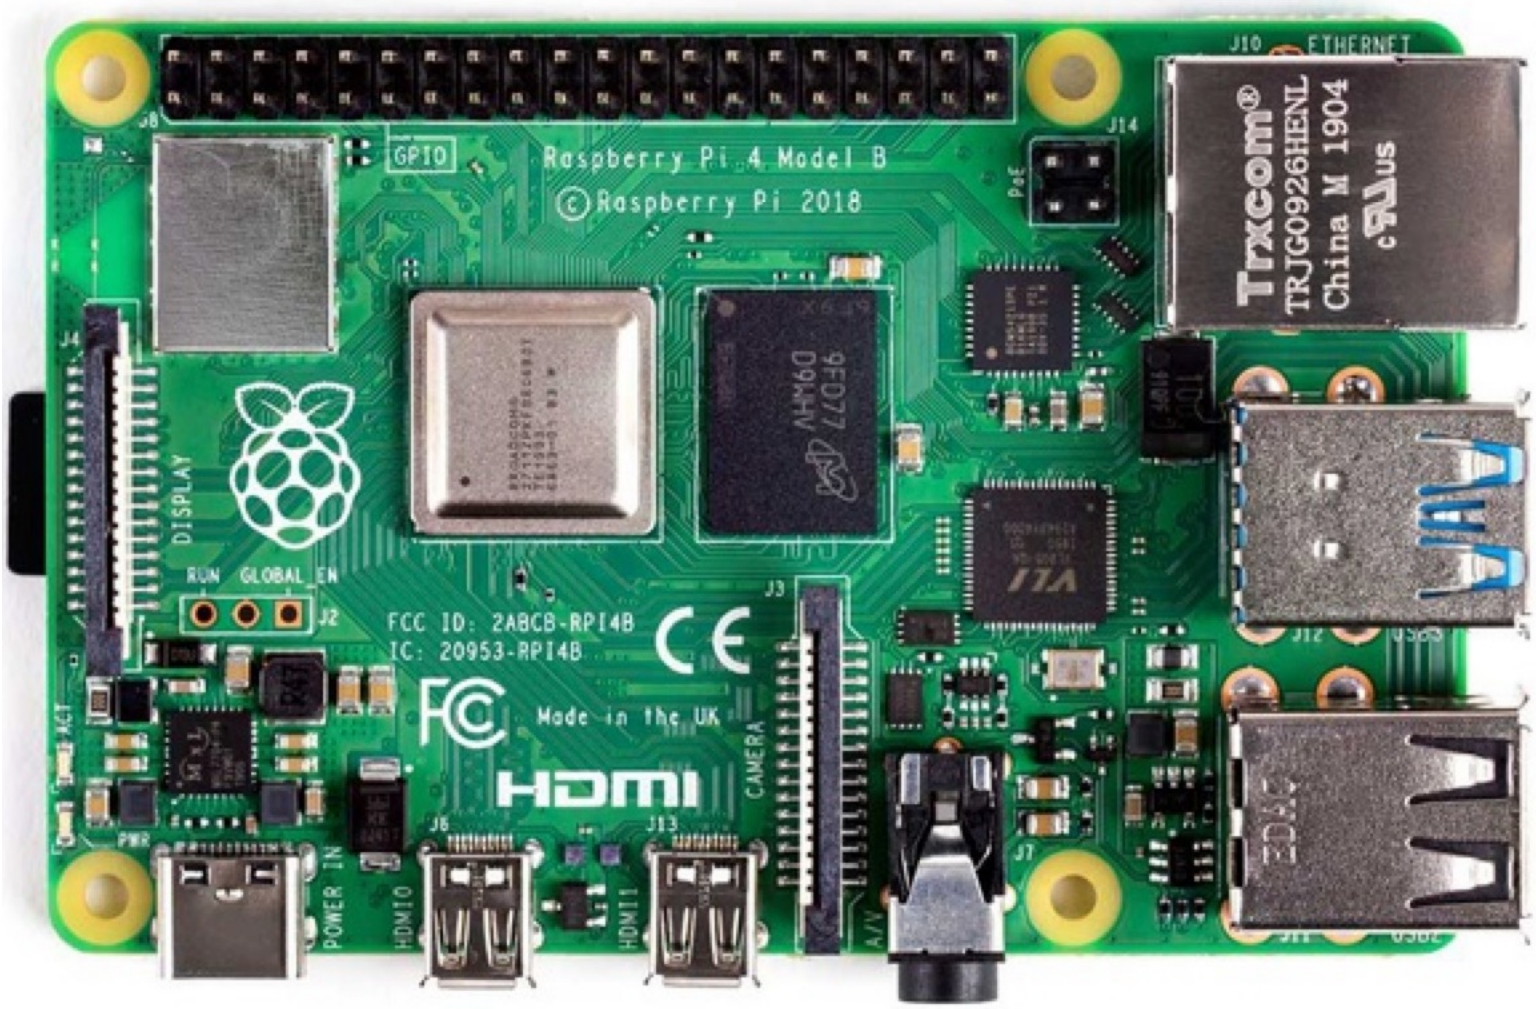

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Minis

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jeff Drobman

DEC PDP-11

DEC PDP/VAX

1st **LSI**-chip Computer | 1970 | Wiki

PDP-11/40. The processor is at the bottom. A TU56 dual DECtape drive is installed above it.

Mag tape

DEC PDP/VAX

DEC PDP-8e

1st **MSI**-chip Computer | 1965

DEC PDP/VAX

Here is a picture of one of the more popular versions, the PDP8/e. (By Florian Schäffer - Own work, CC BY-SA 4.0, File:Digital pdp8-e2.jpg)

DEC PDP/VAX

DEC PDP-11

1970 Wiki

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

PDP-11

From Wikipedia, the free encyclopedia
(Redirected from DEC PDP-11)

This article is about the PDP-11 series of minicomputers. For the PDP-11 processor architecture, see PDP-11 archit

The **PDP-11** is a series of 16-bit minicomputers sold by Digital Equipment Corporation (DEC) from 1970 into the 1990s, one of a succession of products in the PDP series. In total, around 600,000 PDP-11s of all models were sold, making it one of DEC's most successful product lines. The PDP-11 is considered by some experts[1][2][3] to be the most popular minicomputer ever.

The PDP-11 included a number of innovative features in its instruction set and additional general-purpose registers that made it much easier to program than earlier models in the PDP series. Additionally, the innovative Unibus system allowed external devices to be easily interfaced to the system using direct memory access, opening the system to a wide variety of peripherals. The PDP-11 replaced the PDP-8 in many real-time applications, although both product lines lived in parallel for more than 10 years. The ease of programming of the PDP-11 made it very popular for general-purpose computing uses as well.

The design of the PDP-11 inspired the design of late-1970s microprocessors including the Intel x86[] and the Motorola 68000. Design features of PDP-11 operating systems, as well as other operating systems from Digital Equipment, influenced the design of other operating systems such as CP/M and hence also MS-DOS. The first officially named version of Unix ran on the PDP-11/20 in 1970. It is commonly stated that the C programming language took advantage of several low-level PDP-11–dependent programming features,[4] albeit not originally by design.[5]

An effort to expand the PDP-11 from 16 to 32-bit addressing led to the VAX-11 design, which took part of its name from the PDP-11.

DEC PDP/VAX

DEC PDP-11

1970 Wiki

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

No dedicated I/O instructions [edit]

Early models of the PDP-11 had no dedicated bus for input/output, but only a system bus called the Unibus, as input and output devices were mapped to memory addresses.

An input/output device determined the memory addresses to which it would respond, and specified its own interrupt vector and interrupt priority. This flexible

Interrupts [edit]

The PDP-11 supports hardware interrupts at four priority levels. Interrupts are serviced by software service routines, which could specify whether they themselves could be interrupted (achieving interrupt nesting). The event that causes the interrupt is indicated by the device itself, as it informs the processor of the address of its own interrupt vector.

Interrupt vectors are blocks of two 16-bit words in low kernel address space (which normally corresponded to low physical memory) between 0 and 776. The first word of the interrupt vector contains the address of the interrupt service routine and the second word the value to be loaded into the PSW (priority level) on entry to the service routine.

Instruction set orthogonality [edit]

See also: PDP-11 architecture

The PDP-11 processor architecture has a mostly orthogonal instruction set. For example, instead of instructions such as *load* and *store*, the PDP-11 has a *move* instruction for which either operand (source and destination) can be memory or register. There are no specific *input* or *output* instructions; the PDP-11 uses memory-mapped I/O and so the same *move* instruction is used; orthogonality even enables moving data directly from an input device to an output device. More complex instructions such as *add* likewise can have memory, register, input, or output as source or destination.

Most operands can apply any of eight addressing modes to eight registers. The addressing modes provide register, immediate, absolute, relative, deferred (indirect), and indexed addressing, and can specify autoincrementation and autodecrementation of a register by one (byte instructions) or two (word instructions). Use of relative addressing lets a machine-language program be position-independent.

DEC PDP/VAX

DEC PDP-11

1970 Wiki

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

LSI-11 [edit]

The LSI-11 (PDP-11/03), introduced in February 1975[10] is the first PDP-11 model produced using large-scale integration; the entire CPU is contained on four LSI chips made by Western Digital (the MCP-1600 chip set; a fifth chip can be added to

The CPU microcode includes a debugger: firmware with a direct serial interface (RS-232 or current loop) to a terminal. This lets the operator do debugging by typing commands and reading octal numbers, rather than operating switches and reading lights, the typical debugging method at the time. The operator can thus examine and modify the computer's registers, memory, and input/output devices, diagnosing and perhaps correcting failures in software and peripherals (unless a failure disables the microcode itself). The operator can also specify which disk to boot from.

DEC "Fonz-11" (F11) Chipset

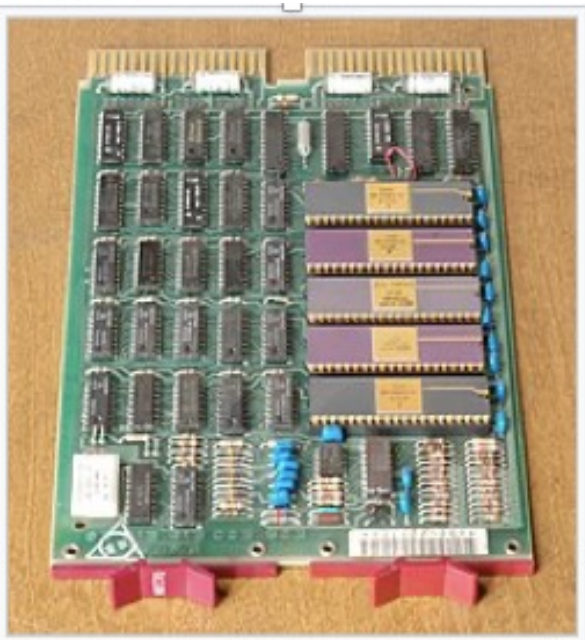

Q-Bus board with LSI-11/2 CPU

DEC System 10

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Decsystem 10. Joe Mabel, CC BY-SA 3.0, via Wikimedia Commons

K-202

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Mini 1972

A printer, camera, radar

Jacek Karpinski, creator of the K-202 minicomputer, at the Poznań International Fair, photo: Aleksander Jalosiński

K-202

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Mini 1972

Photo on an advertising leaflet handed out in 1972 at the Poznań International Fair

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

PC's

Early PC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1974

MITS Altair 8800

Altair 8800 Computer with 8-inch floppy disk

The **Altair 8800** is a microcomputer designed in 1974 by MITS and based on the Intel 8080 CPU at 2 MHz

Apple I Computer

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1976

Robert Mudry, Retired Silicon Valley engineering geek.
Shared Feb 6

A Hand-Built, Original Apple 1 Computer Is Yours for Just 1.5 Million Dollars

Early PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Commodore PET 1977

The Commodore PET 2001-8 alongside its rivals, the Apple II and the TRS-80 Model I

Early PC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Commodore PET

1977

A Commodore PET 2001

Manufacturer	Commodore International
Type	Personal computer
Release date	January 1977; 44 years ago[1]
Introductory price	US$795 (equivalent to $3,354 in 2019)[2]
Discontinued	1982; 39 years ago[3]
Operating system	Commodore BASIC 1.0 ~ 4.0
CPU	MOS Technology 6502 @ 1 MHz
Memory	4–96 KB
Storage	cassette tape, 5.25-inch floppy, 8-inch floppy, hard disk
Display	40×25 or 80×25 text
Graphics	monochrome character graphics
Sound	none or beeper
Successor	Commodore VIC-20

The chiclet keyboard of the PET 2001 series

Commodore 8028 Daisy wheel printer

IBM Datamaster

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

IBM's System 23 Datamaster, pictured here at the IBM Hursley Museum, cost $9,000 US

IBM PC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1981

- Designed and built by IBM (Boca Raton, FL) in 1981
- CPU: **i8088** (i8086 with an 8-bit data bus)
- Speed: 4.77 MHz
- Memory – DRAM: 64KB (48KB soldered, 16KB socketed)
- Disk: 5.25 inch floppy – 2 drives (A:, B:)
- OS: MS/PC-DOS (derived from CP/M)

The IBM PC

IBM PC Chips

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1981

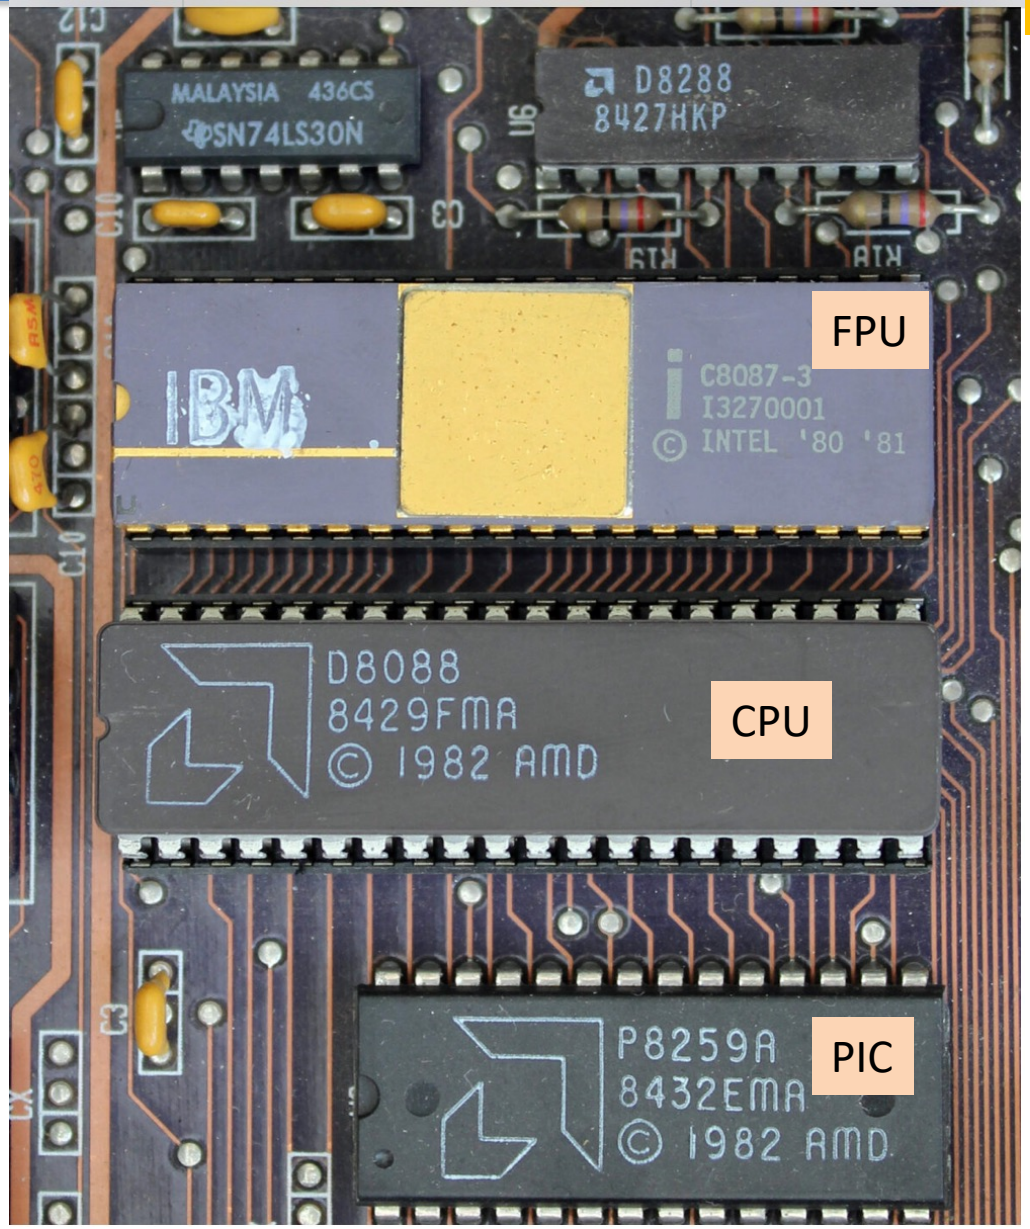

IBM PC – Large Motherbd

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

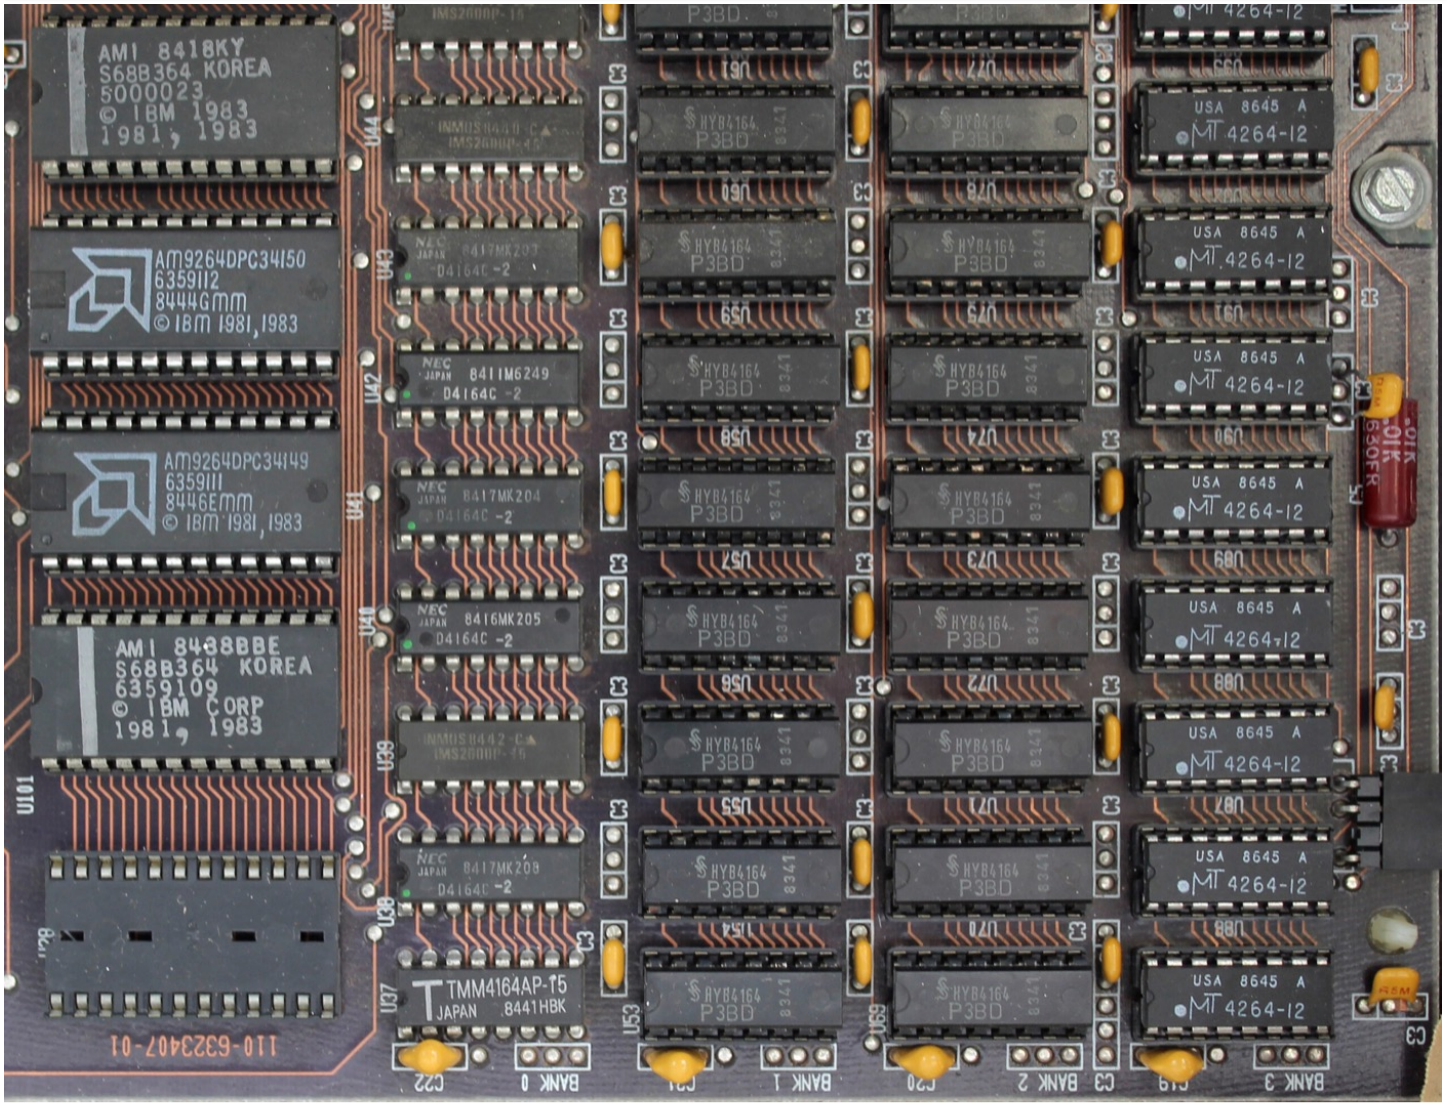

The memory is organised in four banks in the bottom right corner of the motherboard – in this case there are four 64KB banks, adding up to a total of 256KB

Osborne 1 Portable PC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1981

Osborne 1

From Wikipedia, the free encyclopedia

The **Osborne 1** is the first commercially successful portable microcomputer, released on April 3, 1981 by Osborne Computer Corporation.[1] It weighs 10.7 kg (24.5 lb), cost US$1,795, and runs the CP/M 2.2 operating system. It is powered from a wall socket, as it has no on-board battery, but it is still classed as a portable device since it can be hand-carried when packed.

