CSUN | CALIFORNIA STATE UNIVERSITY NORTHRIDGE

DR JEFF SOFTWARE
INDIE APP DEVELOPER

COMP222

COMP 222

Spring 2023 | Rev 3-8-23

Computer Organization

(Architecture)

Lecture 1B: Status

Dr Jeff Drobman

website ➡ *drjeffsoftware.com/classroom.html*

email ➡ *jeffrey.drobman@csun.edu*

Index (vol. 1B)

State of the Art

Industry Landscape

ISA/SoC Landscape

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)

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Tech Titan Timeline

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Annual Revenue in $B | *Historical* Perspective | As of 3Q2022

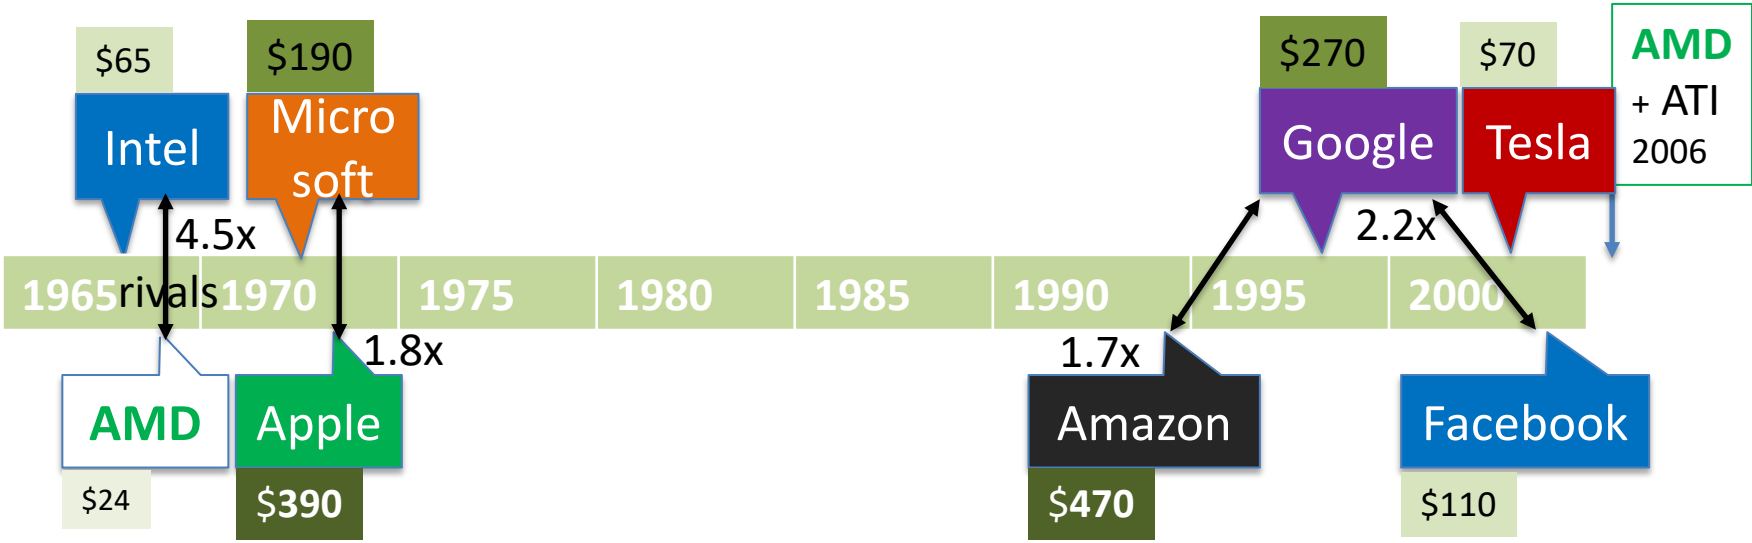

Eras: chips | PCs | Sales | Search | Social

Internet 1969 → delivered → 20 years → WWW 1989

Tech Titan Timeline

Qtr Revenue in $B | Current Perspective | **Q4** 2022

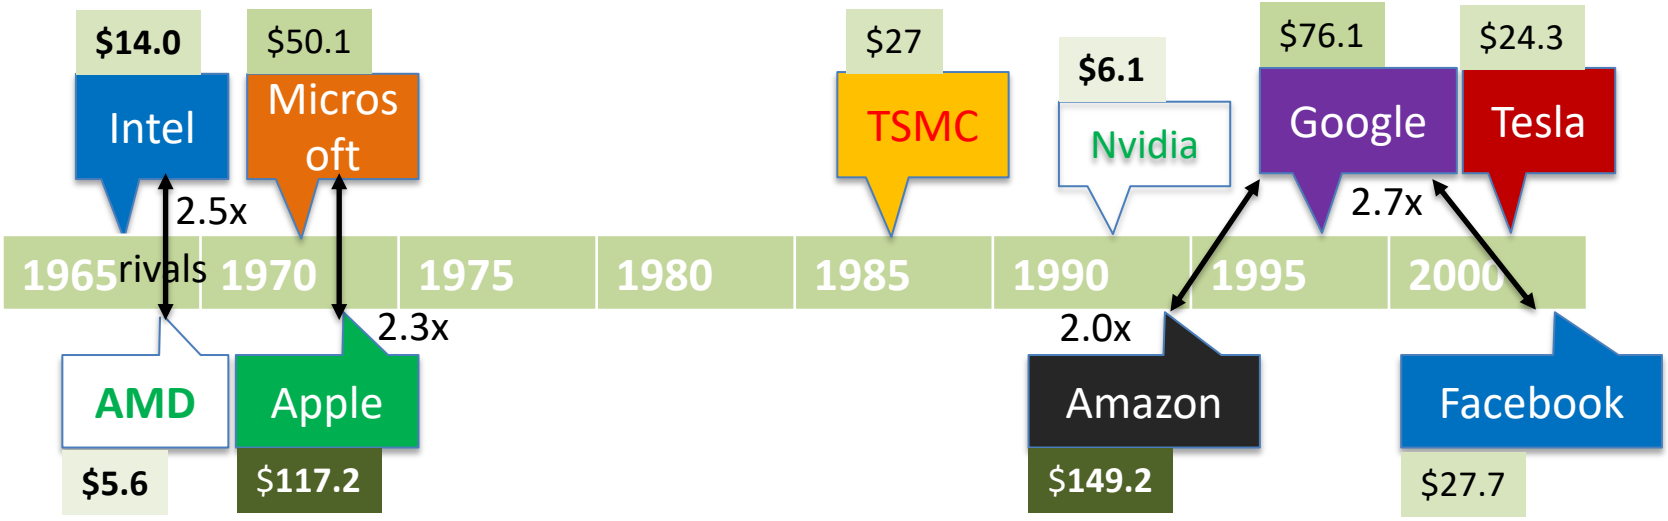

- Other *Industrials*
  - **GM** $43.1 → 2x Tesla
  - **Ford** $37.2
  - **IBM** $16.7 → ~Intel
  - QCOM $9.5
  - TI $4.2
  - NXPI $3.3

- Other *Services*
  - **Netflix** $8.0
  - Visa $7.9
  - **PayPal** $7.4

Top 5 MPU Companies

Leading MPU Suppliers ($B)

2021 Rank	Company	Headquarters	2020	2021	21/20 % Chg	2021 % Marketshare