The computer shipped with a large bundle of software that was almost equivalent in value to the machine itself, a practice adopted by other CP/M computer vendors. Competitors quickly appeared, such as the Kaypro II.

Compaq followed Osborne when founded in Feb 1982, and produced the first widely successful "portable" computer. I recall seeing at the Compaq booth at Comdex in Las Vegas in 1983. It ran Lotus 123 as a demo, both in Compaq's booth and in Lotus's booth. reason was Ben Rosen was VC who funded both companies. Ben Rosen met me there, and introduced me to both company founders, Mitch Kapor and Rod Canion.

Osborne 1 Portable PC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1981

Osborne 1

Developer	Adam Osborne
Manufacturer	Osborne Computer Corporation
Type	Portable computer
Release date	April 3, 1981; 39 years ago
Introductory price	US$1795 (today $5047.94)
Discontinued	1983
Operating system	CP/M
CPU	Zilog Z80 @ 4.0 MHz
Memory	64 KB RAM

Osborne 1 Portable PC

1981

Lee Felsenstein · Updated July 25, 2019
Electronic design engineer, PC pioneer

Designer

The laptop was invented by Adam Osborne in 1981. Why is Adam Osborne not as well known as Steve Jobs and/or Steve Wozniak?

I designed the Osborne-1, and was a founder of the company (part of my compensation for the design). It was intended not so much as a portable computer to be used "on the go" but as a self-contained CP/M computer which could hold its software disk library internally and make minimal demands on its owner for configuration and customization.

The weight was not 55 or 60 pounds - it was 23.5 pounds, and it created the class of computers that become known as "luggables". Someone once commented that most were only carried from one side of the office to the other. We never considered it a "laptop", though we were in discussions with multiple Japanese companies for years working toward a laptop.

I was aware of Alan Kay's "Dynabook" conceptual design at the time (1980) and attempted to advocate making that our design goal, but adequate flat-panel displays were not then available (just look at the small display on the Epson HX-20) and the cost of memory ICs did not then permit a graphic display.

I describe the Osborne-1 as "the first commercially successful portable computer" - others had been sold before with handles on them and IBM called their 1975 5100 a "Portable computer" (with cartridge tape storage and a price tag in the several thousands of dollars) but it did not provide an interoperable operating system).

Osborne Computer Corporation announced their first product in March 1981 and entered chapter 11 bankruptcy in September of 1983. If the company had been managed properly it might still be in business today. Innovators are not the ones in

Apple

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1984

The MacIntosh

Apple II

Wozniak

CPU: MT6502
(8-bit)

Steve Jobs

CPU: M68000
(16-bit)
RAM: 1MB

GUI: Xerox PARC (Alan Kay)

Microsoft's New PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Microsoft's Surface Laptop 3 is 3x more powerful than a MacBook Air

See this

Microsoft's New PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

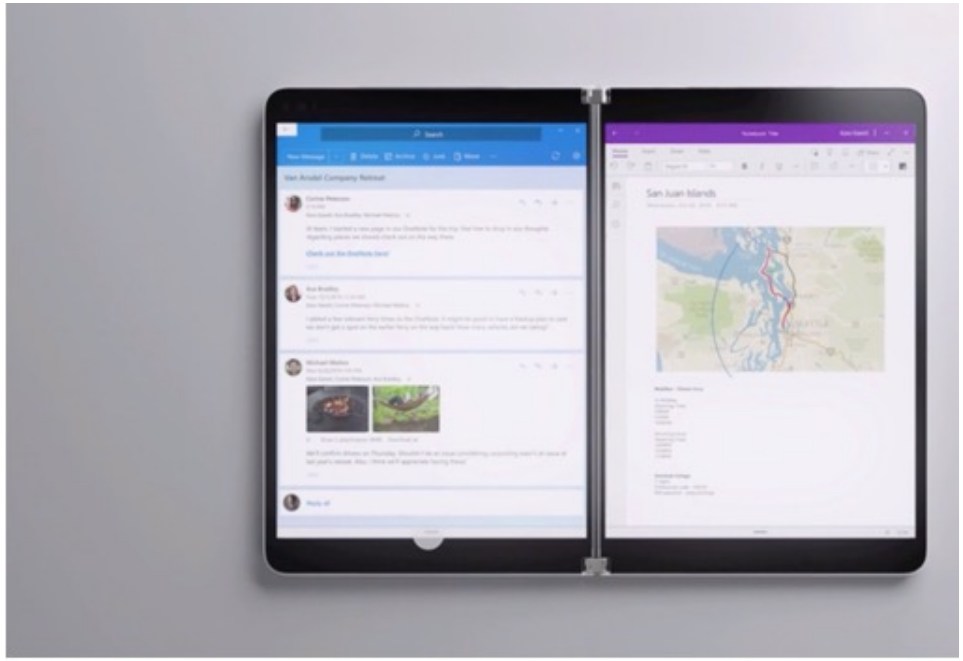

'This is Surface Duo': Yes, Microsoft will release a dual-screen phone

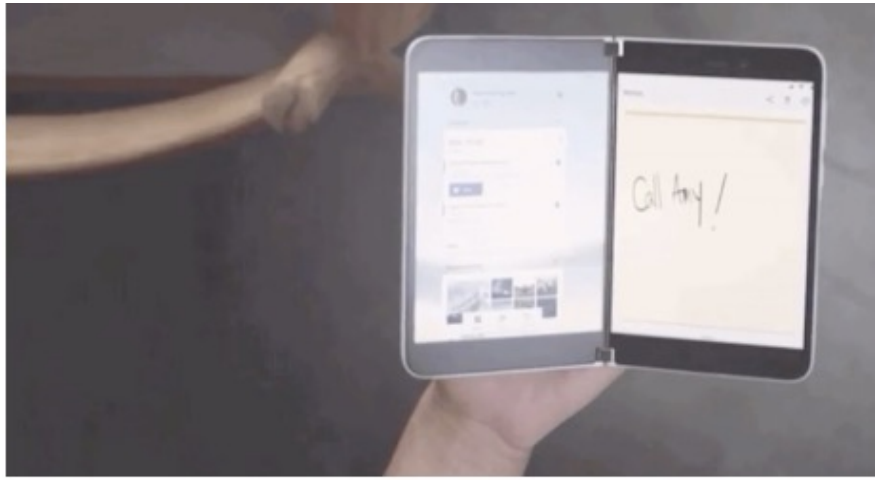

The new Microsoft phone, the Surface Duo.

Microsoft

Microsoft is making a phone. Again.

Microsoft's New PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Microsoft is making a phone. Again.

The tech giant, which stopped producing phones years ago, is hoping to stage a comeback with the Surface Duo. As the name suggests, the device has two screens, connected by a hinge. (Here's how the Surface Duo compares to the Galaxy Fold.)

"This product brings together the absolute best of Microsoft, and we're partnering with Google to bring the absolute best of Android in one product," said Microsoft Product Chief Panos Panay. "This is industry-pushing technology."

Now playing: Microsoft unveils Surface Duo, a foldable Android phone

▶ 4:58

The whole thing is a bit of a surprise, considering Microsoft eventually gave up making phones after its troubled 2014 purchase of phone giant Nokia for more than $7 billion. This time, Microsoft says its new innovations, like the dual-screen folding display and special technology it built for the Google Android software that powers the gadget, will help make the difference. We won't know for sure until the Surface Duo is released next year.

Microsoft's New PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Everything announced

- **Surface Duo**: Microsoft is making a dual-screen Android phone called Surface Duo. Yes, that Microsoft.
  The company known for its Windows operating system is getting back into smartphones by embracing its rival's ecosystem.
- **Surface Neo**: Dual-screen Microsoft Surface Neo is coming, eventually.
  Like other Windows 10 X systems, this dual-screen Surface isn't likely to be in stores till the 2020 holiday season.
- **Windows 10 X**: Windows 10 X OS will work with new dual-screen Surface Neo devices.
  The dual-screen Surface Neo gadgets will be out next year, says Microsoft.
- **Surface Laptop 3**: Microsoft announced a laptop 3x more powerful than the MacBook Air.
  Microsoft's Surface Laptop 3 has USB-C, a bigger screen and a modular design. It comes in 13.5- and 15-inch models and starts at $999.
- **Surface Pro X and Surface Pro 7**: Microsoft unveils $999 Surface Pro X, a tablet with a phonelike Windows experience.
  We also got a minimal update for the Surface Pro 7, now with USB-C.
- **SQ1 custom Arm chip**: Microsoft tries Windows on Arm chips again with the SQ1-powered Surface Pro X.
  The chips are designed to consume less power than those from Intel, Microsoft's traditional partner.

Microsoft's New PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

The Surface Laptop 3 has a custom Ryzen Surface Edition processor on the 15-inch model, while the Surface Pro X goes the ARM-powered route with a new SQ1 processor co-engineered with Qualcomm. It's a big change for the Surface line, even if Intel will still power the Surface Pro 7 and the smaller 13-inch Surface Laptop 3 models.

ARM ⟵ New Windows!

AMD Ryzen

On the AMD side, this Ryzen processor will be available exclusively in the 15-inch model of the Surface Laptop 3, a notebook that also has a metal finish instead of the fabric we've seen on previous Surface Laptop models. Microsoft has worked closely with AMD to add an additional graphics core on the 12nm Ryzen 5 and Ryzen 7 Surface parts that are built on Zen+, and to optimize the chip to fit inside the slim-and-light chassis it uses for the Surface Laptop 3.

Inside Microsoft's new custom Surface processors with AMD and Qualcomm

Surface Ryzen Edition and SQ1 processors have been co-engineered

By Tom Warren on October 2, 2019 11:30 am

Microsoft's New PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

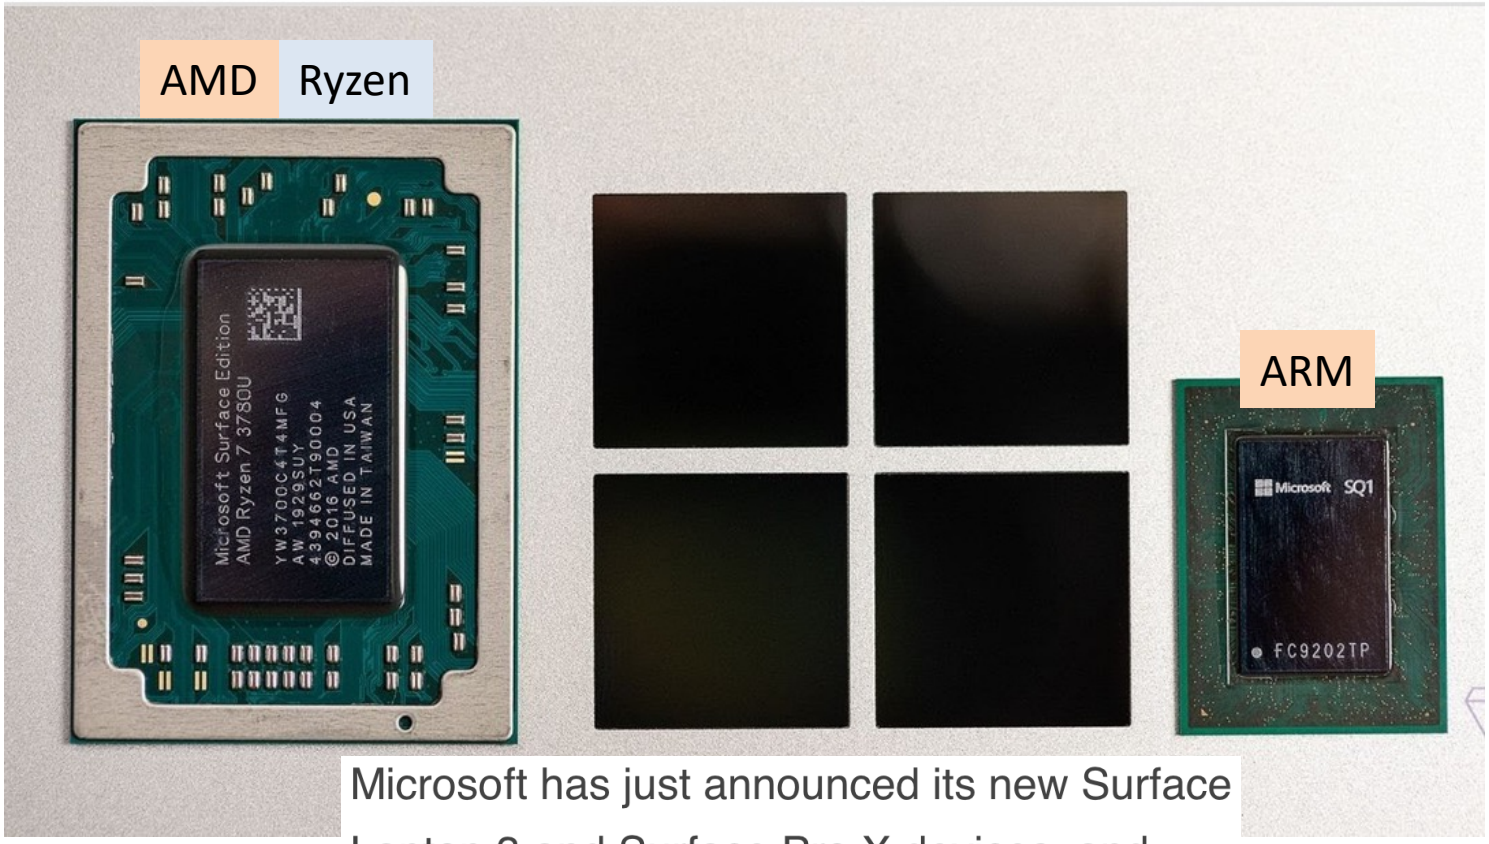

Microsoft has just announced its new Surface Laptop 3 and Surface Pro X devices, and neither will come with an Intel processor. The software giant is diversifying its silicon for Surface this year by partnering closely with AMD and Qualcomm, respectively, to create custom processors for its Surface line.

Prices for Top CPU's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

THE BEST CPUS AT THE BEST PRICE

AMD Ryzen 7 2700X Processor with Wraith Prism LED Cooler – YD270XBGAFBOX

amazon.com

~~$329.00~~
$236.46

VIEW

Intel Core i5-8500 Desktop Processor 6 Core up to 4.1GHz Turbo LGA1151 300 Series 65W

amazon.com

$239.97

VIEW

Intel Core i9-9900K Desktop Processor 8 Cores up to 5.0 GHz Turbo unlocked LGA1151 300 Series 95W

amazon.com

$494.99

VIEW

Intel Systems: NUC 11

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Intel Core i5/7 CPU's + GPU's

The Intel Phantom Canyon NUC 11 Extreme features a Tiger Lake-U 28W Core i7 or Core i5 processor. Both processors will be equipped with the latest generation of integrated Intel GPU Xe graphics processors.

Intel Systems: NUC 11

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

CPU

NUC 11 Extreme has a CPU with 4 cores and 8 threads, which has a clock base of 2.3 GHz and a maximum boost of 4.4 GHz. When tested with 3DMark, this CPU scores about 4,590 points, which is slightly lower than the Core i7-1165G7 CPU.

GPU

The Intel Canyon NUC 11 Extreme will use a discrete GPU in addition to the Intel Xe GPU. The American company plans to embed an NVIDIA GTX 1660 TI 80W GPU with 1,536 CUDA cores and a GPU clock frequency of around 1600 MHz.

RAM

Users will also receive RAM support up to DDR4-3200 with a capacity of up to 64 GB. Users will also receive 2 M.2 slots (1x 22×80 / 110 & 1x 22×80) and a PCIe x4 Gen 3 NVMe connector.

I/O

Front and rear extensions are equally plentiful, including two Thunderbolt 3, HDMI 2.0, MiniDP 1.4, and many USB extensions that also support an 8K output. Also available on the rear is a 2.5 Gbps Gigabit LAN (Intel solution) with IEEE 802.11ax (Intel Wireless-AX 201), WiFi 6 Wireless, Bluetooth 5.0 and more.

Industrial PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Introducing...
LITESHIELD 2.0
our newest industrial product line.

NO FANS. NO PROBLEM.

Windows

> VIEW THIS PRODUCT

Why fanless? >

Up to intel i9 capability.
8 cores/16 threads
wi-fi.
4k on dual displays.
Rigid aluminum enclosure.
litional Internal 2.5" SSD slot.
100% fanless. 100% silent.

Industrial PC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Nuvo Series

- Intel® 9th/8th-Gen CPU
- Core™ i7/i5/i3
- Fanless Box PC
- Multi-PoE & GbE

Rugged Computer

Nuvo VTC Series

- Intel® 9th/8th-Gen CPU
- Isolated CAN bus
- Isolated DIO
- E-Mark & EN50155

In-vehicle PC

Nuvis Series

- Core™ i7/i5
- Deterministic I/O
- Real-time Control
- NVIDIA® 75W GPU

Machine Vision

Nuvo GC Series

- Intel® 9th/8th-Gen CPU
- 250W GPU or Tesla T4
- M.2 2280 NVMe socket
- Wide-temp Operation

GPU Computing

Intel Core i3/5/7 CPU's

+Nvidia GPU's

Computers

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

HPC

*see separate slide set

Hardware

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quantum Computers

*see separate slide set

Quantum Computers (QC)

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

- Google
- IBM
- Intel
- Microsoft

Outlook

Businesses are hoping the advancement of quantum computers—by tech giants such as Google, IBM, and Intel, as well as startups such as Rigetti Computing—will lead to unprecedented scientific and technical breakthroughs in the coming years. They're eyeing applications from new chemical reactions for the development of drugs, fertilizers, and batteries, to the improvement of optimization algorithms and mathematical modeling.

Computers & QC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

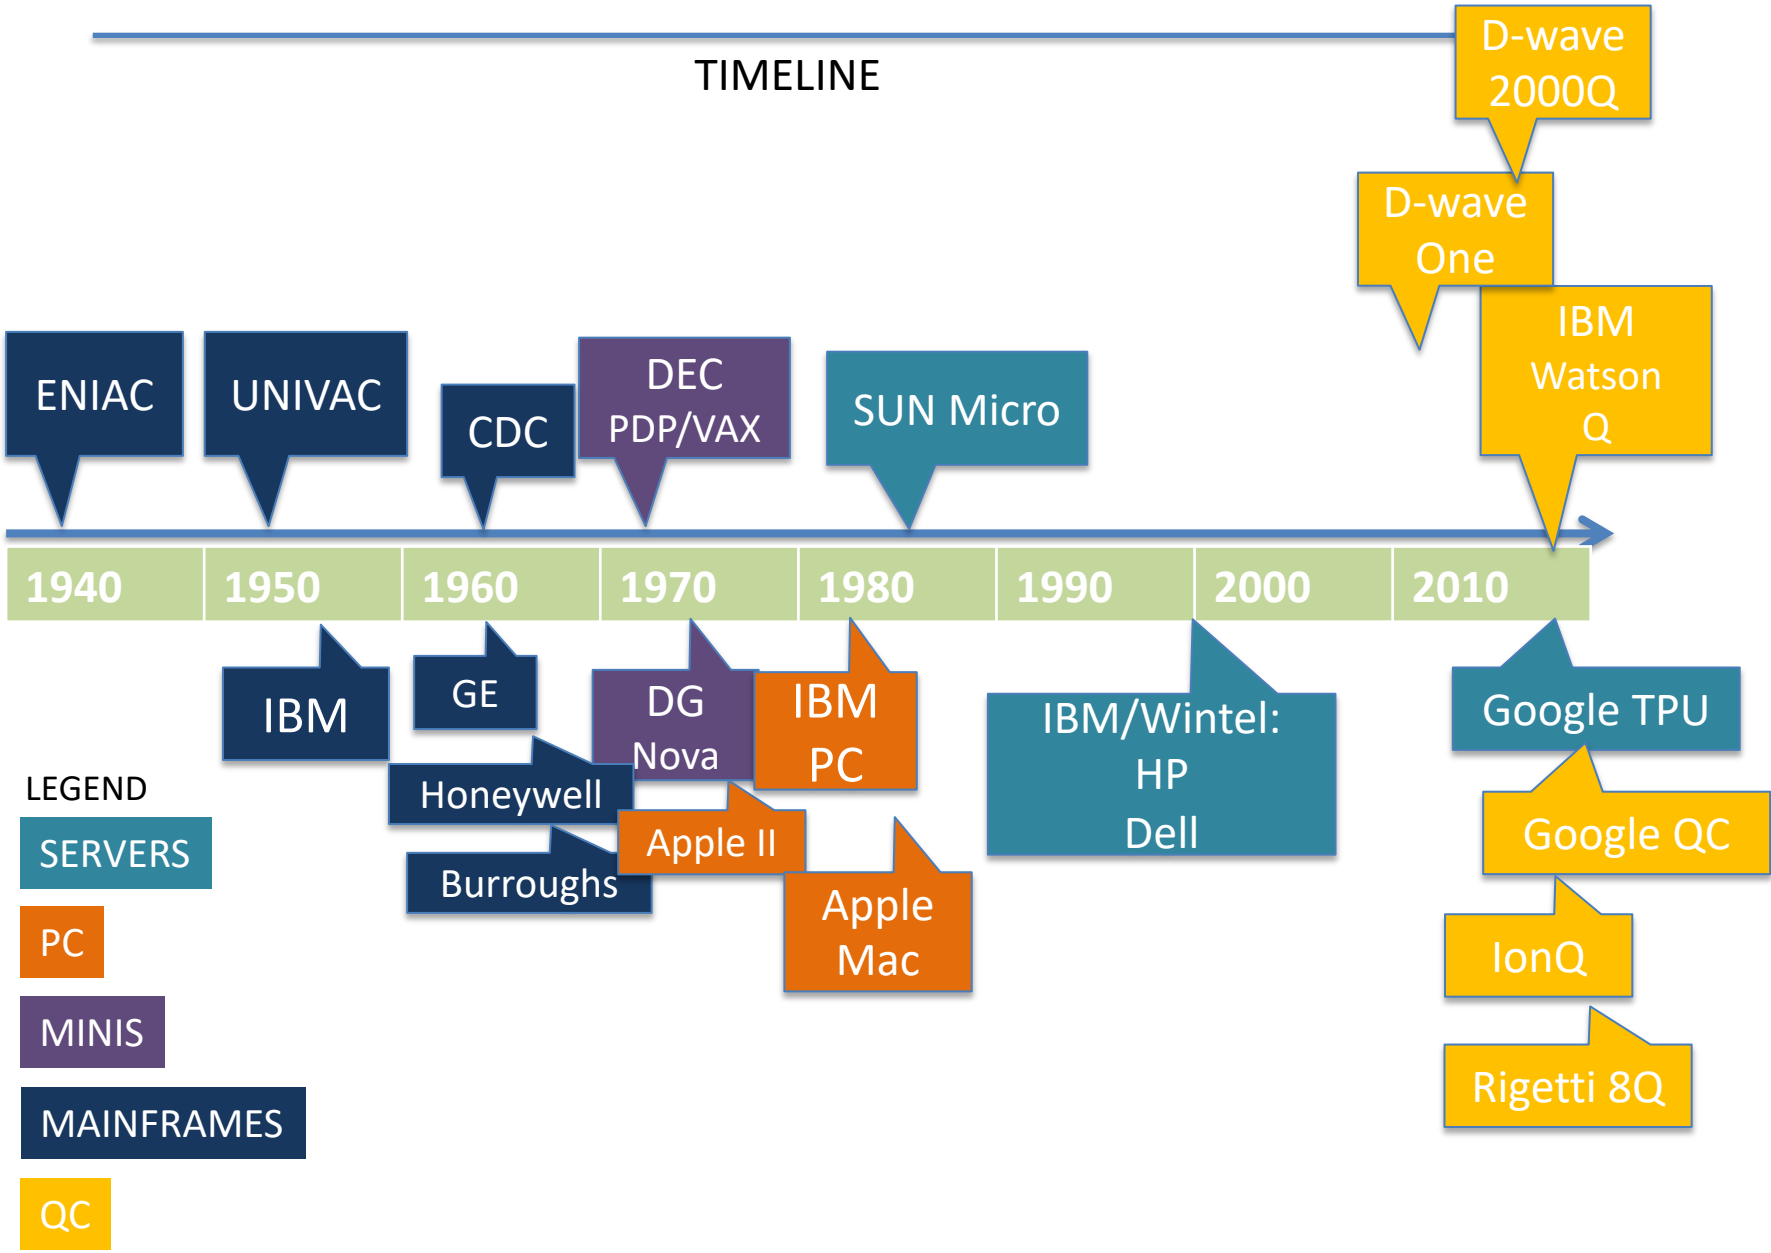

Commercial QC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

List of quantum processors

From Wikipedia, the free encyclopedia

Circuit-based quantum processors [edit]

These QPUs are based on the quantum circuit and quantum logic gate-based model of computing.

Manufacturer	Name/Codename/Designation	Architecture	Layout	Socket	Fidelity	Qubits	Release date
Google	N/A	Superconducting	N/A	N/A	99.5%[1]	20 qb	2017
Google	N/A	Superconducting	7×7 lattice	N/A	99.7%[1]	49 qb[2]	Q4 2017 (planned)
Google	Bristlecone	Superconducting	6×12 lattice	N/A	99% (readout) 99.9% (1 qubit) 99.4% (2 qubits)	72 qb[3][4]	5 March 2018
Google	Sycamore	Nonlinear superconducting resonator	N/A	N/A	N/A	54 transmon qb 53 qb effective	2019

Commercial QC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

IBM	IBM Q 5 Tenerife	Superconducting	bow tie	N/A	99.897% (average gate) 98.64% (readout)	5 qb	2016[1]
IBM	IBM Q 5 Yorktown	Superconducting	bow tie	N/A	99.545% (average gate) 94.2% (readout)	5 qb	
IBM	IBM Q 14 Melbourne	Superconducting	N/A	N/A	99.735% (average gate) 97.13% (readout)	14 qb	
IBM	IBM Q 16 Rüschlikon	Superconducting	2×8 lattice	N/A	99.779% (average gate) 94.24% (readout)	16 qb[5]	17 May 2017 (Retired: 26 September 2018)[6]
IBM	IBM Q 17	Superconducting	N/A	N/A	N/A	17 qb[5]	17 May 2017
IBM	IBM Q 20 Tokyo	Superconducting	5x4 lattice	N/A	99.812% (average gate) 93.21% (readout)	20 qb[7]	10 November 2017
IBM	IBM Q 20 Austin	Superconducting	5x4 lattice	N/A	N/A	20 qb	(Retired: 4 July 2018)[6]
IBM	IBM Q 50 prototype	Superconducting	N/A	N/A	N/A	50 qb[7]	
IBM	IBM Q 53	Superconducting	N/A	N/A	N/A	53 qb	October 2019

Commercial QC's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Rigetti	8Q Agave	Superconducting	N/A	N/A	N/A	8 qb	4 June 2018[11]
Rigetti	16Q Aspen-1	Superconducting	N/A	N/A	N/A	16 qb	30 November 2018[11]
Rigetti	19Q Acorn	Superconducting	N/A	N/A	N/A	19 qb[12]	17 December 2017
IBM	IBM Ourense[13]	Superconducting	T	N/A	N/A	5 qb	03 July 2019
IBM	IBM Vigo[13]	Superconducting	T	N/A	N/A	5 qb	03 July 2019
IBM	IBM London[13]	Superconducting	T	N/A	N/A	5 qb	13 September 2019
IBM	IBM Burlington[13]	Superconducting	T	N/A	N/A	5 qb	13 September 2019
IBM	IBM Essex[13]	Superconducting	T	N/A	N/A	5 qb	13 September 2019

Annealing quantum processors [edit]

These QPUs are based on quantum annealing.

Manufacturer	Name/Codename/Designation	Architecture	Layout	Socket	Fidelity	Qubits	Release date
D-Wave	D-Wave One (Ranier)	Superconducting	N/A	N/A	N/A	128 qb	11 May 2011
D-Wave	D-Wave Two	Superconducting	N/A	N/A	N/A	512 qb	2013
D-Wave	D-Wave 2X	Superconducting	N/A	N/A	N/A	1152 qb	2015
D-Wave	D-Wave 2000Q	Superconducting	N/A	N/A	N/A	2048 qb	2017
D-Wave	D-Wave Advantage	Superconducting	N/A	N/A	N/A	5000 qb	2020

Computer Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Modern Microprocessors

AMD

Modern CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

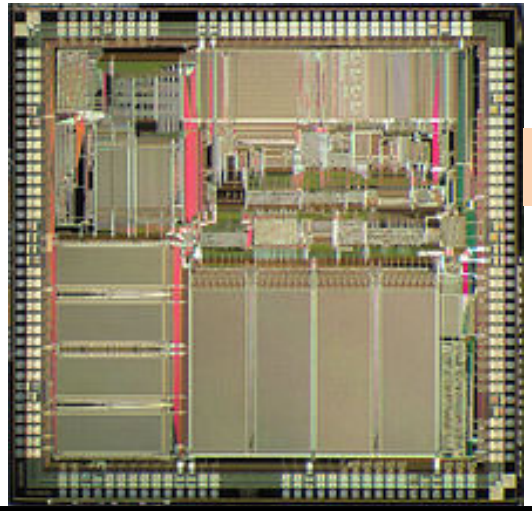

ARM 610

AMD “Mullins” APU

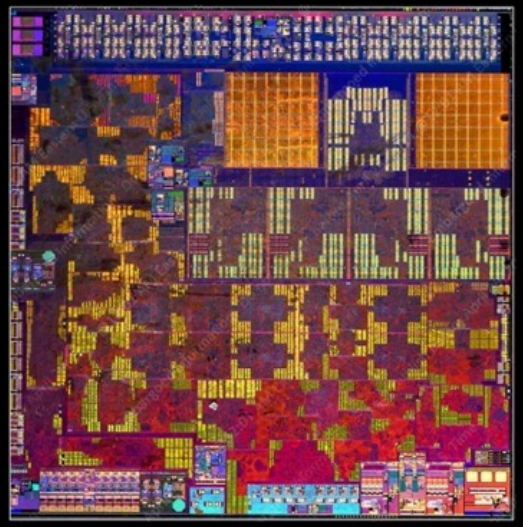

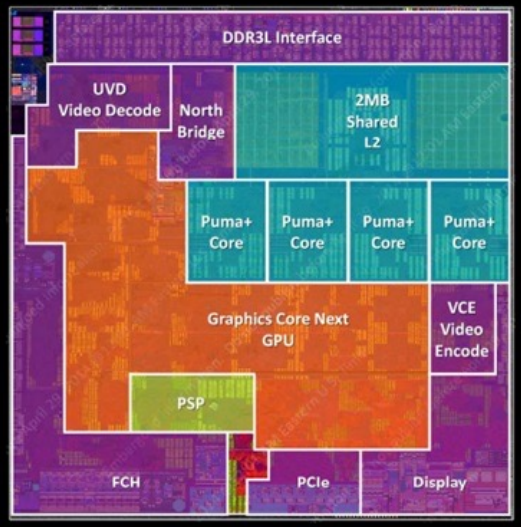

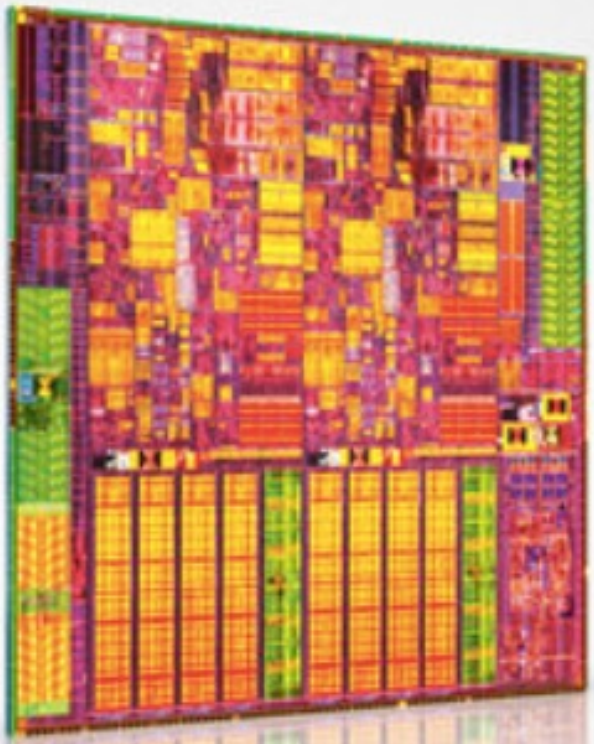

Intel “Westmere” CPU

The Westmere Die, a processor introduced by Intel in 2010. Intel

AMD vs. Intel

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

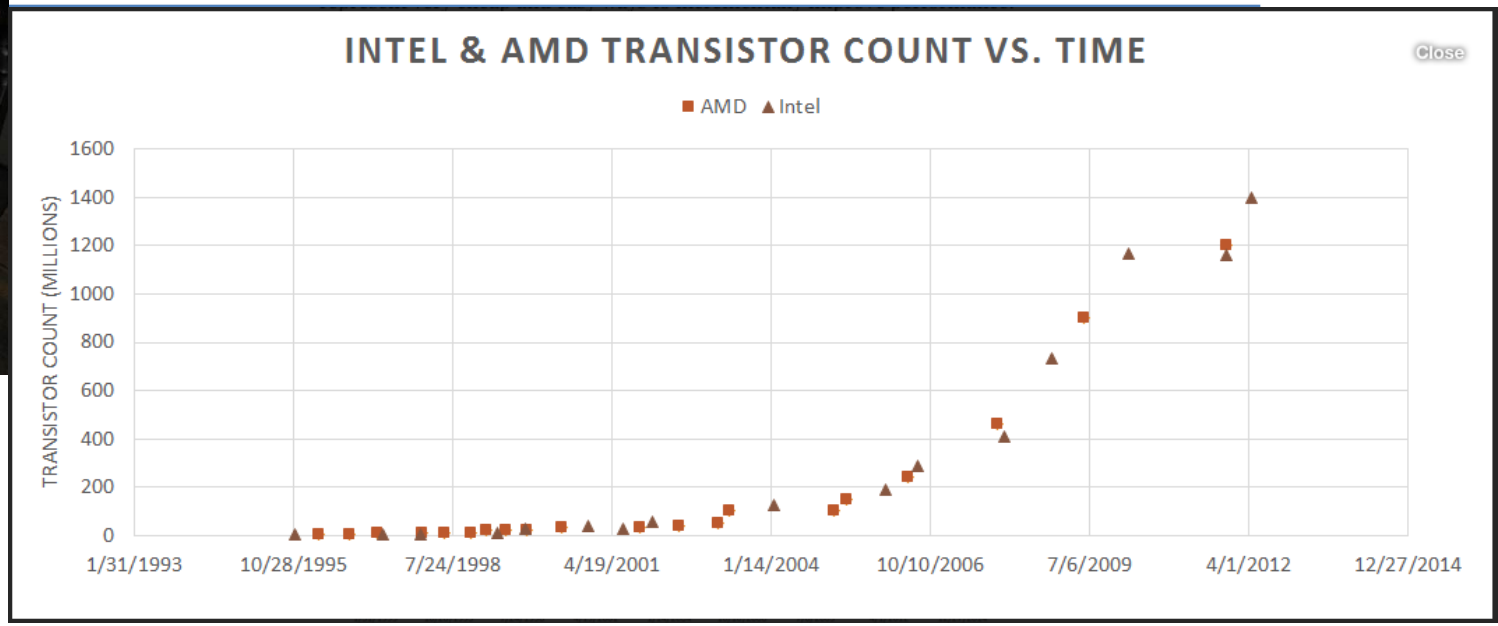

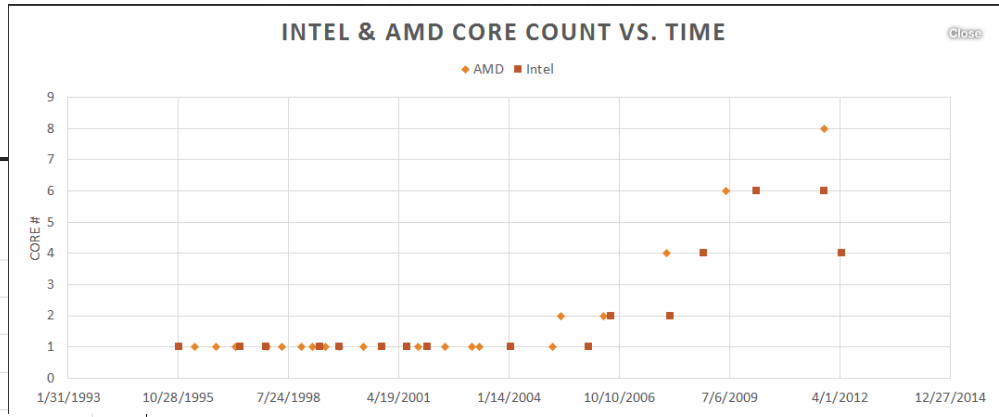

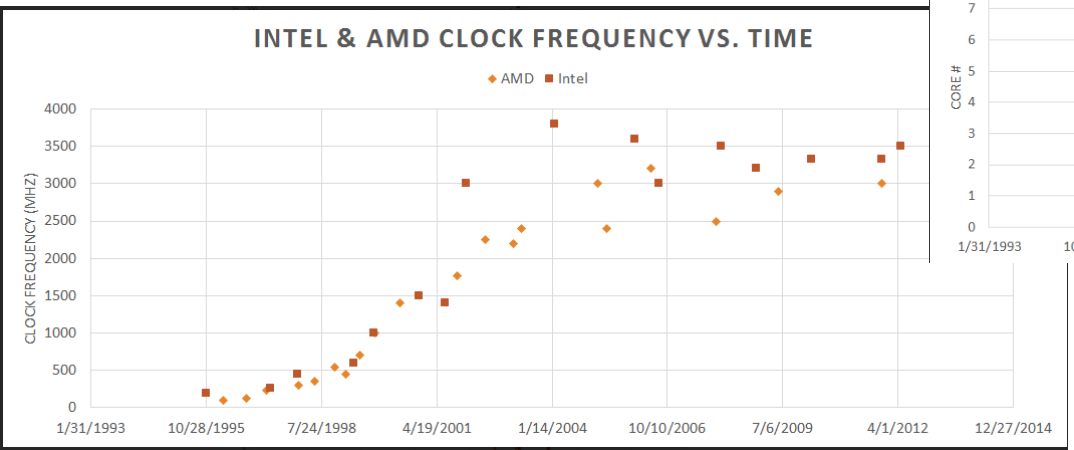

AMD CPU's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD x86

V·T·E	AMD processors	
Lists	Microprocessors · Microarchitectures · Chipsets · Sockets · Duron · Athlon (XP) · Athlon 64 (X2) · Sempron · Phenom · Ryzen ·	
Microarchitectures	IA-32 (32-bit)	K5 · K6 · Athlon/K7
	x86-64 desktop	K8 · K9 · K10 (aka 10h) · 15h (Bulldozer · Piledriver · Steamroller · Excavator) · Zen (Zen+ · Zen 2)
	x86-64 low-power	Bobcat (aka 14h) · 16h (Jaguar · Puma)
	ARM64	K12 (aka 12h)
Current products	IA-32 (32-bit)	Geode
	x86-64 (64-bit)	APU · Athlon X4 · FX · Ryzen · Epyc · Opteron
Discontinued	Early x86 (16-bit)	Am286
	IA-32 (32-bit)	Am386 · Am486 · Am5x86 · K5 · K6 · K6-2 · K6-III · Duron · Athlon (XP · MP)
	x86-64 (64-bit)	Sempron · Athlon 64 (X2 · II) · Phenom (II) · Turion
	Other	Am9080 · **Am2900** (list) · Am29000 · Alchemy (MIPS32)

AMD's Zen

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD News
AMD's new Zen Processor

AMD Zen
stock news usa (2016)
(Click image to view full size)

AMD made radical alterations to its Zen design while keeping itself distant from an ugly past. The company knew it had to make the changes to become a force to reckon with in the server and PC markets. So when the designers of the chip sat down to map the Zen design, they had two priorities: To boost CPU performance to maximum and to stabilize power efficiency.

According to a company spokesperson, the chips will come with 8 to 32 cores. The 32-core chips may come in the quad-CPU configurations although those details haven't been finalized yet.

5-19B Transistors

- CPU performance
- Power efficiency

8-32 cores

Source: stocknewsusa (2016-08-26)

Inside AMD's Production Of The Zen CPU

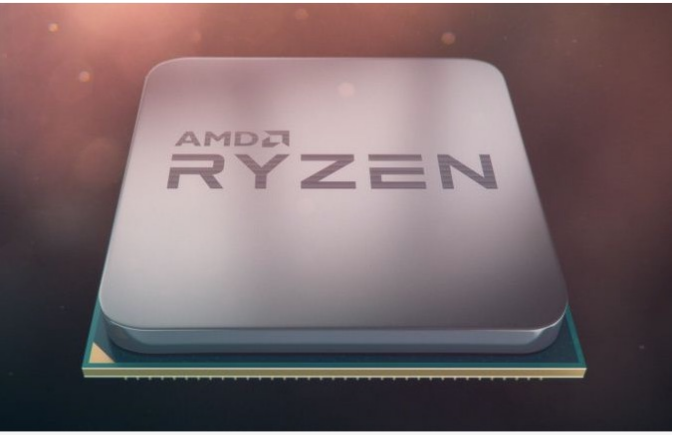

Ryzen 7 will have three CPUs to start, all having eight cores and supporting simultaneous multi-threading:

- *Ryzen 7 1800X: 8C/16T, 3.6 GHz base, 4.0 GHz turbo, 95W, $499*
- *Ryzen 7 1700X: 8C/16T, 3.4 GHz base, 3.8 GHz turbo, 95W, $399*
- *Ryzen 7 1700: 8C/16T, 3.0 GHz base, 3.7 GHz turbo, $329*

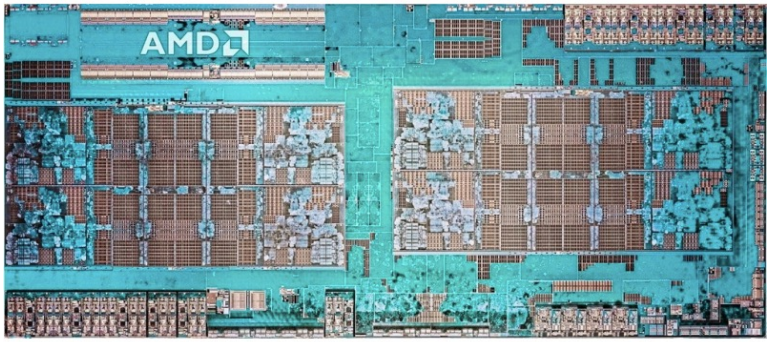

Ryzen

4.8 billion transistors and more than 2,000m of signal wire

AMD RyZen

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

128GiB DRAM

The big old Threadripper in my desktop PC only supports 128GiB DRAM. But processors built for servers, Intel® Xeon® and AMD EPYC, will support 1TB+ DRAM. The EPYC 7742 will support 4TiB DRAM, and the price of the DRAM won't seem so shocking when you learn the price of your 64-core CPU is about $7,500!

AMD CPU's in PS5

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jeff Drobman

7nm Ryzen

PS5 Specs: Rumors and Current Predictions

Before we dive into our in-depth look at potential PS5 specs, lets start with the current predictions. The system's architect, Mark Cerny, revealed some official details in April 2019, but the exact specs still remain a mystery. Here is what we know from that reveal:

- PS5 will support up to 8K resolutions
- The system will include an SSD that will drastically reduce load times
- Ray tracing (a powerful graphics technique) is supported by PS5
- The system uses a variation of AMD's third generation Ryzen with eight cores of the new 7nm Zen 2 microarchitecture
- GPU is a variation of the Radeon Navi family
- The system provides 3D audio without any additional hardware

AMD CPU's in PS5

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

7nm Ryzen

These are our current spec predictions based on rumors and what was revealed by Mark Cerny:

- CPU: 8 core/16 threads at 3.2Ghz with a Zen2 architecture
- GPU: Navi-based with AMD next-gen features at 12.6 to 14.2 teraflops
- Memory: 24GB total with reportedly 20GB GDDR6 at 880GB/S and 4GB DDR4 for the operating system
- 2 TB SSD

AMD Over-clocked

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD Ryzen 9 vs. Intel i9

AMD Update

AMD News

Overclocked Ryzen Sets New Performance Record.

A liquid nitrogen-cooled AMD Ryzen 9 3900X processor has just set a new overclocking world record in the wPrime benchmark, beating the previous title holder, the Intel Core i9-7920X.

The feat was performed by Australian overclocker jordan.hyde99, who got the AMD Ryzen 9 3900X to reach speeds of 5,625 MHz to complete the tests in just 35 seconds and 517 milliseconds.