1	Intel	U.S.	50.6	52.3	3%	50.9%
2	Apple*	U.S.	10.5	13.4	27%	13.0%
3	Qualcomm	U.S.	7.4	9.4	26%	9.1%
4	AMD	U.S.	5.9	9.2	56%	8.9%
5	MediaTek	Taiwan	2.7	4.1	51%	4.0%

*Custom designs for Apple's products that are made by IC foundries.

Source: Company reports, IC Insights

Top 10 IC Designers

Table 1: Global Top Ten IC Design Company Revenue Ranking, 1Q22 (Unit: US$1 Million)

1Q22 Rank	1Q21 Rank	Company	1Q22 Revenue	1Q21 Revenue	YoY
1	1	Qualcomm	9,548	6,281	52%
2	2	NVIDIA	7,904	5,173	53%
3	3	Broadcom	6,110	4,849	26%
4	5	AMD	5,887	3,445	71%
5	4	MediaTek	5,007	3,805	32%
6	9	Marvell	1,412	821	72%
7	6	Novatek	1,281	929	38%
8	8	Realtek	1,044	822	27%
9	-	Will Semiconductor	744	815	-9%
10	-	Cirrus Logic	490	294	67%
	7	Xilinx	-	851	-
-	10	Dialog	-	366	-
Total Revenue			39,427	27,342	44%

Notes

1. This top ten ranking only accounts for companies ahead of public financial reporting.
2. Qualcomm revenue only includes QCT; NVDIA excludes OEM/IP revenue; Broadcom revenue only includes semiconductors; Will Semiconductor revenue only includes semiconductor design and sales.
3. NT$:US$ exchange rate: 1Q22 - 28.50:1; 1Q21 - 28.39:1
4. RMB:US$ exchange rate: 1Q22 - 6.336:1; 1Q21 - 6.483:1