Liquid nitrogen-cooled AMD Ryzen 9 3900X is the new world overclocking champ

Source: Tech radar (16 Dec 2019)

Computer Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Modern Microprocessors

Intel

Intel Pentium “M”

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

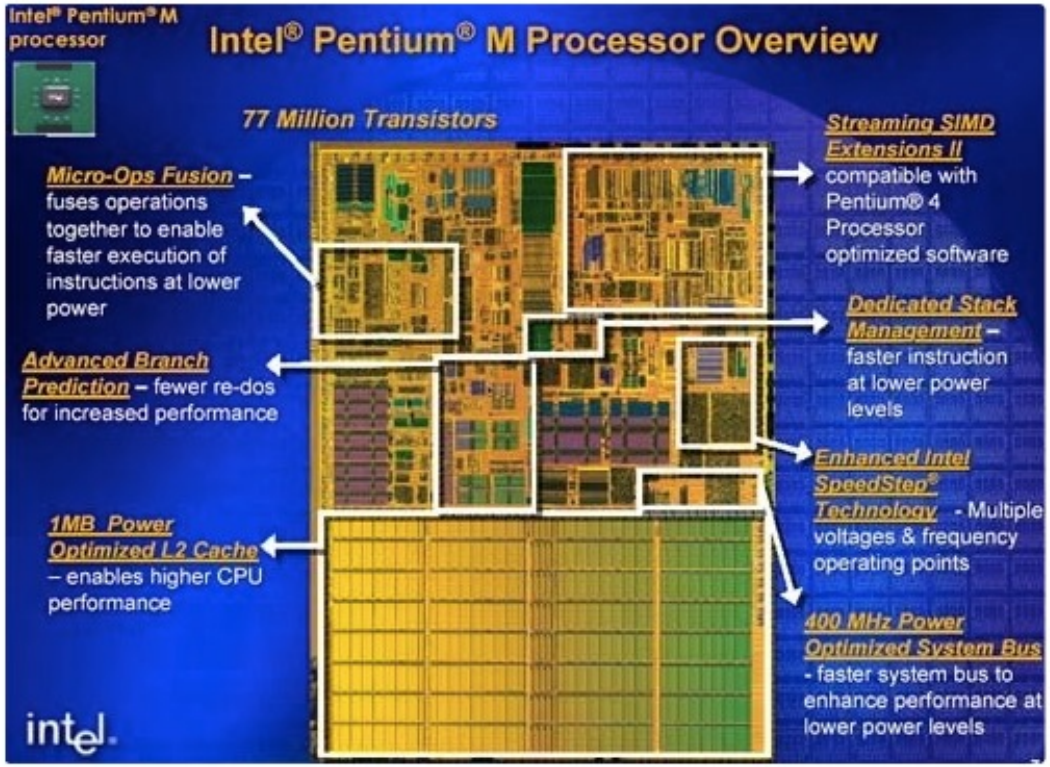

The arrival of the Pentium M on the market immediately gave Intel back the performance lead in the mobile segment. A second revision, the “Dothan,” increased this lead. The Pentium M was so superior to even desktop Pentium 4’s that many gamers bought them and ran them on desktop motherboards released for exactly that purpose. Of course, it wasn’t long before Intel started contemplating the release of a real desktop version of the Pentium M. When the next-generation Netburst design, called “Tejas,” proved to a disaster (engineering samples used 150W at 2.8 GHz), Intel cancelled it and decided to make the Pentium M design the basis of all future Intel releases.

Intel Core i3/5/7/9 + Xeon

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Intel® Core™

- Intel® Core™ X-Series
- 9th Gen Intel® Core™ i9
- 10th Gen Intel® Core™ i7
- 10th Gen Intel® Core™ i5
- 10th Gen Intel® Core™ i3
- 8th Gen Intel® Core™ m3
- 9th Gen Intel® Core™ vPro™

1 or 2 **Threads**/Core

1.75-3B Transistors

Intel® Xeon®

- Intel® Xeon® Scalable
- Intel® Xeon® D
- Intel® Xeon® W
- Intel® Xeon® E

Intel Atom®

- Intel Atom® C
- Intel Atom® E
- Intel Atom® X

Pentium®

Celeron®

Itanium®

Intel® Quark™

Bottom of an LGA1151 CPU.

Top of an Intel Core i7-6700K (6th Gen).

Intel Xeon

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

>60 Cores

Xeon Phi 7200 Series	sSpec Number	Cores (Threads)	Clock (MHz)		L2 Cache	MCDRAM Memory		DDR4 Memory		Peak DP Compute	TDP (W)	Soc-ket	Release Date	Part Num
			Base	Turbo		Quantity	BW	Quantity	BW					
Xeon Phi 7210[81]	SR2ME (B0)	64 (256)	1300	1500	32 MB	16 GB	400+ GB/s	384 GB	102.4 Gbit/s	2662 GFLOPS	215	SVLCLGA3647	20 June, 2016	HJ8066702859300
	SR2X4 (B0)													
Xeon Phi 7210F[82]	SR2X5 (B0)										230			HJ8066702975000
Xeon Phi 7230[83]	SR2MF (B0)										215			HJ8066702859400
	SR2X3 (B0)													
Xeon Phi 7230F[84]	SR2X2 (B0)										230			HJ8066702269002
Xeon Phi 7250[85]	SR2MD (B0)	68 (272)	1400	1600	34 MB					3046 GFLOPS[86]	215			HJ8066702859200
	SR2X1 (B0)													
Xeon Phi 7250F[87]	SR2X0 (B0)										230			HJ8066702268900
Xeon Phi 7290[88]	SR2WY (B0)	72 (288)	1500	1700	36 MB					3456 GFLOPS	245			HJ8066702974700
Xeon Phi 7290F[89]	SR2WZ (B0)										260			HJ8066702975200

Multi-Cores + L2/L3

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD

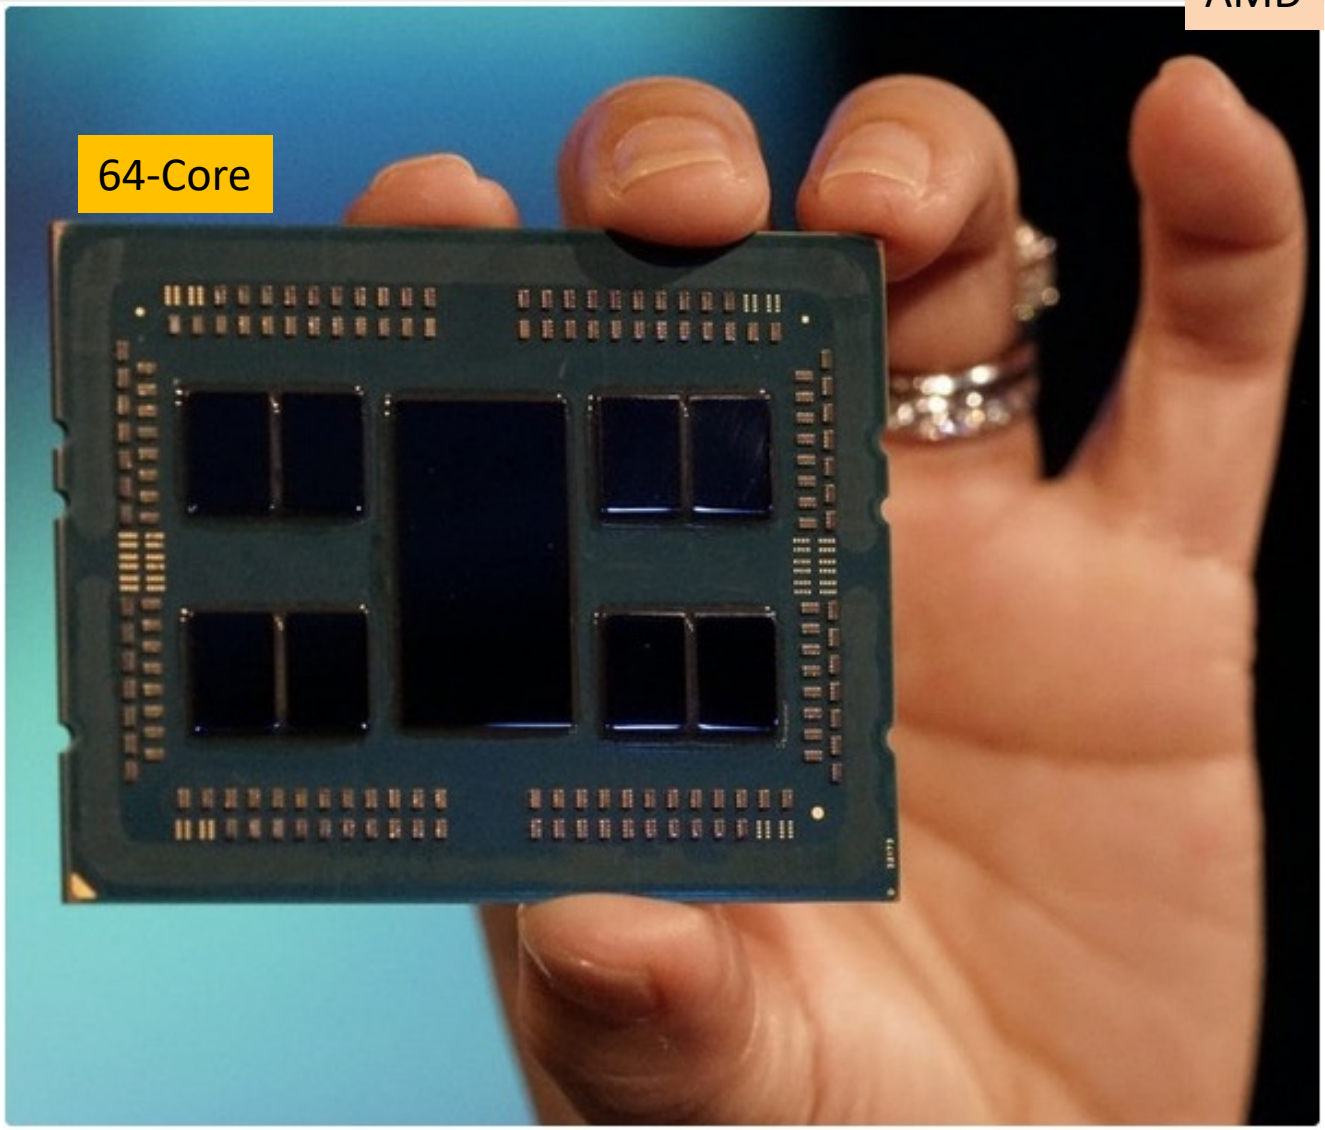

Lisa Su proudly shows off her 64-core EPYC monster at CES.

Multi-Cores + L2/L3

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Intel

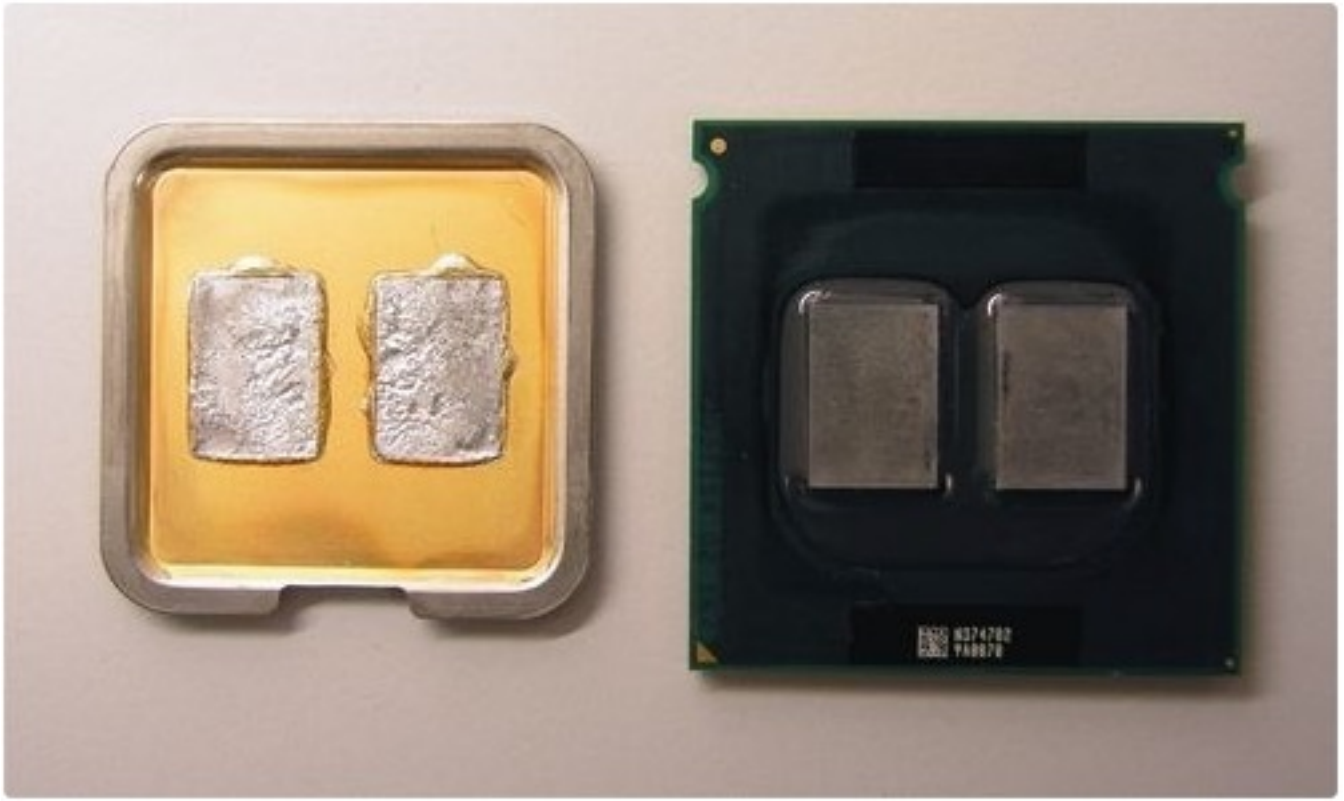

This is a delided Core 2 Quad with the two Core 2 Duo dies.

Multi-Cores + L2/L3

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Intel

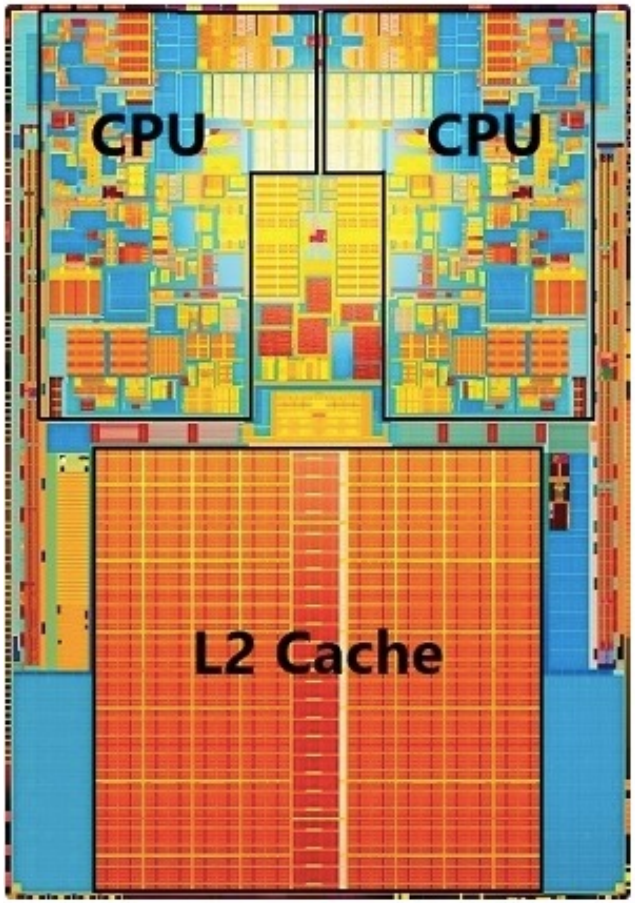

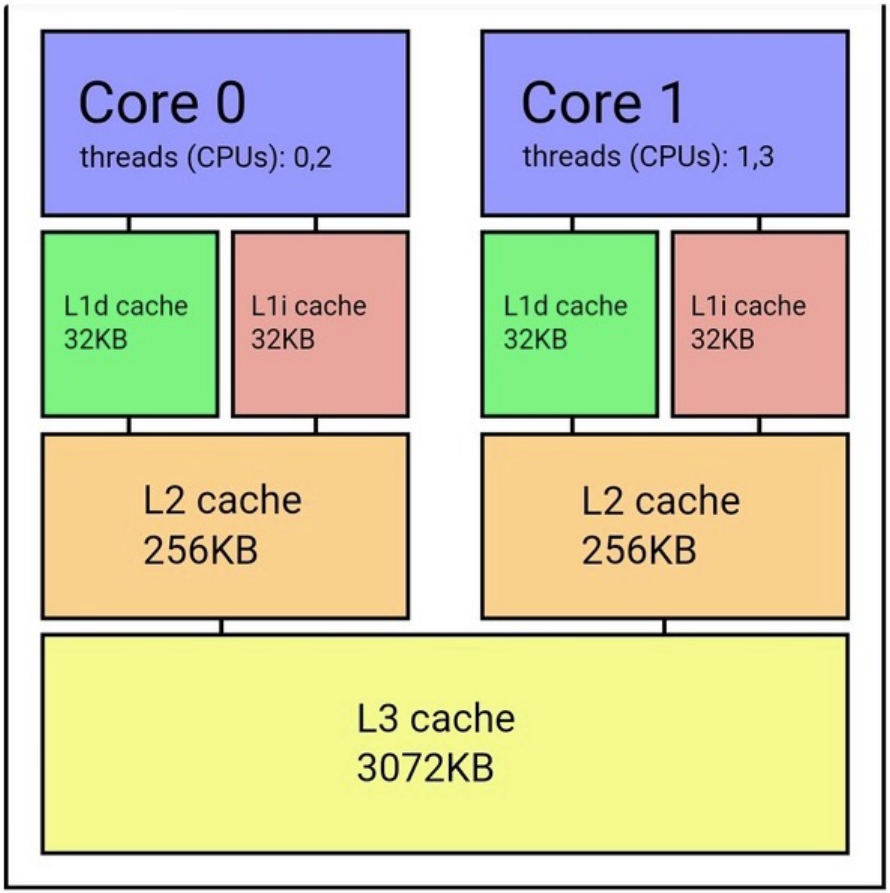

This is a late Core 2 Duo (Wolfdale), note how massive the uncore L2 is compared to the actual CPU cores. (uncore = outside the CPU)

Multi-Cores + L2/L3

Intel

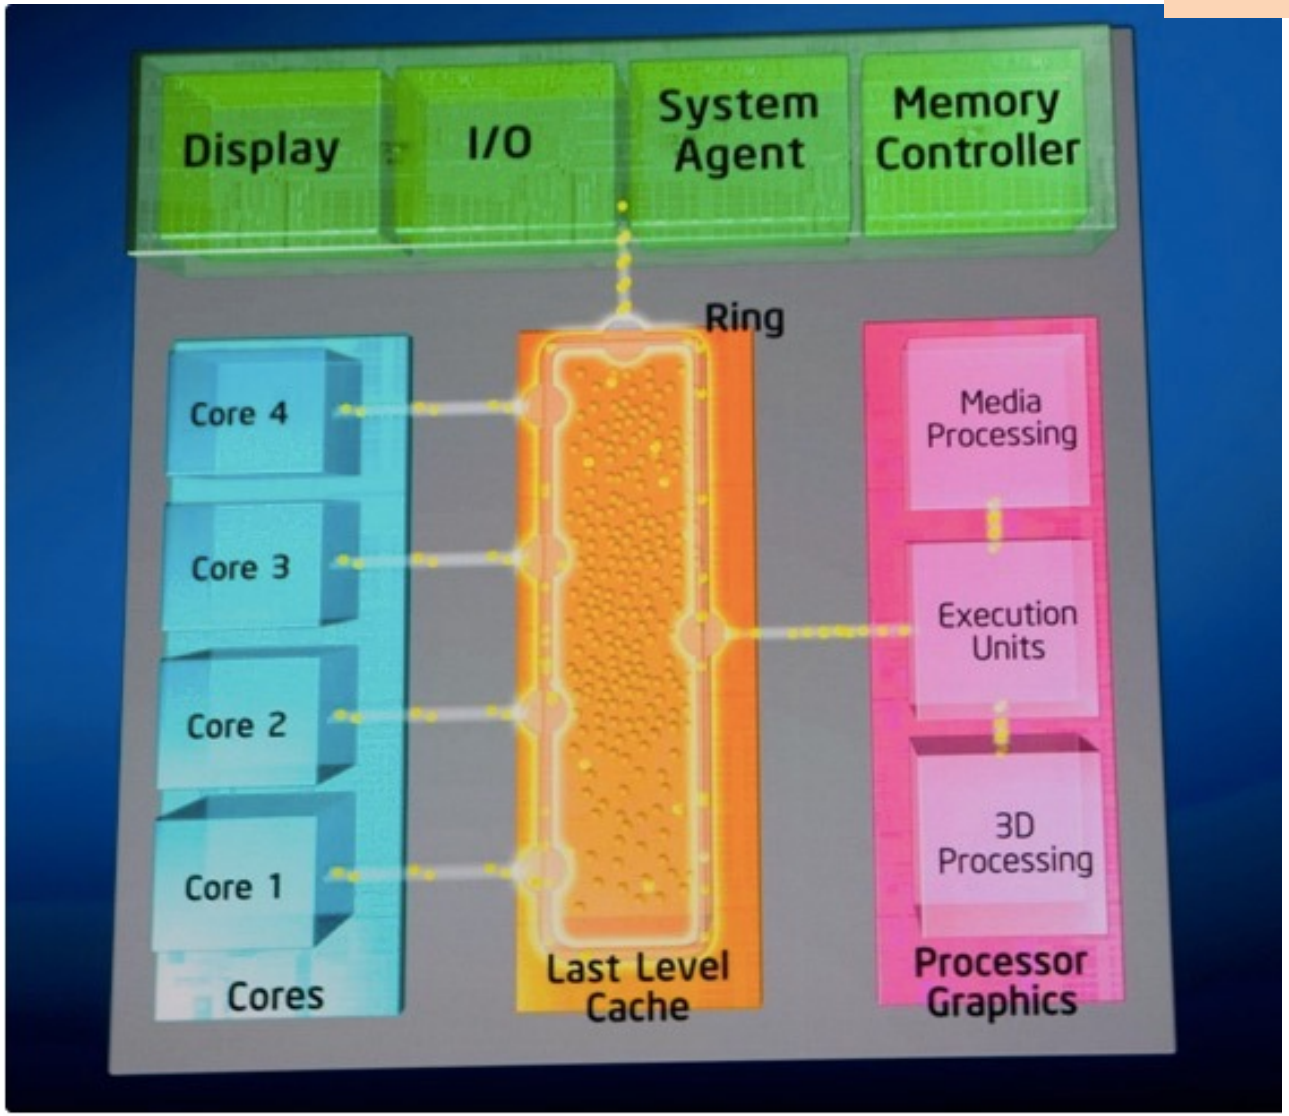

This special 'glue' remains in use today in all non-HEDT Intel CPUs.

Multi-Cores + L2/L3

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Intel Core 4

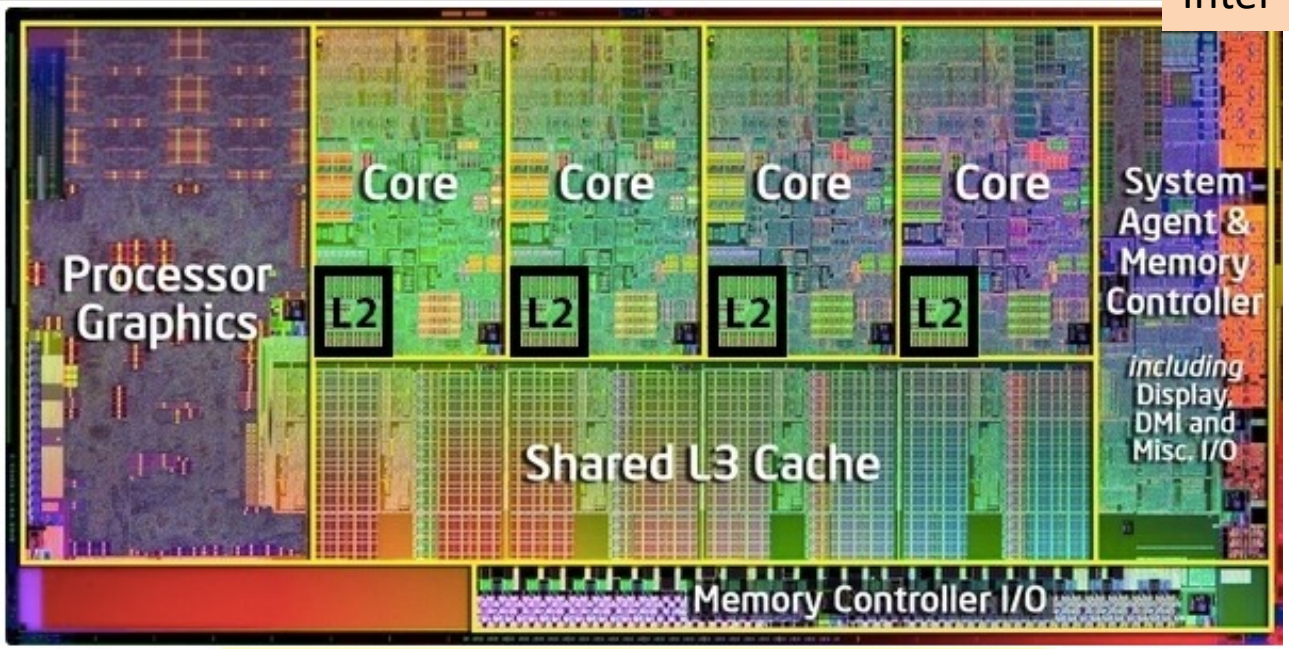

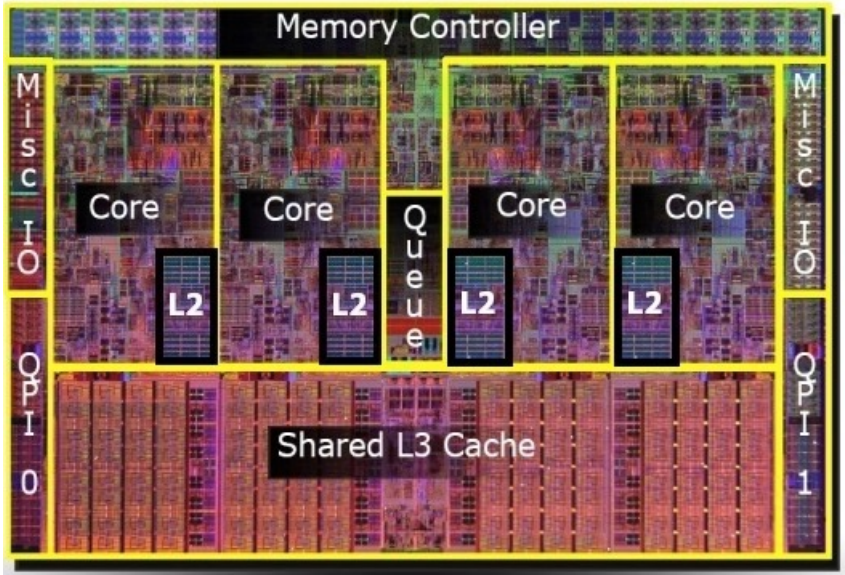

Multi-Cores + L2/L3

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD | Zen 2

AMD is taking a creative approach with its "7nm" production by keeping the most problematic part of the CPU at 14nm and putting the scalable "core" segments on 7nm octa-core "chiplets" that can be used alone or doubled up with another octa-core chiplet or (shhh...) a Navi GPU.

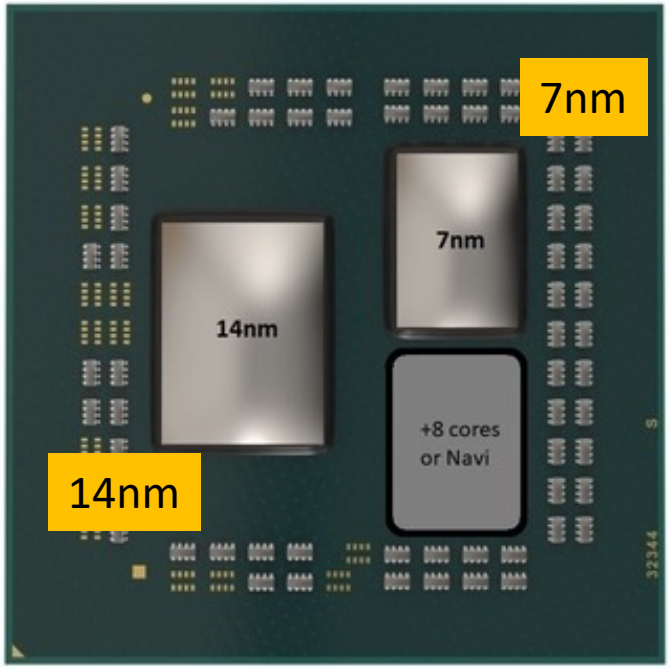

AMD kills two birds with one stone thinking outside the box with this creative Zen 2 layout: silicon yield at 7nm is vastly improved by making smaller chips, and the final product can be reconfigured to produce either workhorse 16-core CPU's or excellent SOC chips with a massive Navi GPU.

Multi-Cores + L2/L3

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

AMD

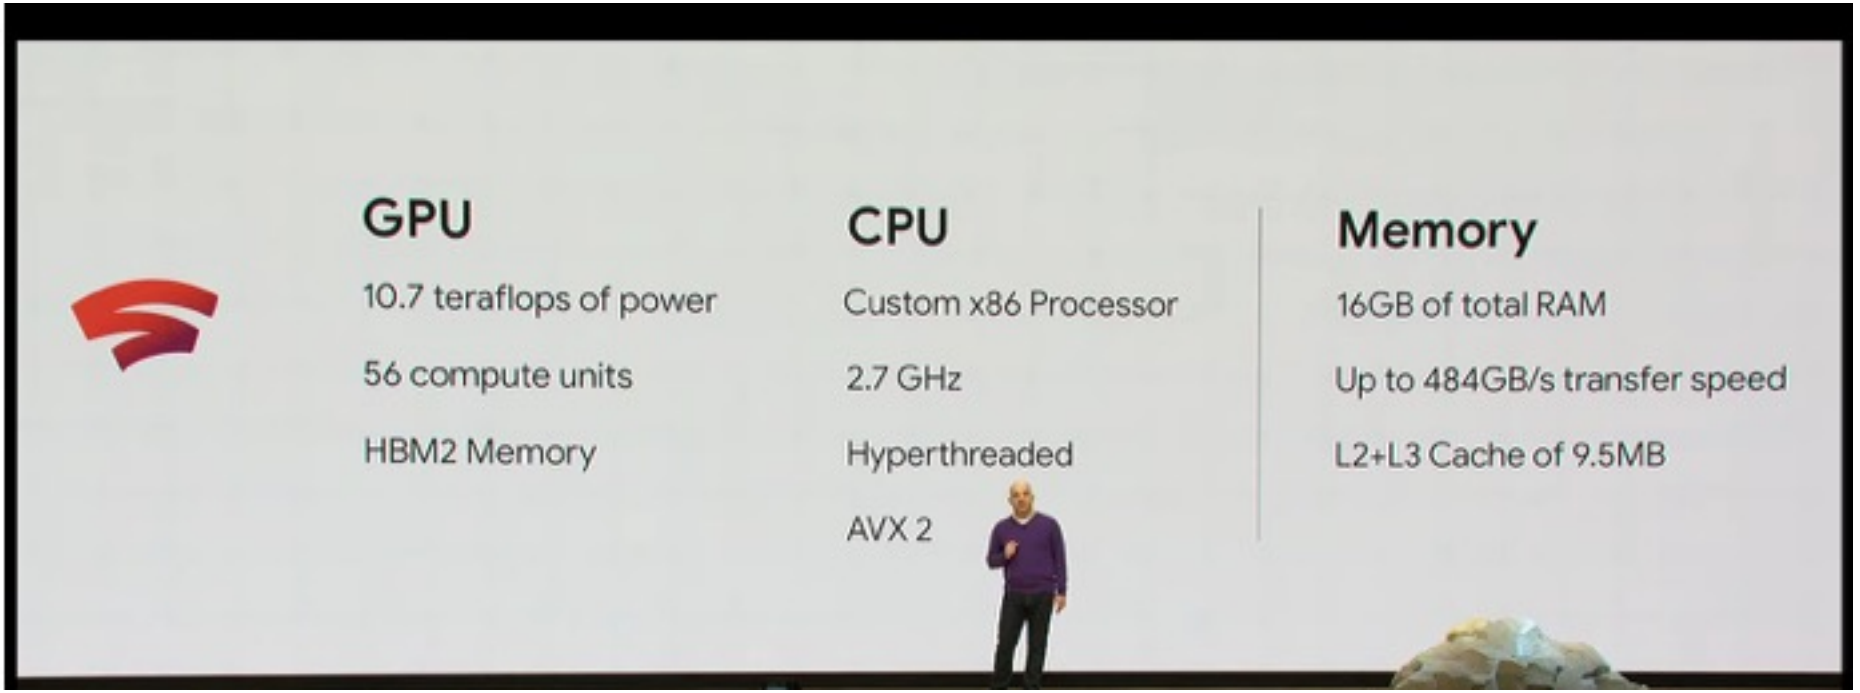

Intel Future CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Xeon

Intel

Source: Intel

Next-Gen Intel Xeon Scalable Processors to Deliver Breakthrough Platform Performance

This week, Intel announced its next-generation Intel Xeon Scalable processor platform. Codenamed Cooper Lake, the latest addition to the processor family boasts up to 56 processor cores per socket, higher memory bandwidth, and built-in AI inference and training acceleration. Learn more about the high-core-count Cooper Lake processor platform, which will be available in the first half of 2020.

<= 56 cores in 1H 2020

Intel Future CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Ice Lake | Intel

Source: TechRadar

Intel Ice Lake Has Landed, and It's Smarter, Faster and More Efficient Than Ever

Built on 10nm technology, Intel's Ice Lake platform is the first step toward Project Athena—an initiative that aims to leverage next-generation tech for advanced mobile computing experiences. Ice Lake's deep-learning capabilities will deliver optimized performance while the inclusion of Gen11 graphics will make for more sophisticated gaming experiences. Here's what else you can expect from Ice Lake.

Intel Future CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Intel

Source: Intel IT Peer Network

Exascale Computing Will Redefine Content Creation

In line with Intel's pursuit of exponential increases in computing capabilities, Intel's Architecture, Software, and Graphics group has been challenged to deliver a 1000x workflow improvement for content creators. Enter exascale computing. With the ability to process a quintillion (that's a billion billion, for those who can't count that high) calculations per second, exascale computing will enable content creators to push the limits for high-quality animation.

Exascale = 10^18 IPS

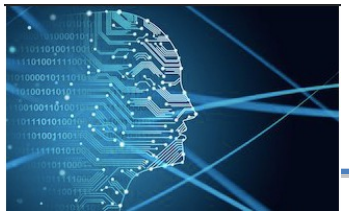

Intel Future CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Hot Chips 2019 | Intel

AT HOT CHIPS, INTEL PUSHES 'AI EVERYWHERE'

What's New: At Hot Chips 2019, Intel revealed new details of upcoming high-performance artificial intelligence (AI) accelerators: Intel® Nervana™ neural network processors, with the NNP-T for training and the NNP-I for inference. Intel engineers also presented technical details on hybrid chip packaging technology, Intel® Optane™ DC persistent memory and chiplet technology for optical I/O.

Source: Intel

Intel Future CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Hot Chips 2019 | Intel

What Intel Presented at Hot Chips 2019:

Intel Nervana NNP-T: Built from the ground up to train deep learning models at scale: Intel Nervana NNP-T (Neural Network Processor) pushes the boundaries of deep learning training. It is built to prioritize two key real-world considerations: training a network as fast as possible and doing it within a given power budget. This deep learning training processor is built with flexibility in mind, striking a balance among computing, communication and memory. While Intel® Xeon® Scalable processors bring AI-specific instructions and provide a great foundation for AI, the NNP-T is architected from scratch, building in features and requirements needed to solve for large models, without the overhead needed to support legacy technology. To account for future deep learning needs, the Intel Nervana NNP-T is built with flexibility and programmability so it can be tailored to accelerate a wide variety of workloads – both existing ones today and new ones that will emerge. View the presentation for additional technical detail into Intel Nervana NNP-T's (code-named Spring Crest) capabilities and architecture.

Intel Nervana NNP-I: High-performing deep learning inference for major data center workloads: Intel Nervana NNP-I is purpose-built specifically for inference and is designed to accelerate deep learning deployment at scale, introducing specialized leading-edge deep learning acceleration while leveraging Intel's 10nm process technology with Ice Lake cores to offer industry-leading performance per watt across all major datacenter workloads. Additionally, the Intel Nervana NNP-I offers a high degree of programmability without compromising performance or power efficiency. As AI becomes pervasive across every workload, having a dedicated inference accelerator that is easy to program, has short latencies, has fast code porting and includes support for all major deep learning frameworks allows companies to harness the full potential of their data as actionable insights. View the presentation for additional technical detail into Intel Nervana NNP-I's (code-named Spring Hill) design and architecture.

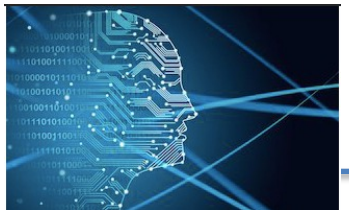

Intel Future CPUs

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Hot Chips 2019 | Intel

TeraPHY: An in-package optical I/O chiplet for high-bandwidth, low-power communication: Intel and Ayar Labs demonstrated the industry's first integration of monolithic in-package optics (MIPO) with a high-performance system-on-chip (SOC). The Ayar Labs TeraPHY* optical I/O chiplet is co-packaged with the Intel Stratix 10 FPGA using Intel Embedded Multi-die Interconnect Bridge (EMIB) technology, offering high-bandwidth, low-power data communication from the chip package with determinant latency for distances up to 2 km. This collaboration will enable new approaches to architecting computing systems for the next phase of Moore's Law by removing the traditional performance, power and cost bottlenecks in moving data. View the presentation for additional technical detail and design decisions on creating processors with optical I/O.

Intel Optane DC persistent memory: Architecture and performance: Intel Optane DC persistent memory, now shipping in volume, is the first product in the memory/storage hierarchy's entirely new tier called persistent memory. Based on Intel® 3D XPoint™ technology and in a memory module form factor, it can deliver large capacity at near-memory speeds, latency in nanoseconds, while also natively delivering the persistence of storage. Details of the two operational modes (memory mode and app direct mode) as well performance examples show how this new tier can support a complete re-architecting of the data supply subsystem to enable faster and new workloads. View the presentation for additional architectural details, memory controller design, power fail implementation and performance results for Intel Optane DC persistent memory.

Intel Moore's Tick Tock

Pipelined Architecture

staggered it allows a once a year cadence of new product introduction. One year the new chip is a process shrink of an existing architecture. The next year the new chip uses the same process node but introduces a new architecture.

The beauty of this chip development model is that it is a business innovation that takes advantage of a classic pipeline that hides latency to achieve throughput.

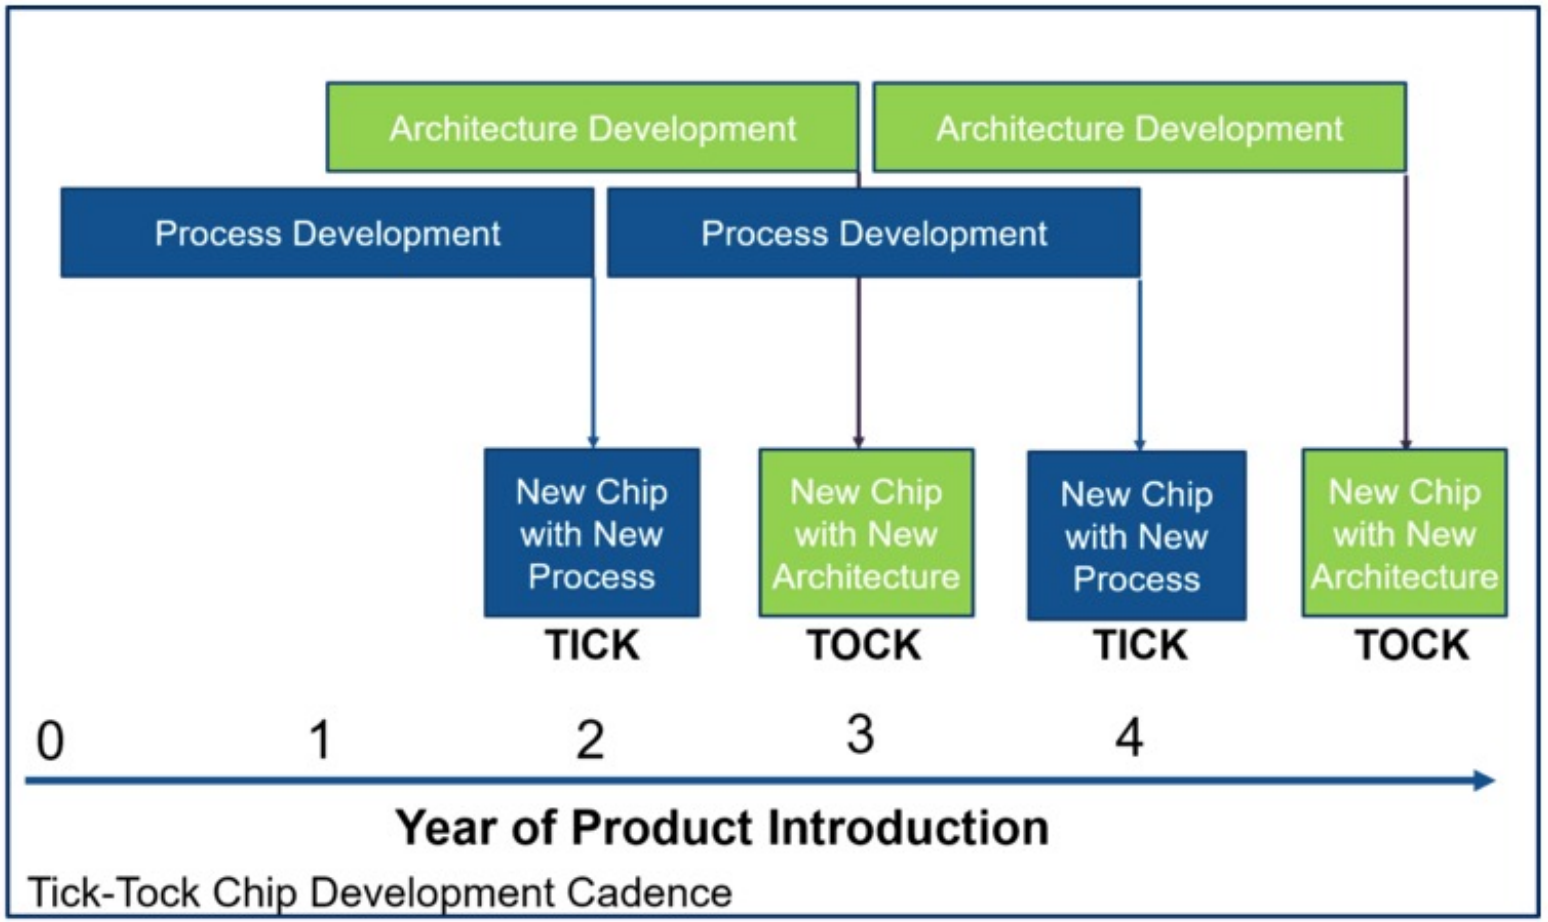

Tick-Tock Chip Development Cadence

Computer Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Modern Microprocessors

ARM → Apple

Apple ARM A Series

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quora post

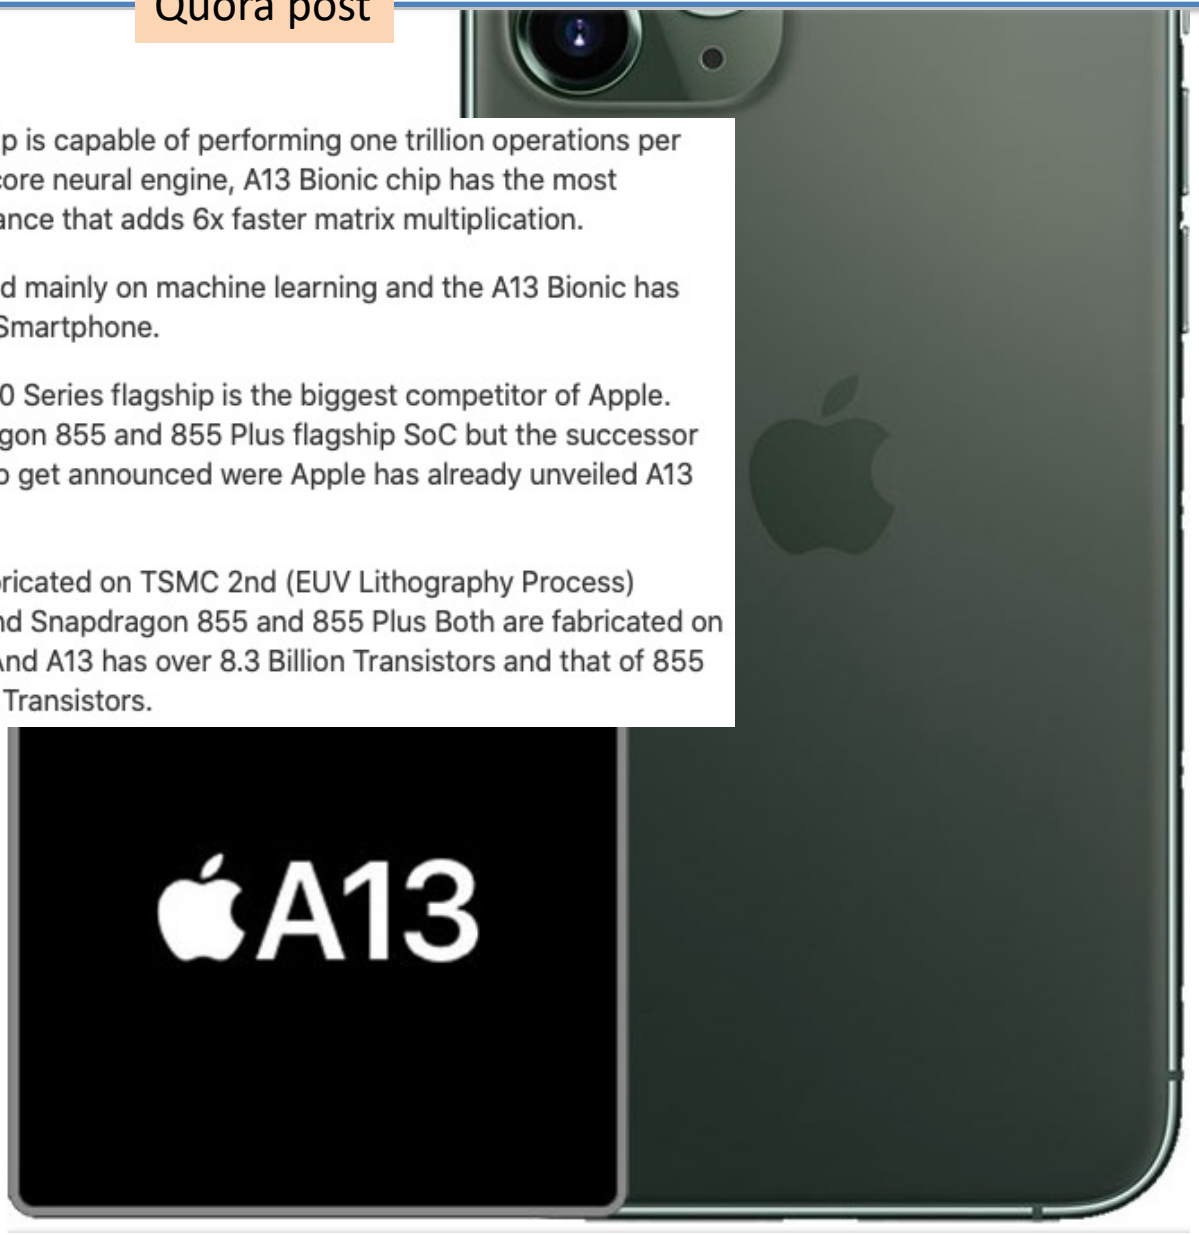

According to Apple, the chip is capable of performing one trillion operations per second and with an eight-core neural engine, A13 Bionic chip has the most Machine Learning performance that adds 6x faster matrix multiplication.