Source: TrendForce, Jun. 2022

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Revenue Growth

Semi Wiki

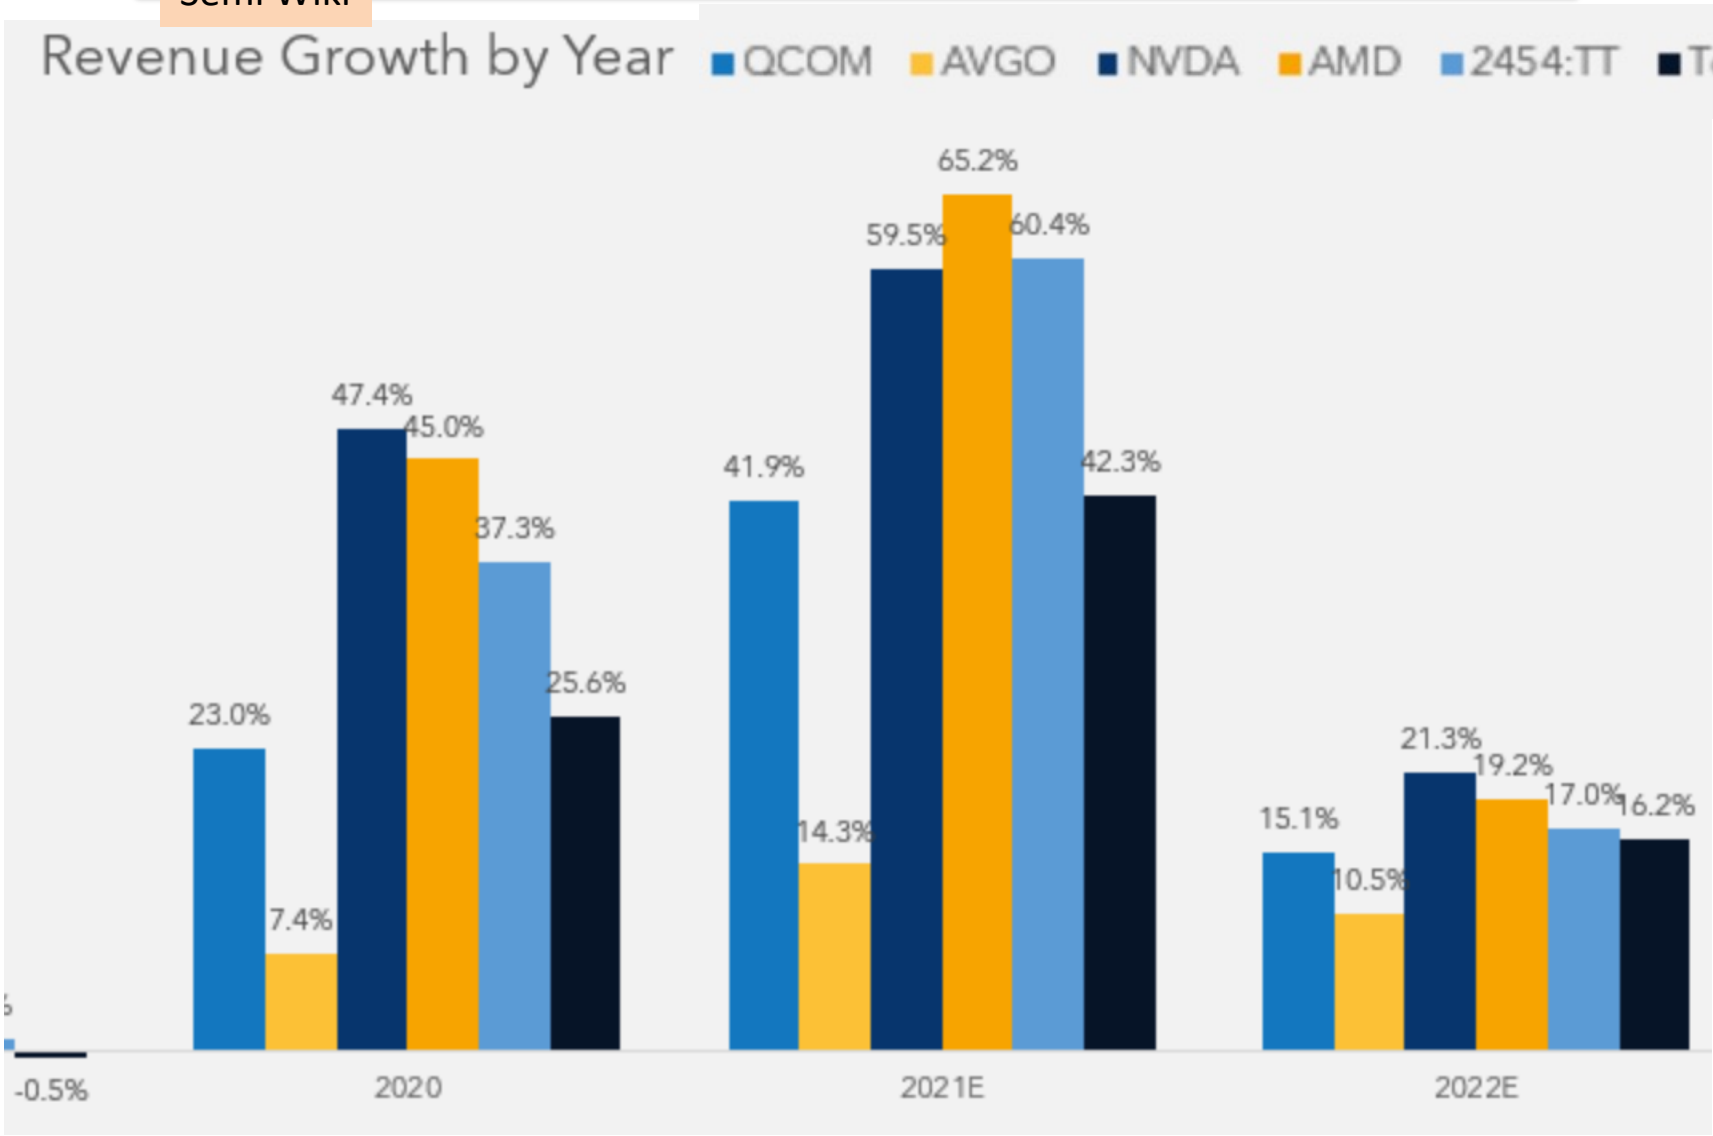

State of the Art

x86

AMD

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

CHM
David Laws

AMD History

AMD 10-Yr Ad 1979

AMD History

In 1983 the annual AMD Christmas party was held in to San Francisco's largest convention space, the Moscone Center. These extravaganzas attracted so much TV and press coverage that their cost of half a million dollars or more was justified on the basis of PR impact.