This time Apple has focused mainly on machine learning and the A13 Bionic has Fastest CPU and GPU in a Smartphone.

Qualcomm Snapdragon 800 Series flagship is the biggest competitor of Apple. Qualcomm has its Snapdragon 855 and 855 Plus flagship SoC but the successor to Snapdragon 855 is yet to get announced were Apple has already unveiled A13 Bionic.

The Apple A13 Bionic is fabricated on TSMC 2nd (EUV Lithography Process) Generation 7nm process and Snapdragon 855 and 855 Plus Both are fabricated on TSMC 7nm DUV process. And A13 has over 8.3 Billion Transistors and that of 855 Snapdragon has 6.9 Billion Transistors.

Apple ARM A*12/13* SoC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quora Home Answer 121 Spaces Notifications 45 Sea

3 Answers

Matthew J. Stott, Senior Systems & Mac Engineer (1996-present)
Answered Sep 30

It's called a SoC - System on Chip. It means the CPU package includes a lot more than just the CPU cores. What's changed with the A13 is even more power management abilities to shut off unused parts of the A13 but also right down to individual transistors as well. It is the most advanced power management in use right now. It is responsible for the excellent battery life of the 11, 11 Pro, 11 Pro Max iPhones. Yes, they increased the battery capacity a bit at the same time but that is just improved battery engineering.

Add to Yowan's A12X the Image Processing Core, a couple of Machine Learning accelerator cores and a bit less on the GPU with the A13 Bionic SoC. It is expected there will be an A13X for upgrade iPad Pros coming soon.

Apple ARM A*12/13* SoC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

- 8 CPU
- 7 GPU
- 1 NPU
- 2 MCU

This is a block diagram of the Apple A12X with 10 billion transistors. Of that amount, only 25% is dedicated to the two CPU clusters:

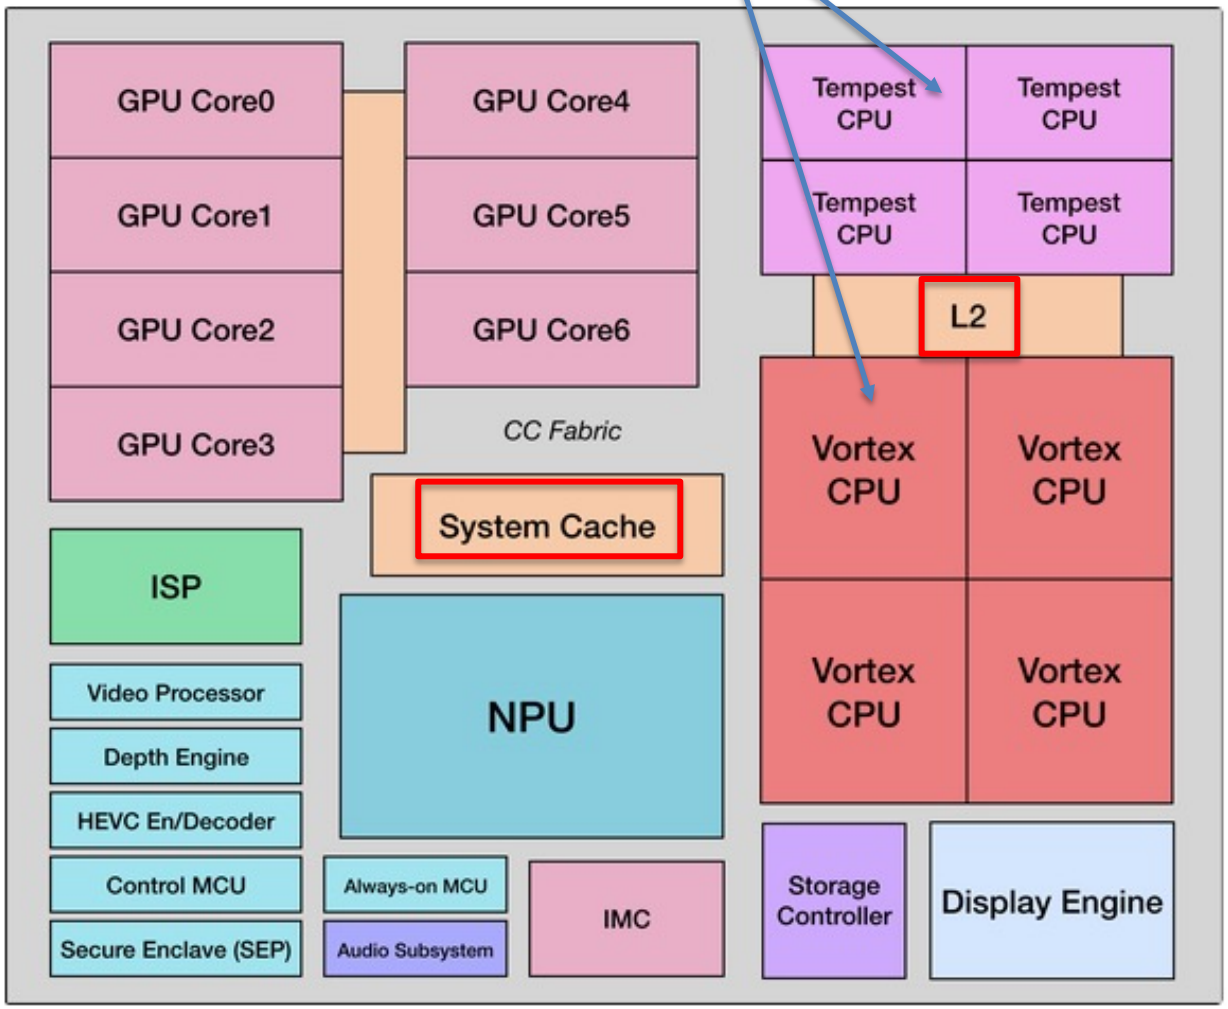

Image source: A12X Bionic - Apple - WikiChip

Apple A13

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quora post

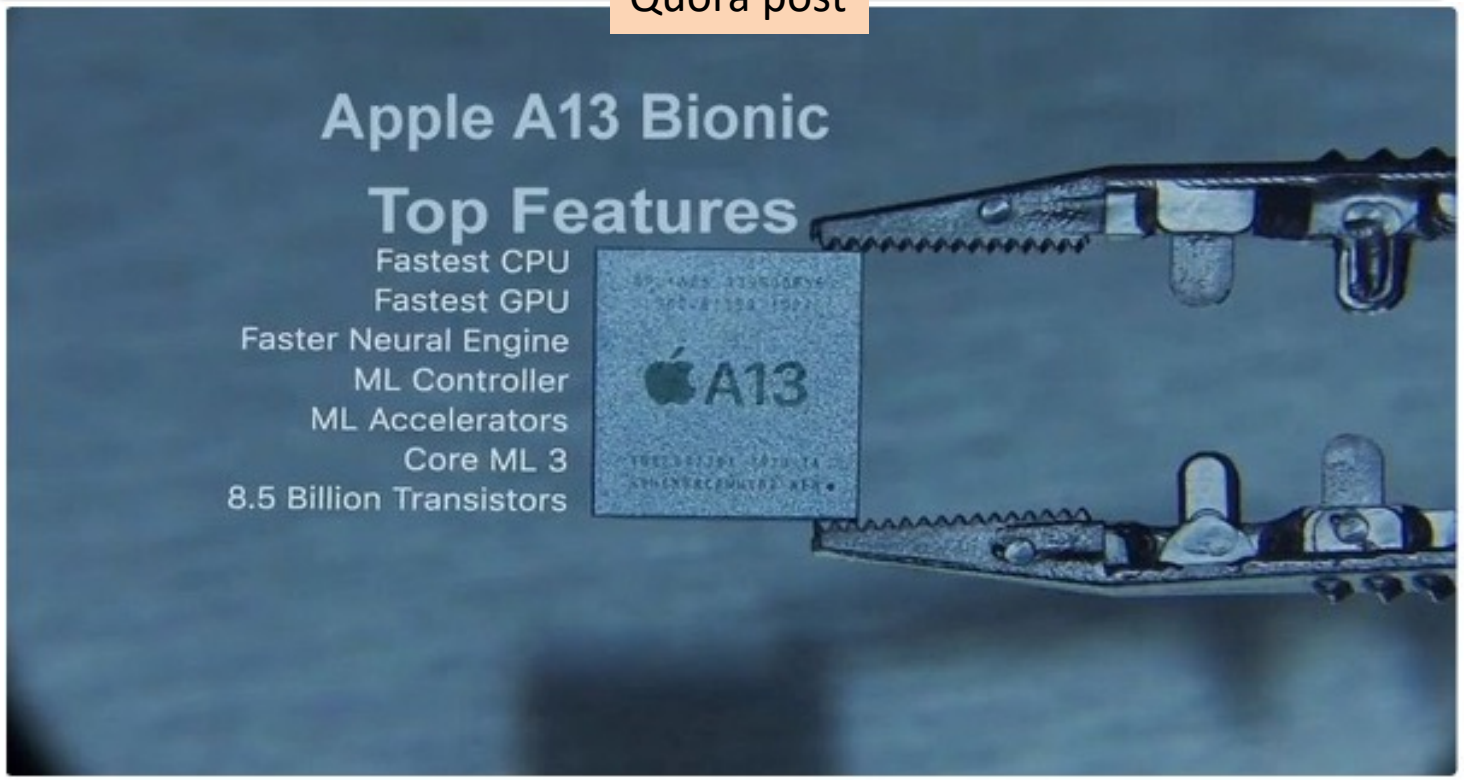

In Steve Jobs Studio, Apple launched three new iPhones - iPhone 11, iPhone 11 Pro and iPhone 11 Pro Max. All these three new iPhones are based on the new Apple A13 Bionic chipset. The company launched this SoC with these iPhones. The A13 Bionic is a successor of last year's A12 Bionic SoC. The last year's A12 Bionic is ahead of Snapdragon 855 in terms of performance. According to the company, Apple A12 Bionic is at least two years ahead of other Android smartphones in the race for fast processors. After all, what is new in the A13 Bionic that makes iPhones so powerful? Don't worry, here we ... (more)

Apple A13

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Quora post

And I think iPhones are going to be more power-efficient than Snapdragon 855 powered Android Phones. If you are asking about CPU, then the A13 Bionic is based on 64-bit Fusion Architecture. It is a Hexa-Core CPU with 2 Performance cores and 4 Efficiency cores. And it consumes 40% less power than the A12 Bionic. Coming to 855 Snapdragon, both Snapdragon 855 and 855 Plus is an ARM 64-bit SoC with Kryo 485 Octa-Core CPU. And it has Three CPU Clusters: 1 Cortex-A76 Prime Core, 3 Cortex-A76 Performance Cores and 4 Cortex-A55 Efficiency Cores. From these it seems Snapdragon 855 will definitely be a strong competitor for the Apple A13 Bionic Chip.

Moreover, while talking about GPU, for Apple, it is an Apple-designed Quad Core GPU and Snapdragon 855 has Adreno 640 GPU. And I don't think the Snapdragon will beat the performance of Apple's A13 bionic chip.

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Modern Microprocessors

IBM z15

IBM CPU's: z15

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

z15 floorplan:

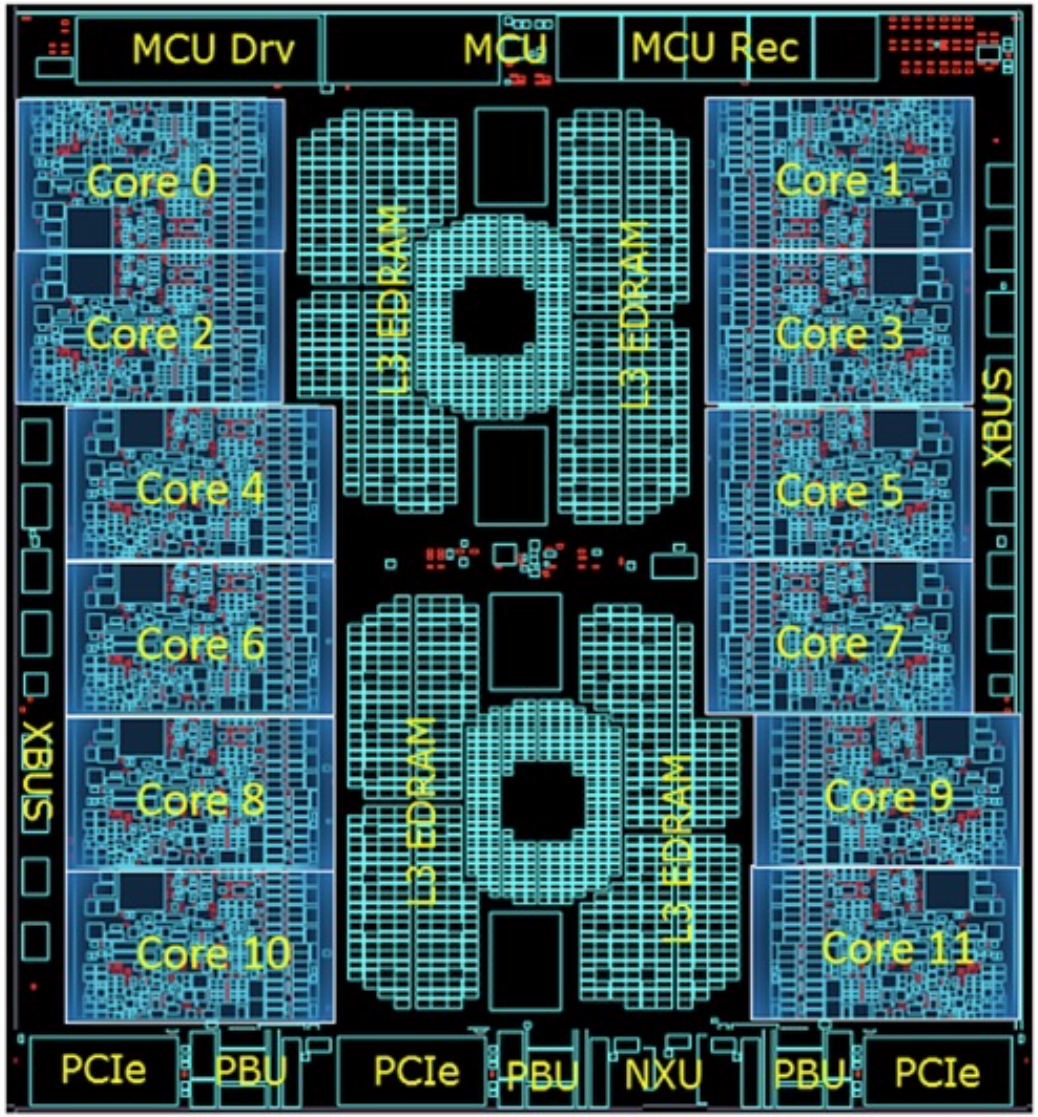

IBM CPU's: z15

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

The z15 is almost exactly as large as Intel's Skylake XCC config. 695 mm² vs ~ 694 mm² for the Intel one. Intel places 28 cores vs the z15's 12. That alone tells you that IBM is placing a lot more resources per core than Intel or AMD. (That's not good or bad per se, just an observation)

An individual z15 core:

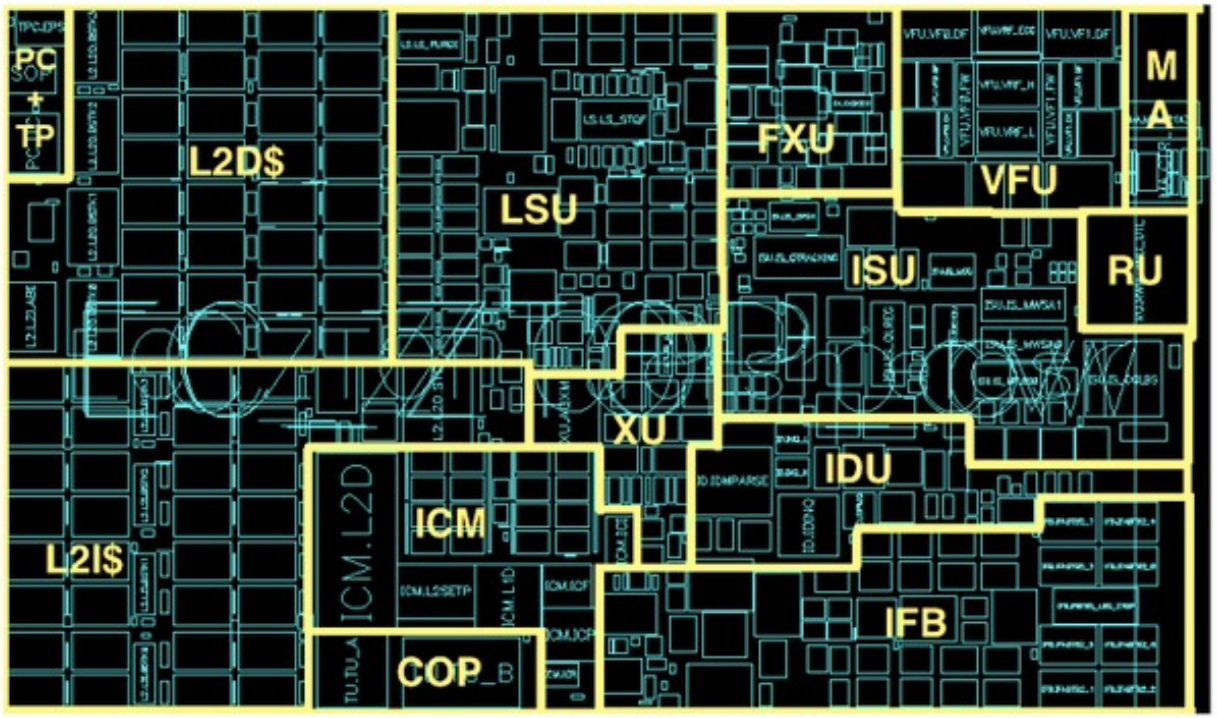

Source: WikiChip

Again, note the large L2 caches. In other CPUs that amount of real estate would typically be allocated for L3 cache.

IBM CPU's: ISA

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Instruction Set Architecture

While AMD and Intel are using the x86–64 Instruction Set Architecture (ISA), IBM mainframe CPUs use the z/Architecture ISA, a tradidtional CISC architecture, which is a descendant of the original S/360 ISA released all the way back in 1964. IBM has another RISC ISA called IBM POWER and produces processors using it, but the two aren't related. Nothing too special about that, there are plenty of ISAs out there.

Microarchitecture

When looking at the microarchitecture itself - or at least at what IBM is willing to publish - you'll see that IBM is trying to get the best of all worlds: A rich CISC instruction-set with over two thousand instructions and addressing modes, some of them extremely specific and complex. Combine that with very conservative instruction break-ups in the decode stage into µOPs, meaning that most instructions are actually implemented directly in hardware without the use of microcode (IBM calls it *millicode* for whatever reason). Add in very deep pipelines and you basically have a *braniac* CPU core, that tries to extract as much instruction and data-level parralelism as possible and that is both a *speed-demon*, running at insane clockspeeds. IBM manufactures the highest-clockspeed CPUs in the industry. The IBM z12 clocked in at a whopping 5.5 GHz for a six-core CPU, and that was sustained under all circumstances for all cores. No Turbo mode here. 5.5 GHz for six cores. All the time. The newest z15 manages 5.2 GHz across twelve cores. Under all circumstances, all the time.