AMD History

4K SRAM

Advanced Micro Devices has developed, introduced and is volume-producing the first family of 4K static RAM's. They're yours now. Off the shelf. Sorry, Intel.

The Am9130 and Am9140.

Fully Static—no refresh required. No dynamic nodes.
Single 5V power supply.
High speed: Access times to 200 nanoseconds.
Two organizations: Am9140-4Kx1, Am9130-1Kx4.
All input and output logic levels identical to TTL—full 400 mV noise immunity.
Low power dissipation: 250 mW typ.
High output drive: 3.2mA @ 0.4V.

Full military range available: −55C° to +125C° ambient.
DC standby mode. Reduces power dissipation by 80%.
Memory status signal. Indicates when data are valid, allows improved overall performance, and simplifies timing. (And if you don't want to use it, just ignore it. It won't affect conventional memory operation.)
MIL-STD-883, of course.

READY

Advanced MOS/LSI

NMOS

CSUN CALIFORNIA STATE UNIVERSITY NORTHRIDGE

CHM

COMP222 David Laws

AMD History

At one-point AMD was selling more units of its improved version of the Intel 8080 microprocessor than the original supplier.

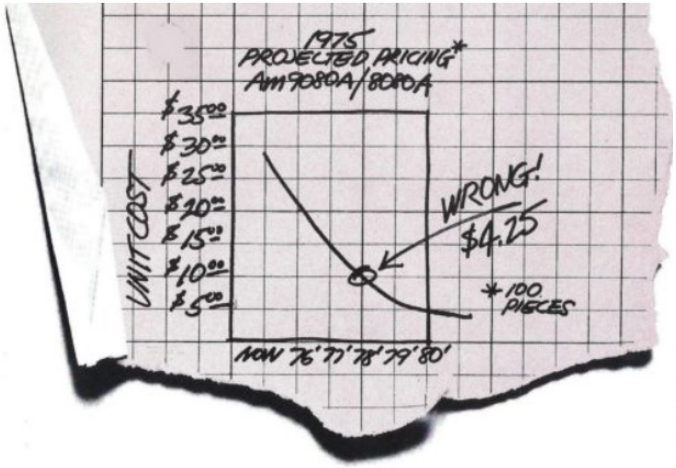

Three years ago we predicted our 8080A, the Am9080A, would be $10 each in 1978.

It's not. It's $4.25.*

And we're talking about the best 8080A you can buy. The one with the best power, speed and reliability records in the business. (The only one that's MIL-STD-883-for-free.)

And we didn't stop there. The Am9080A has an entire family of equally terrific, equally MIL-STD-883-for-free support and peripheral circuits to ith it.

Am9080A's for $4.25.

How do you like that? The best part in the business at the best price in town.

8080A's FOR $4.25.

Advanced Micro Devices

Multiple technologies. One product: excellence.
901 Thompson Place, Sunnyvale, California 94086
Telephone (408) 732-2400

AMD History

Z8000

Lacking access to the next generation Intel 8086 architecture, in 1978 AMD signed an agreement with Zilog to second source the Z8000 processor.

powerful instructions. It can even accommodate more data types. And the AmZ8000 has a lot higher throughput using standard NMOS than the 8086 using HMOS.

To demonstrate the capability of the AmZ8000, we developed a fully assembled and tested Evaluation Board with a memory, an I/O and a monitor. Ask for it by name: AMC 96/4016. You can also get a full ASCII keyboard/display and an assembler.

One last thing: we know it hurts to drop Intel for somebody else. We went through it ourselves. But it's going to hurt a lot more next year. By then, your competitors could be so far ahead of you, you might never catch up.

Call Advanced Micro Devices. We'll send you all the latest information on the AmZ8000, System 8/8 and the AMC 96/4016. Or, we'll line you up for our next 4-day seminar.

When you've looked at all the facts, one fact is going to come through loud and clear: The AmZ8000 is better.