IBM CPU's: SFU's

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Decimal

Special Function Units

IBM zSeries CPUs also contain dedicated executions units not found on many other architectures. One example would be decimal floating point and fixed point units, in addition to the more traditional binary FP units found in Intel or AMD machines. The reason for that is that IBM mainframes are traditionally used in the banking and insurance industries, both of which deal with handling high volume computations that represent actual money. A problem with binary floating point units is that due to the way they work, some fractions of real numbers cannot be accurately represented using a binary format. Just as the fraction 1/3 cannot be accurately represented in the decimal system, producing 0.3333333... and being usually truncated to 0.333...334 in decimal representation, the value 1/10 cannot be accurately represented in binary. It too produces an infinite sequence of repeating numbers, that have to be truncated. Problem: Dollars aren't measured in "thirds", but "tenths" and "hundreths" of a Dollar. These values routinley come up in interest or premium or tax calculations. You can now live with ever more compundung errors or implement your money-calculations using some type of software library, which costs a lot of performance, or you can use an architecture that has dedicated decimal floating point units handling all of that in hardware. Another SFU would be dedicated compression and encryption units per core. While both AMD and Intel have implemented some form of AES instructions, IBMs are much more comprehensive, handling more encryption standards and compression operations in hardware.

IBM CPU's: Cache

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

MLM

Cache Hierarchy

One unusual feature of IBM Mainframe CPUs is their cache hierarchy. Both in scope and implementation. L1D and L1I caches are fairly large at 128 kB each. L2 caches are huge in comparison to Intel and AMD CPUs, and they are also split into dedicated L2D and L2I caches, a whopping 4MB each per core. Compare that to the usual 1 - 2 MB unified L2 per core for Intel or AMD. The L3 cache is also special in that instead of using SRAM cells, the L3 cache is implemented using embedded DRAM, which increases latency, but allows IBM to place much more L3 cache into a given silicon area. For the z15, that is 256 MB of unified L3 cache per CPU chip. Really unique to the IBM zSeries cache hierarchy is an additional off-chip unified L4 cache, also implemented using eDRAM, in case of the z15 that's 960 MB of L4, shared by four CPU chips. All of that cache is used to feed those beefy and fast execution units with as much data and instructions as possible to avoid bubbles in the deep pipelines.

IBM CPU's: Cache

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

L4 OFF-CHIP

L4 Cache Chip floorplan:

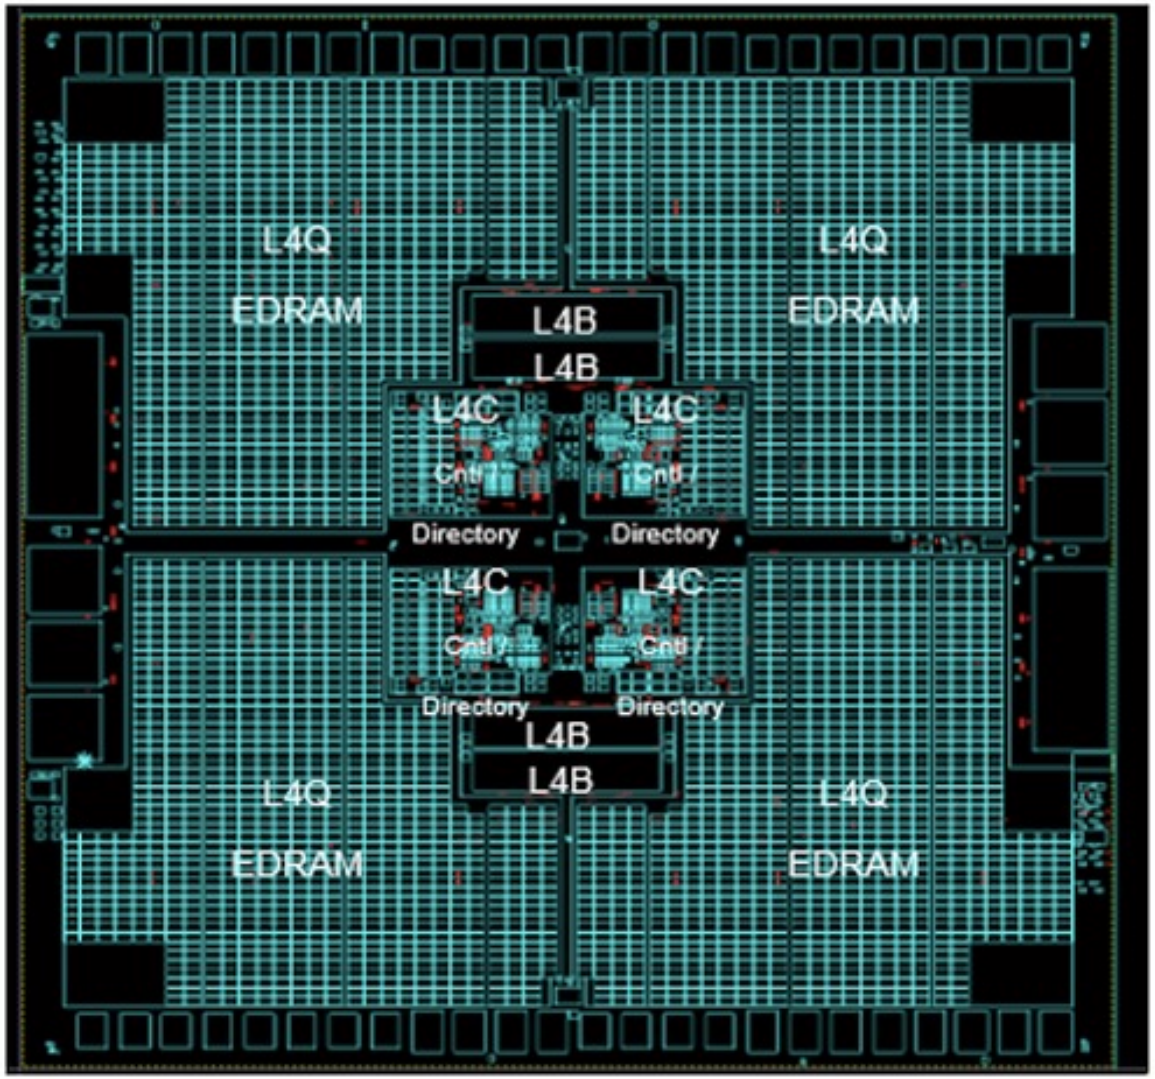

IBM CPU's: Memory ECC

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Memory Controllers

The memory controllers are almost one-of-a-kind: multi-channel memory controllers are nothing new or fancy, but the way IBM has implemented them in the zSeries is almost unique: each memory controller - one per chip - can handle 5 channels. The special sauce in this case is that those five channels are kind of RAIDed together. IBM calls it Redundant Array of independent Memory (RAIM). In addition to the usual SECDED ECC algoritms to protect against single chip memory errors, RAIM can tolerate an entire memory channel failure on the fly, without any OSes or applications noticing. Losing a memory channel on pretty much any other architecture will lock up your machine for good. AMD and Intel use memory mirroring to protect against that, but if you enable it you'll lose half your memory capacity and cut your memory bandwith in half. IBMs approach means that you'll only be limiting yourself to 4/5-ths of the installed capacity and bandwith. Pretty neat. The last general purpose CPU that could do this was the Alpha CPU, but they were phased out by HP in 2003 in favor of the ill-fated Itanium.

RAS features

Anyone producing a server-grade, general-purpose CPU will need to address the problem of soft and hard failures inside their CPUs. DECTED ECC for the caches, SECDED for less critical data structures and at least parity checking for non-critical ones. Error and sanity checkers in critial pathways. That's pretty standard. IBMs zSeries CPUs take that to another level. IBM estimates (whatever that may mean, they don't know?), that they dedicate up to 10% of silicon real estate to error detection and correction. An important part of that are the R-Units, another type of SFU that periodically stores the entire architectural state of the core in question, and upon non-transient failures can transfer that state to one of two spare cores present on all but the smallest mainframe models for them to pick up the load.

Computer Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jeff Drobman

CPU Performance

CPU Performance

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

$$\text{Time} = \text{Seconds/Program} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock cycle}}$$

$$\text{CPU time} = \text{Instruction count} \times \text{CPI} \times \text{Clock cycle time}$$

Clock rate = 1/Clock cycle time

$$\text{CPU time} = \frac{\text{Instruction count} \times \text{CPI}}{\text{Clock rate}}$$

CPU Performance

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

The following table summarizes how these components affect the factors in the CPU performance equation.

Hardware or software component	Affects what?	How?
Algorithm	Instruction count, possibly CPI	The algorithm determines the number of source program instructions executed and hence the number of processor instructions executed. The algorithm may also affect the CPI, by favoring slower or faster instructions. For example, if the algorithm uses more divides, it will tend to have a higher CPI.
Programming language	Instruction count, CPI	The programming language certainly affects the instruction count, since statements in the language are translated to processor instructions, which determine instruction count. The language may also affect the CPI because of its features; for example, a language with heavy support for data abstraction (e.g., Java) will require indirect calls, which will use higher CPI instructions.
Compiler	Instruction count, CPI	The efficiency of the compiler affects both the instruction count and average cycles per instruction, since the compiler determines the translation of the source language instructions into computer instructions. The compiler's role can be very complex and affect the CPI in varied ways.
Instruction set architecture	Instruction count, clock rate, CPI	The instruction set architecture affects all three aspects of CPU performance, since it affects the instructions needed for a function, the cost in cycles of each instruction, and the overall clock rate of the processor.

CPU Performance

Frames/sec

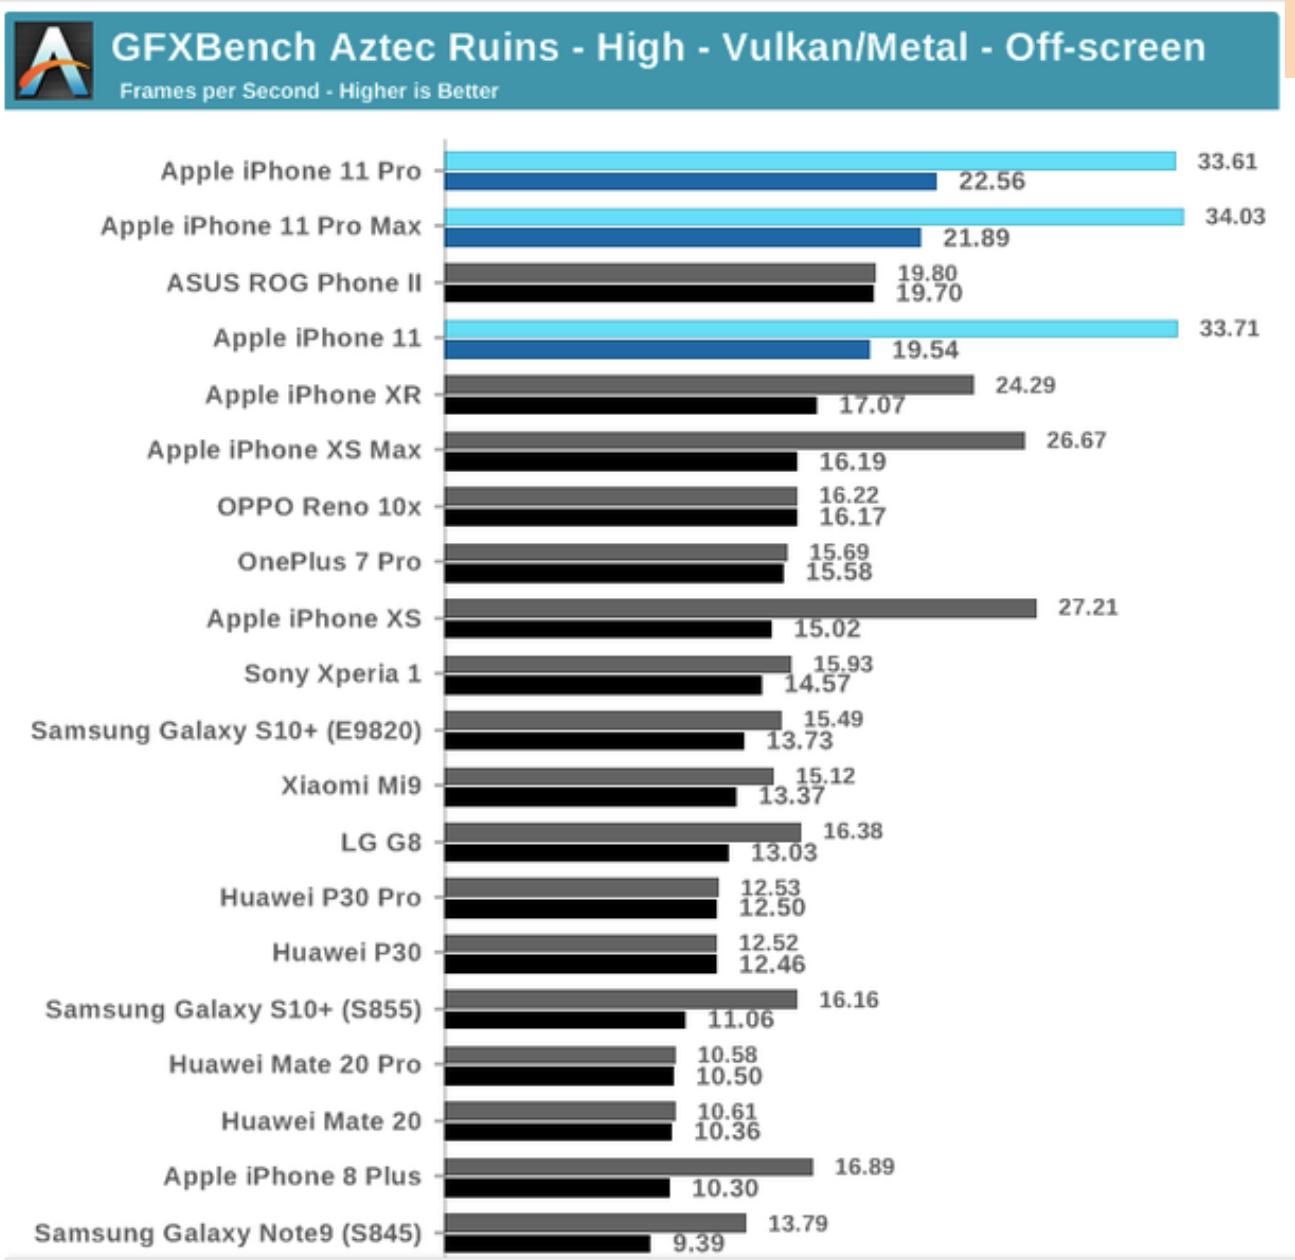

CPU Performance

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Frames/sec

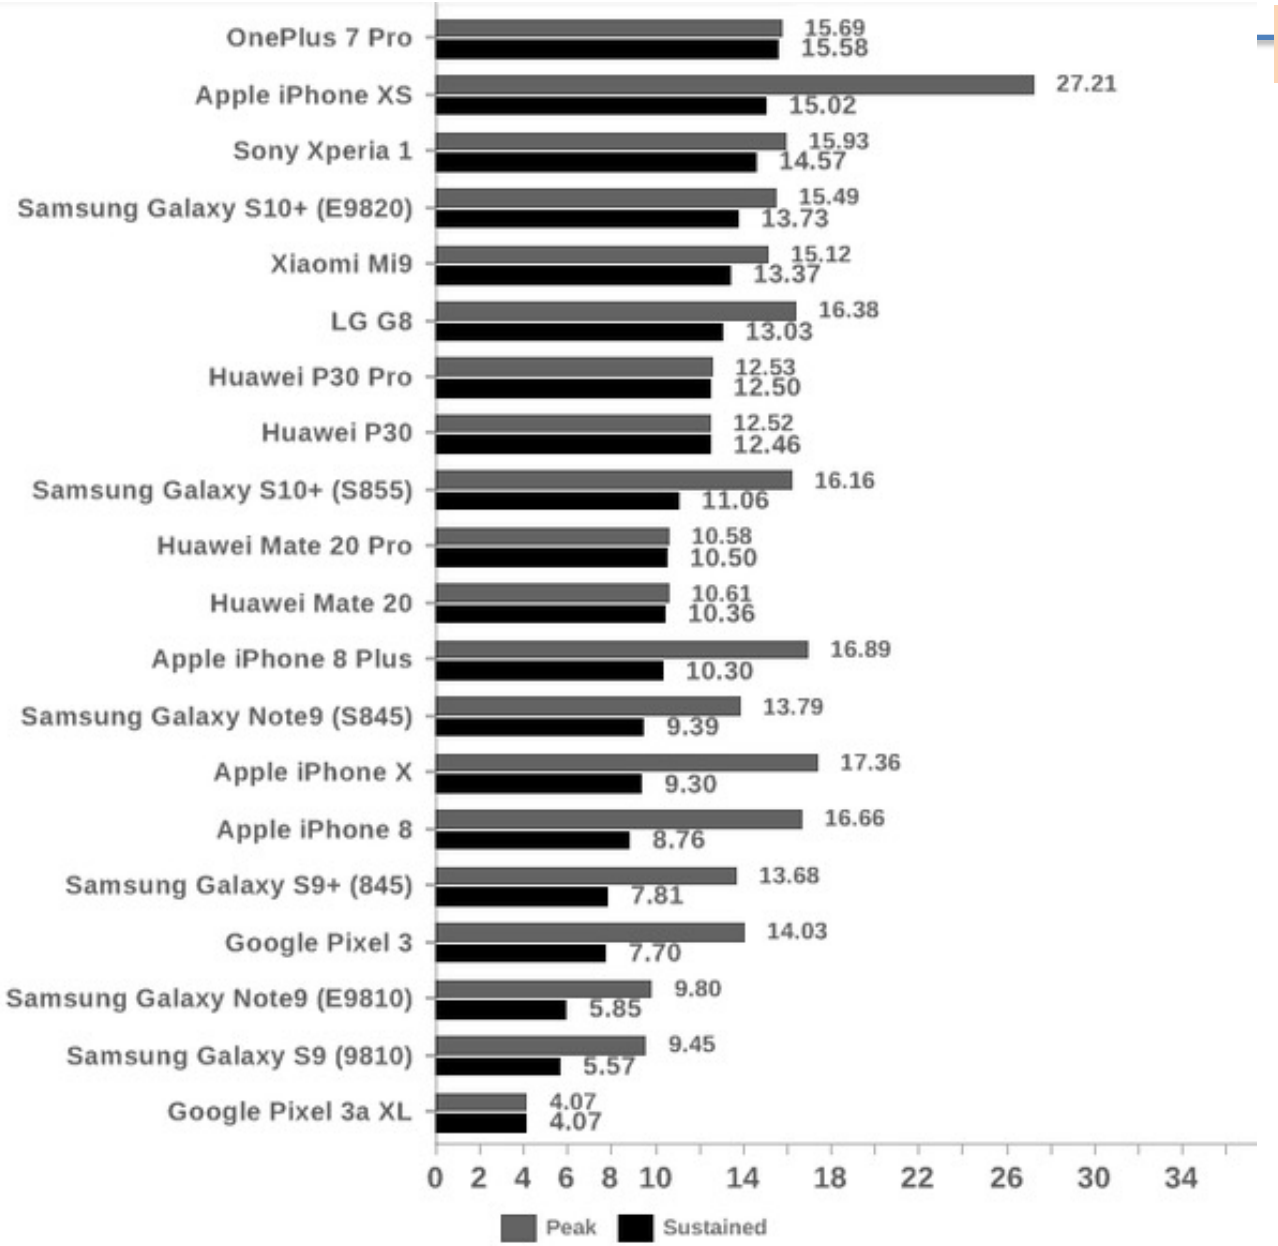

Computer Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jeff Drobman

Parallelism

I-D Parallelism: SIMD

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Flynn Partition

Michael J. Flynn paper (U Illinois (UIUC), Ca 1969)

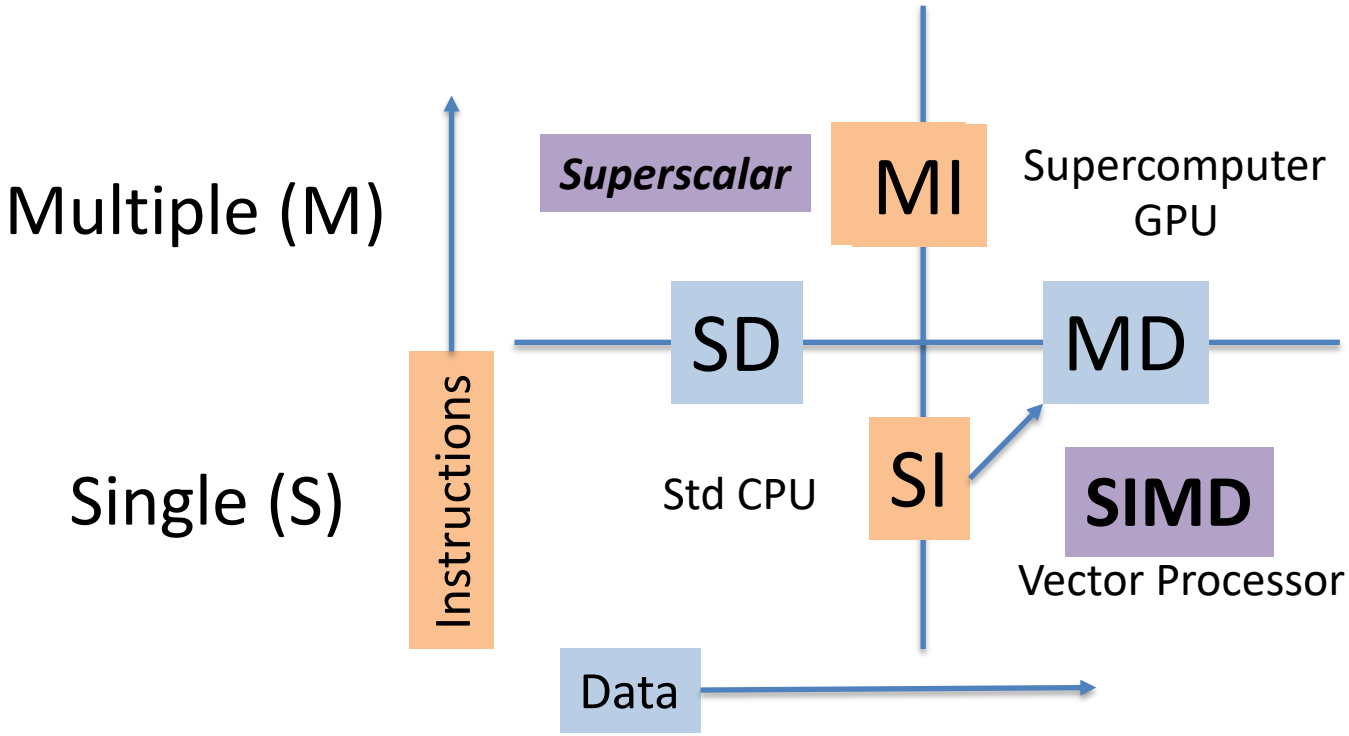

Instruction Level *Parallelism*

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Super- *Pipelining* **SISD**

- Split some pipeline stages (4-5 → 8-11)
- Faster clock cycle → higher *throughput* (*mips*)
- Affect CPI?