Advanced Micro Devices

901 Thompson Place, Sunnyvale, CA 94086

CSUN CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

CHM

David Laws

AMD History

Z8000 → 8086

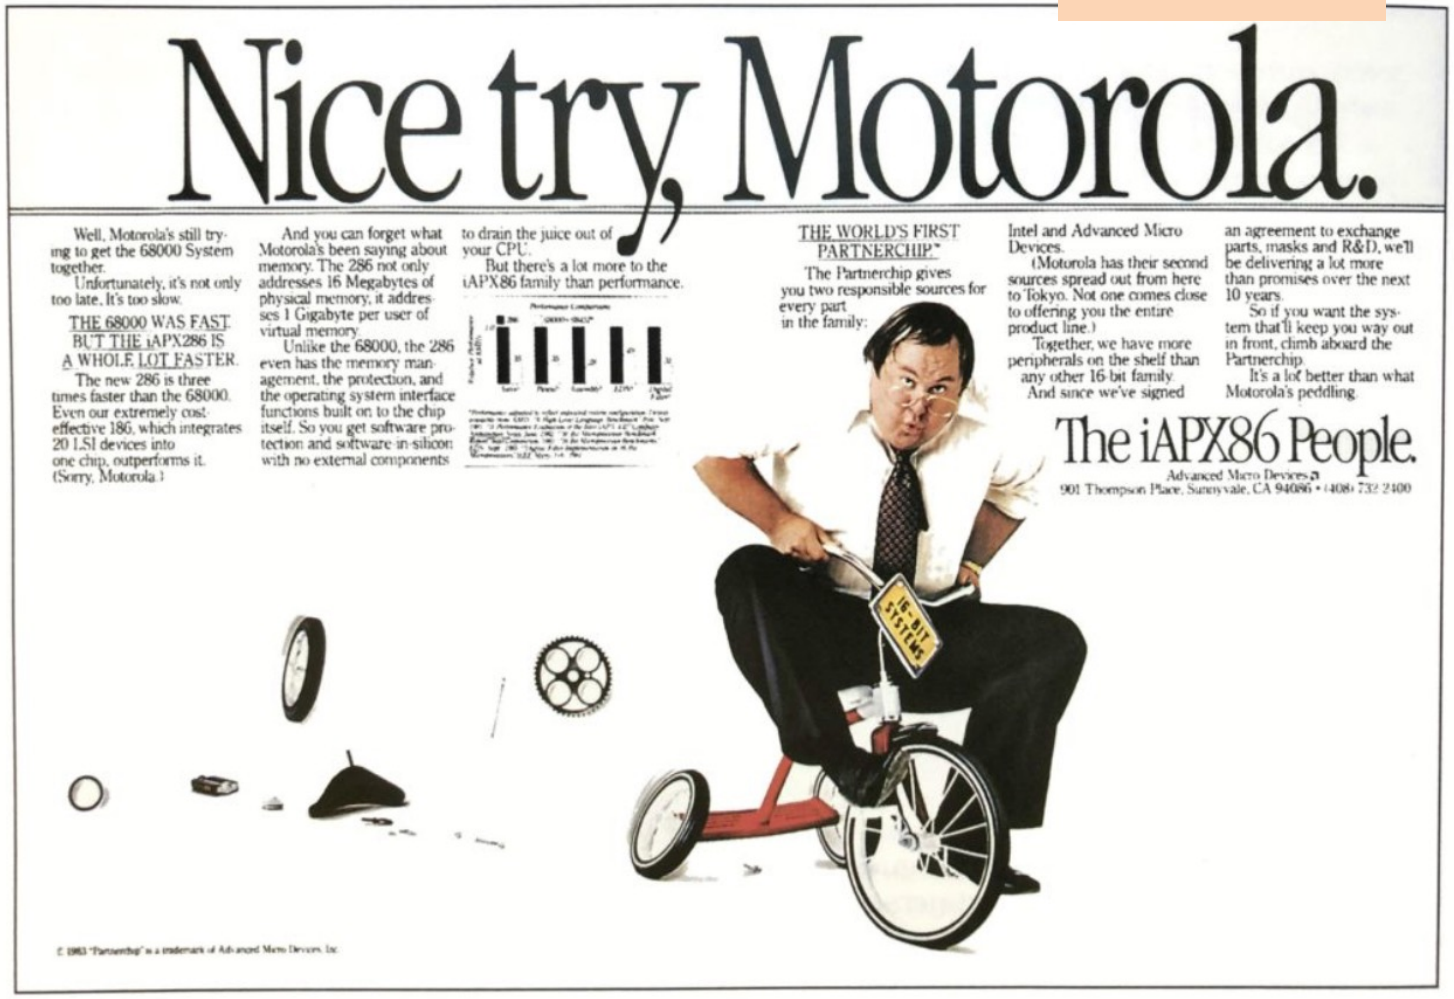

Nice try, Motorola.

Well, Motorola's still trying to get the 68000 System together.

Unfortunately, it's not only too late. It's too slow.

THE 68000 WAS FAST. BUT THE iAPX286 IS A WHOLE LOT FASTER.

The new 286 is three times faster than the 68000. Even our extremely cost-effective 186, which integrates 20 LSI devices into one chip, outperforms it. (Sorry, Motorola.)

And you can forget what Motorola's been saying about memory. The 286 not only addresses 16 Megabytes of physical memory, it addresses 1 Gigabyte per user of virtual memory.

Unlike the 68000, the 286 even has the memory management, the protection, and the operating system interface functions built on to the chip itself. So you get software protection and software-in-silicon with no external components to drain the juice out of your CPU.

But there's a lot more to the iAPX86 family than performance.

THE WORLD'S FIRST PARTNERCHIP™

The Partnerchip gives you two responsible sources for every part in the family:

Intel and Advanced Micro Devices.

(Motorola has their second sources spread out from here to Tokyo. Not one comes close to offering you the entire product line.)

Together, we have more peripherals on the shelf than any other 16-bit family.

And since we've signed an agreement to exchange parts, masks and R&D, we'll be delivering a lot more than promises over the next 10 years.