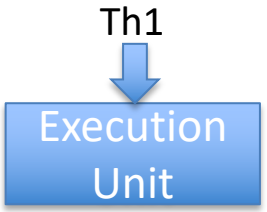

Super- *Scalar* **SISD**

- Multiple ***Execution Units*** (multi-issue pipelines)
  - each EU = ICU+ALU, with shared GR's
- Hardware + compiler *schedules* instruction streams

Multi- *Threading* **SISD**

- Multiple control threads (usually 2, same/dif program)
- Programs can *allocate* code to threads
- Automatic *scheduling* of control threads
- 2 types: *SMT/superscalar* or *temporal* (interleaved: coarse/fine)

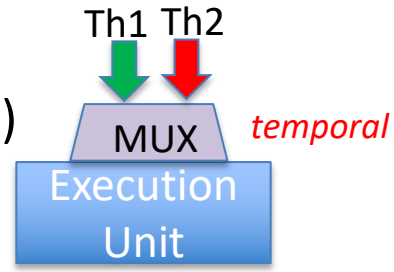

Multi- *Core* **MISD**

- Classic Parallelism: multiple copies of the CPU
- **Multiple L1/L2 caches** (one set per core)

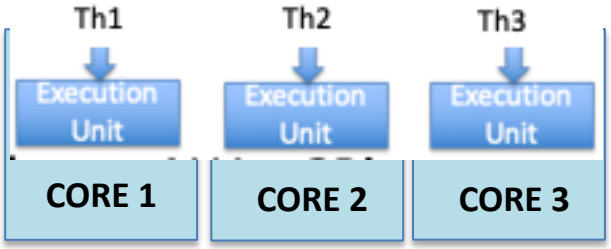

Superscalar

Super- *Scalar* SISD

- Multiple ***Execution Units*** (multi-issue pipelines)
  - each EU = ICU+ALU, with shared GR's
- Hardware + compiler schedules instruction streams

Th1	**EU1**	pipeline1
Th2	**EU2**	pipeline2

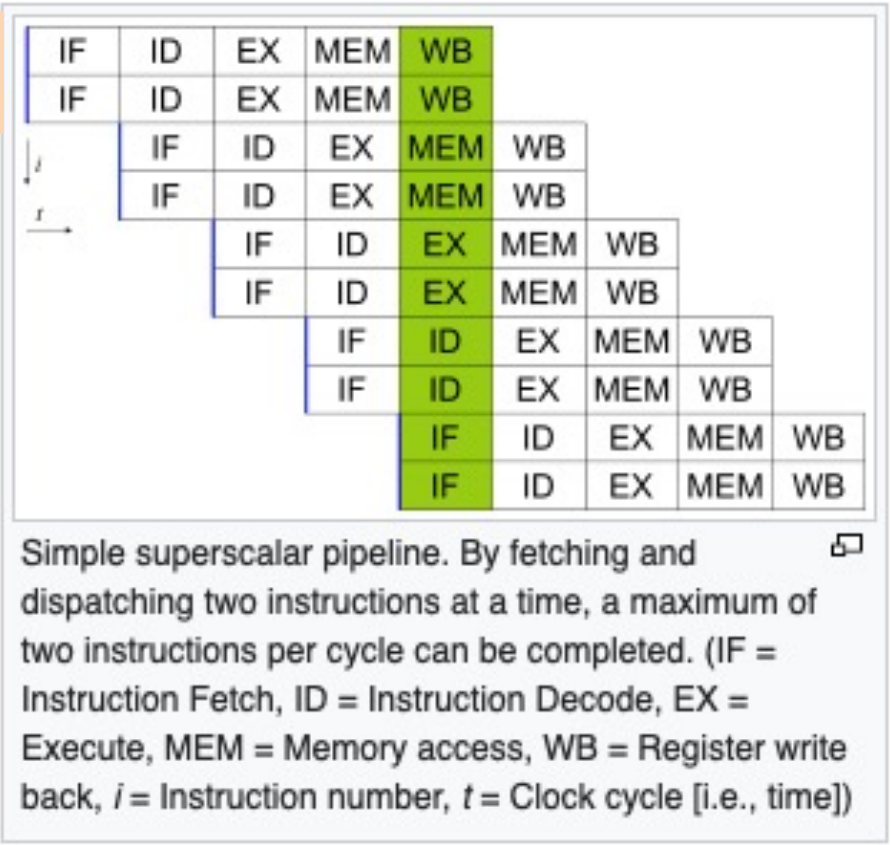

Simple superscalar pipeline. By fetching and dispatching two instructions at a time, a maximum of two instructions per cycle can be completed. (IF = Instruction Fetch, ID = Instruction Decode, EX = Execute, MEM = Memory access, WB = Register write back, *i* = Instruction number, *t* = Clock cycle [i.e., time])

Hyperthreading

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

WIKIPEDIA
The Free Encyclopedia

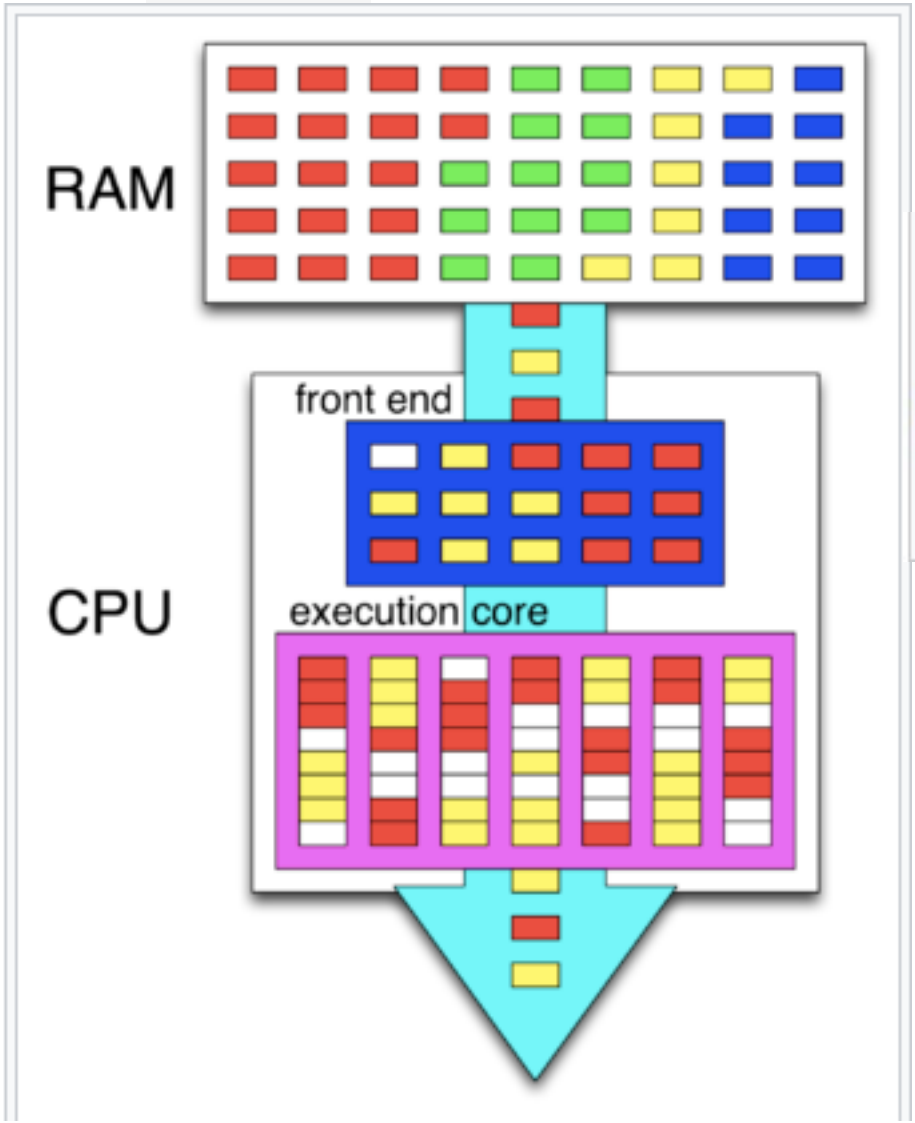

Virtual ***Thread*** Machine

Sits on top of Superscalar Multi-cores

In this high-level depiction of HTT, instructions are fetched from RAM (differently colored boxes represent the instructions of four different processes), decoded and reordered by the front end (white boxes represent pipeline bubbles), and passed to the execution core capable of executing instructions from two different programs during the same clock cycle.[1][2][3]

Intel's proprietary ***HTT***

Computer Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

GPU

❖ See more at “GPU” slide set

Special Processors

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

- MPU → CPU — Processor
- GPU
- CPU + GPU → APU — Graphics
- NPU — Network/Neural Processor
- DSP — DSP engine
- AI: GPU → TPU ?? — AI engine

Single FF (1961)

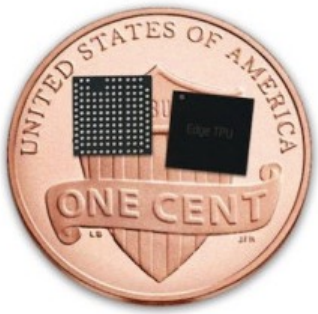

The Edge TPU, shown with a U.S. penny for reference

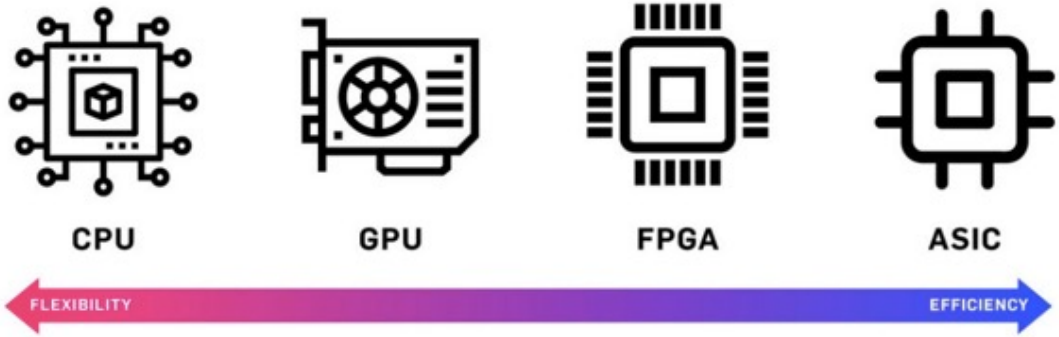

GPU Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

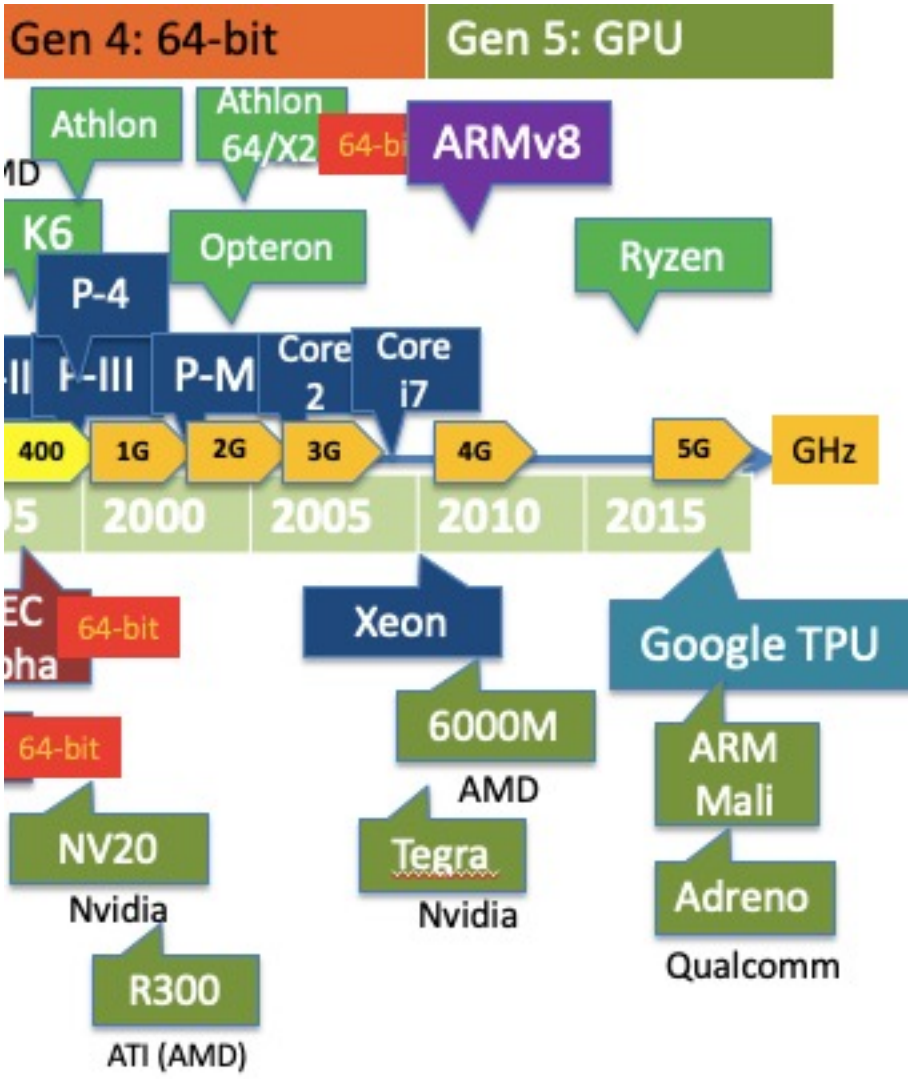

ARM Mali GPU

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Mali (GPU)

From Wikipedia, the free encyclopedia

The **Mali** series of graphics processing units (GPUs) and multimedia processors are semiconductor intellectual property cores produced by ARM Holdings for licensing in various ASIC designs by ARM partners.

Mali GPUs were developed by Falanx Microsystems A/S, which was a spin-off of a research project from the Norwegian University of Science and Technology.[1] Arm Holdings acquired Falanx Microsystems A/S on June 23, 2006 and renamed the company to Arm Norway.[2]

Technical details [edit]

Like other embedded IP cores for 3D rendering acceleration, the Mali GPU does not include display controllers driving monitors, in contrast to common desktop video cards. Instead, the Mali ARM core is a pure 3D engine that renders graphics into memory and passes the rendered image over to another core to handle display.

ARM does, however, license display controller SIP cores independently of the Mali 3D accelerator SIP block, e.g. Mali DP500, DP550 and DP650.[3]

ARM also supplies tools to help in authoring OpenGL ES shaders named *Mali GPU Shader Development Studio* and *Mali GPU User Interface Engine*.

Display controllers such as the ARM HDLCD display controller are available separately.[4]

Mali GPU Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Variants [edit]

The Mali core grew out of the cores previously produced by Falanx and currently constitute:

Model	Micro-architecture	Type	Launch date	Shader core count	Fab (nm)	Die size (mm^2)	Core clock rate (MHz)	L2 cache size
Mali-55/110	?	Fixed function pipeline[5]	2005 [permanent dead link]	1	?	?	?	N/A
Mali-200	Utgard[6]	Programmable pipeline[7]	2007[8]	1	?	?	?	N/A
Mali-300			?	1	40 28	?	500	8 KiB
Mali-400 MP			2008	1–4	40 28	?	200–600	8-256 KiB
Mali-450 MP			2012	1–8	40 28	?	300–750	8-512 KiB
Mali-470 MP			2015	1–4	40 28	?	250–650	8–256 KiB
Mali-T604[9]	Midgard		?	1–4	32 28	?	533	

Mali GPU Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Model	Micro-architecture	Type	Launch date	Shader core count	Fab (nm)	Die size (mm^2)	Core clock rate (MHz)	Max L2 cache size
Mali-G71		Unified shader model + Unified memory + scalar, clause-based ISA	Q2 2016	1–32	16 14 10	?	546-1037	128–2048 KiB
Mali-G52	Bifrost 2nd gen		Q1 2018	1-4 (2 or 3 EU per core)	7 16	?	850	
Mali-G72			Q2 2017	1–32	16 12 10	1.36 mm^2 per shader core at 10 nm[29]	572-800	128–2048 KiB
Mali-G76	Bifrost 3rd gen		Q2 2018	4-20	12 8 7	?	600-800	512–4096 KiB
Mali-G57	Valhall 1st gen	Superscalar engine + Unified memory + simplified scalar ISA	Q2 2019	1-6	7	?	?	64–512 KiB
Mali-G77			Q2 2019	7-16	7	?	850	512–4096 KiB

GPU Architecture

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

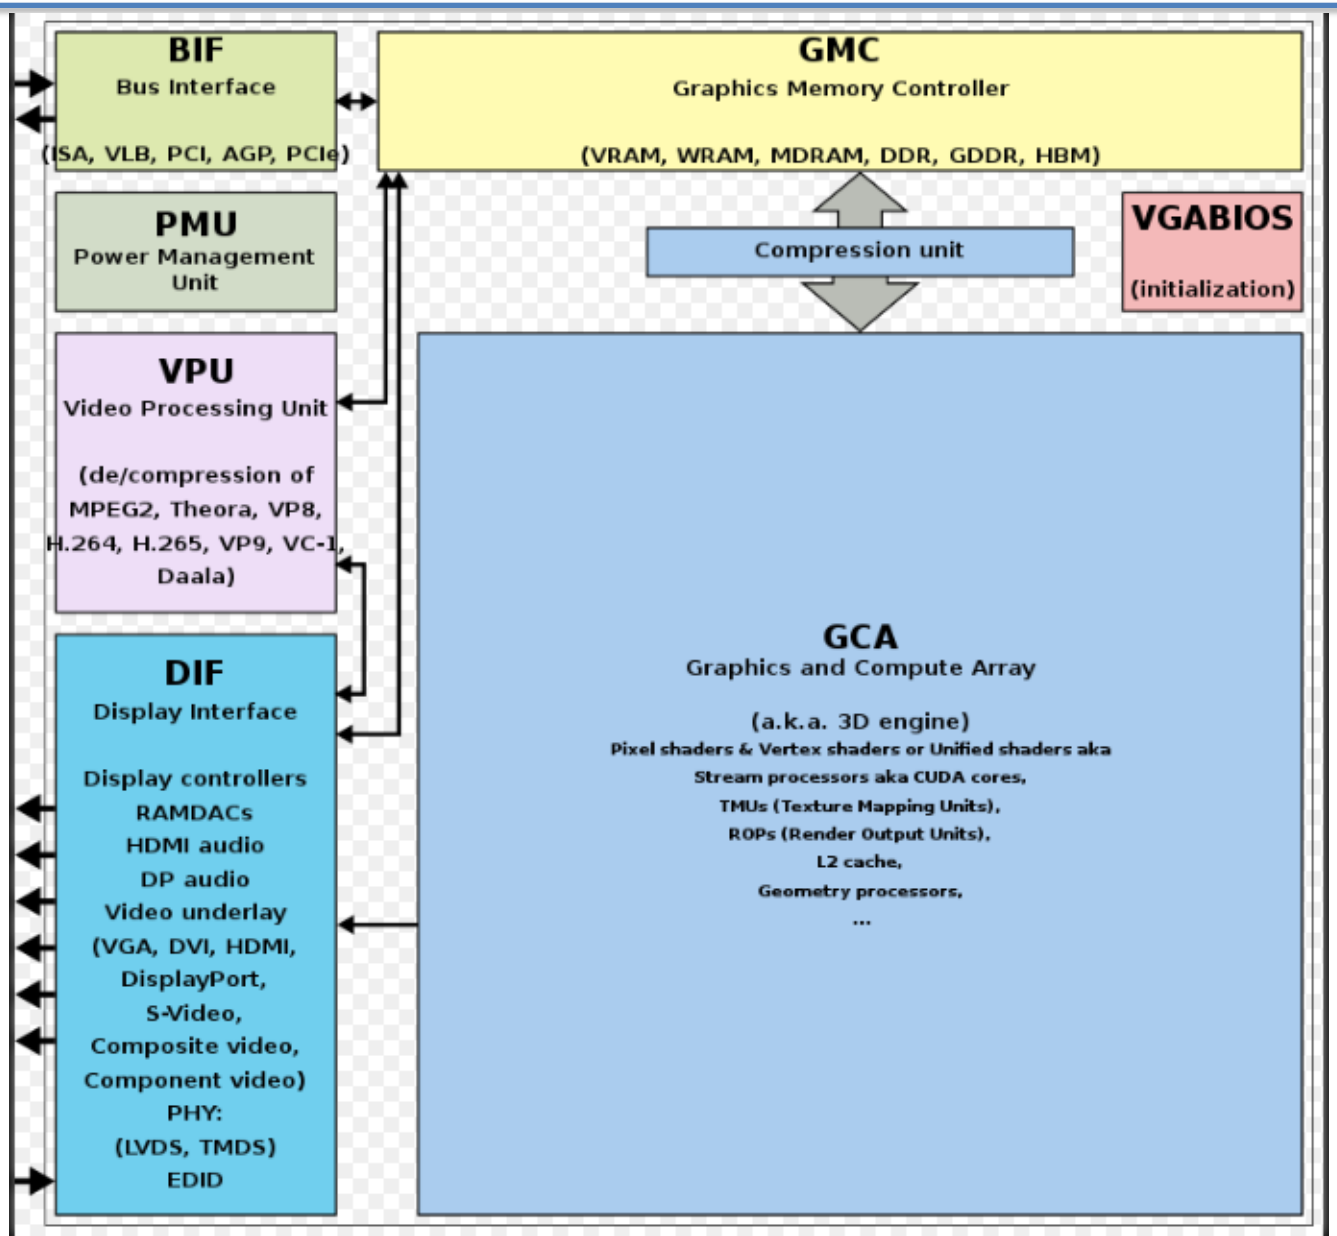

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

Phones

*see slide set on “SoC’s” for phone chips

Old Palm Pilot Phone

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

1st iPhone

DR JEFF SOFTWARE
INDIE APP DEVELOPER
Jeff Drobman

iPhone

ANNOUNCED: Jan. 9, 2007

RELEASED: June 29, 2007

KEY FEATURES:
3.5-inch diagonal screen;
320 x 480 pixels at 163 ppi;
2-megapixel camera

PRICE: 4GB model, $499;
8GB version, $599 (with
a two-year contract)