So if you want the system that'll keep you way out in front, climb aboard the Partnerchip.

It's a lot better than what Motorola's peddling.

The iAPX86 People.

Advanced Micro Devices

901 Thompson Place, Sunnyvale, CA 94086 • (408) 732-2400

16-BIT SYSTEMS

By 1983, AMD had dropped the Z8000 architecture in favor of an agreement with Intel to produce the 8086/8088 family for the IBM PC. The competitor for new designs was now the Motorola 68000.

AMD History

ACKNOWLEDGEMENTS

David Laws

In full disclosure: Beginning in 1975, the author spent 12 of the most rewarding years of his professional life at AMD in roles from marketing manager to vice president of business development. Thank you, Jerry Sanders for creating an extraordinary place to work and grow. To Steve Zelencik for his many years of mentoring and support and also for donating his collection of AMD advertisements, several of which are used in this article, and other documents to the Computer History Museum (CHM). And to John Springer, keeper of the FLAMES website and archive for Former Loyal Advanced Micro Employees.

AMD vs Intel: CPU Families

Market Segment	AMD	Intel
Desktop	Ryzen 3/5/7K	Core i5/i7/i9 (13th 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.

Quora

Brett Bergan · Follow
Building PC's for 25 years ·

AMD Chiplets

Improved Yield and Lower Costs

To me the single stroke of genius came in the form of using 7nm CCX chiplets (of 74mm² in size) on the same package as the 14nm I/O die for huge increases in yield per wafer.

- Perfect 8-core CCX chips went to the 3700X, 3800X, and 3950X
- Flawed chips (with 6 good cores) ended up in the 3600, 3600X, and 3900X
- Chiplets with multiple failure points ended up as quad cores in 3100 and 3300X

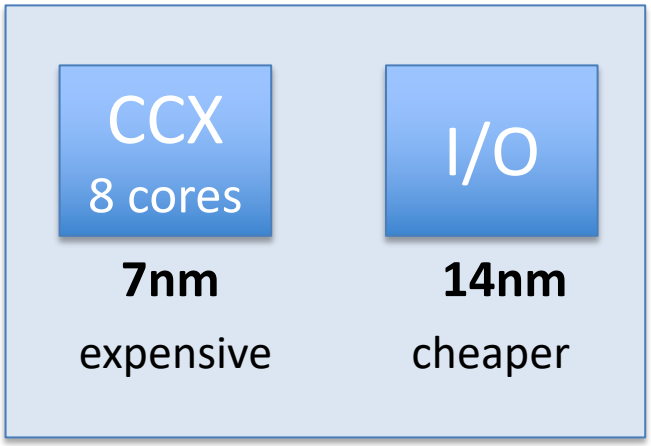

MCM socket

AMD vs Intel

Quora

Drazen Zoric · Follow
Lives in Cork, Ireland · 5h

Overall x86 CPU Share (ALL CPUs)

Overall x86 CPU Share Includes IoT and SoC	2022 Q1 Current Quarter	2021 Q4 Prior Quarter	2021 Q1 Year Ago Quarter
	Share	Share	Share
Intel	72.3%	74.4%	79.3%
AMD	27.7%	25.6%	20.7%
VIA	0.0%	0.0%	0.0%
Total	100.0%	100.0%	100%

AMD vs Intel

Quora

Drazen Zoric · Follow
Lives in Cork, Ireland · 5h

Server CPU Share excluding IoT

Server CPU Share	2022 Q1 Current Quarter	2021 Q4 Prior Quarter	2021 Q1 Year Ago Quarter	Share Change (points)	Share Change (points)
	Share	Share	Share	Quarter	Year
Intel	88.4%	89.3%	91.1%	- 0.9	- 2.7
AMD	11.6%	10.7%	8.9%	+ 0.9	+ 2.7
Total	100.0%	100.0%	100.0%		

Desktop CPU Share excluding IoT

Desktop PC CPU Share	2022 Q1 Current Quarter	2021 Q4 Prior Quarter	2021 Q1 Year Ago Quarter	Share Change (points)	Share Change (points)
	Share	Share	Share	Quarter	Year
Intel	81.7%	83.8%	80.6%	- 2.1	+ 1.1
AMD	18.3%	16.2%	19.3%	+ 2.1	- 1.0
VIA	0.0%	0.0%	0.1%	+ 0.0	- 0.0
Total	100.0%	100.0%	100.0%		

AMD vs Intel

Quora

Drazen Zoric · Follow
Lives in Cork, Ireland · 5h

Mobile CPU Share excluding IoT

Mobile CPU Share	2022 Q1	2021 Q4	2021 Q1	Share	Share
	Current Quarter	Prior Quarter	Year Ago Quarter	Change (points)	Change (points)
	Share	Share	Share	Quarter	Year
Intel	77.5%	78.4%	82.0%	- 0.9	- 4.4
AMD	22.5%	21.6%	18.0%	+ 0.9	+ 4.4
Total	100.0%	100.0%	100.0%		

Yeah, Intel lost 2 - 7% but still sells 4 - 8 times more CPUs. AMD will never be able to close this gap. Things in Intel changed with 12th gen when they have again fastest CPUs. AMD Zen4 might take a lead but in few months 13th gen is out which will be better.
I saw ridiculous AMD Zen4 pricing, insane $800 for 7950X. If Intel lowers 13900 it will regain share.

Next year when Intel switches to HA EUV, 14th gen, it will have also better laptop CPUs which will be lower power.

Question is what is going on with Sapphire Rapids. It has 12 respins and still some 500 bugs. AMD already released Epyc Genoa with 96 cores what will threaten Intel in server and already did in supercomputer segments.

AMD News

Industry News:

- AMD announced it has completed Class B qualification for the company's first space-grade Versal adaptive SoCs, allowing the devices to begin shipping in early 2023. The announcement saw coverage in The Register, SiliconANGLE, HPCWire and others.
- AMD showcased its momentum in high performance computing and latest wins on the Top500 and Green500 lists, as part of SC22. Tom's Hardware, The Next Platform, HotHardware and many others covered the news.
- PC Gamer, Tom's Hardware and Digital Trends shared articles following various events at the "together we advance_gaming" event in Las Vegas.

Zen 4 Benchmark

- Multi-threaded - here Ryzen leads cause it has more cores

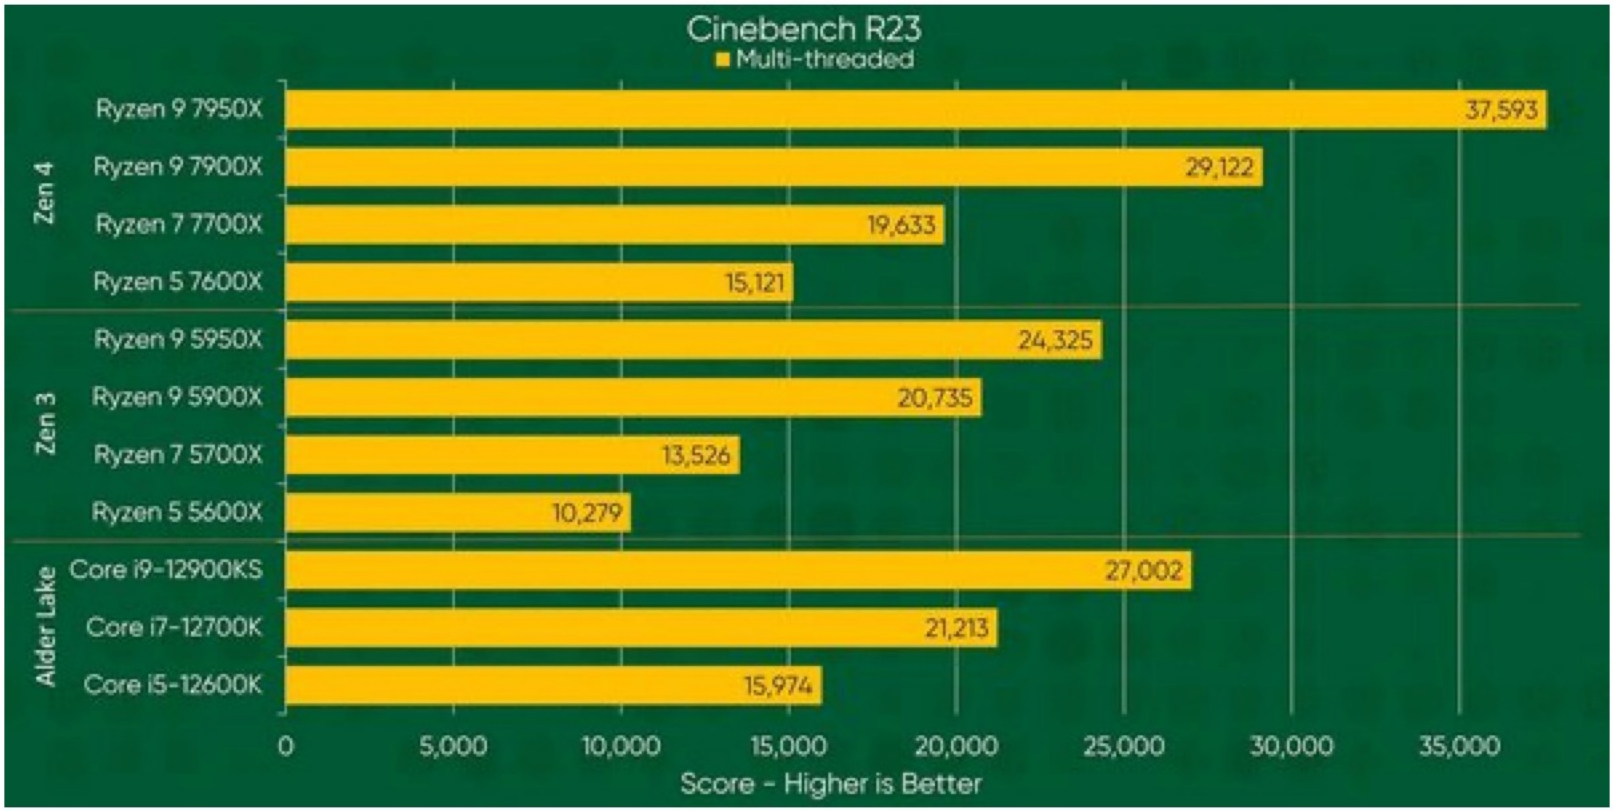

➢ **Zen 4** is 50% better than Zen 3

AMD CES 2023

Jan 2023

AMD CES 2023

Jan 2023

AMD's new Ryzen 7000 mobile processors include a massive 16-core chip

Story by Monica Chin • Yesterday 7:30 PM

Comments

This may seem like an underdog compared to Intel's top 13th-Gen chips, which have 24 cores — Intel has claimed that its Core i9-13980HX is the "world's fastest mobile processor". This chip, however, only has eight performance cores and 16 efficiency cores, while all 16 of the Ryzen's will be going full-speed ahead. This is essentially two full eight-core chips stuffed into one.

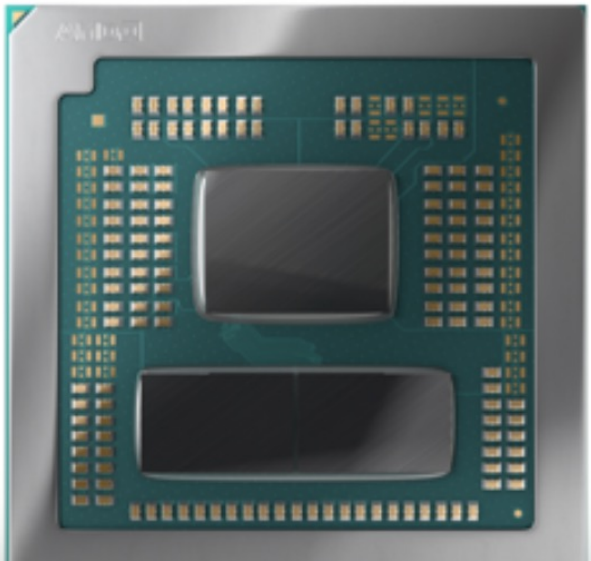

AMD CES 2023

Jan 2023

Zen4

5nm

16 P

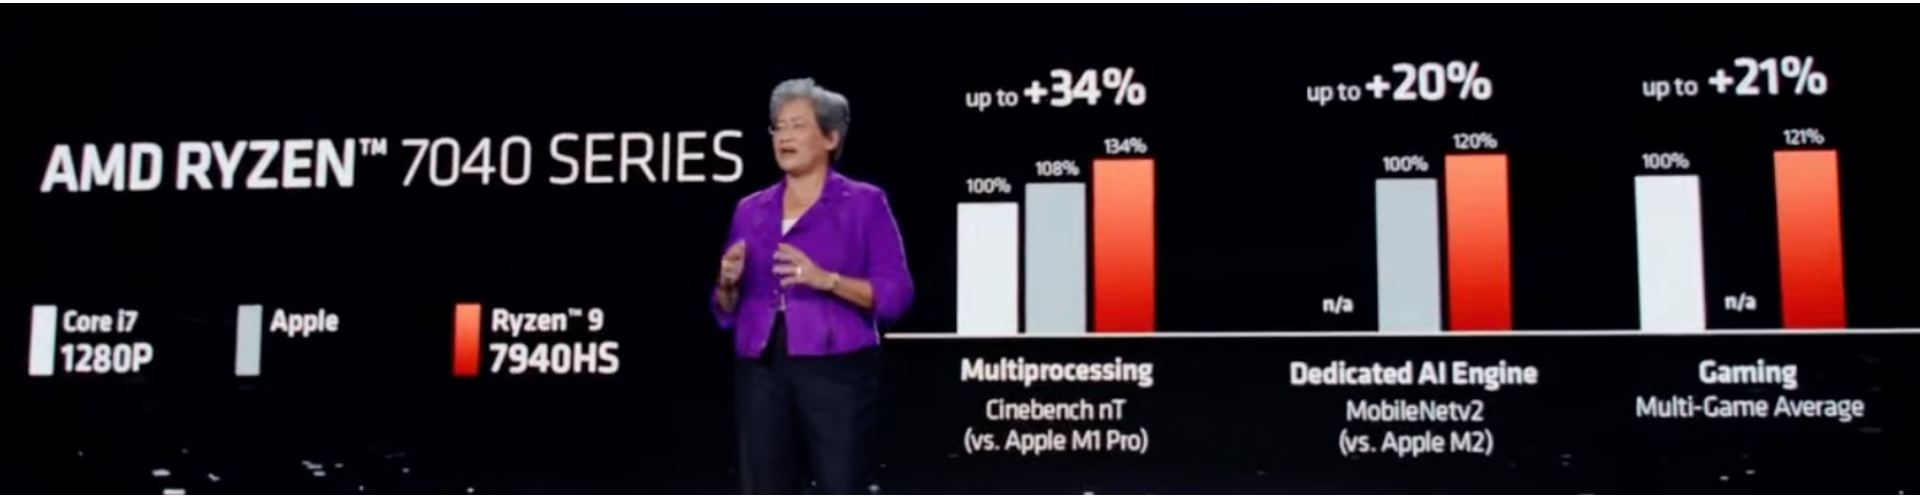

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

AMD CES 2023

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jan 2023

Zen4

5nm

16 P

RDNA3

6nm

32 CU

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

AMD CES 2023

Jan 2023

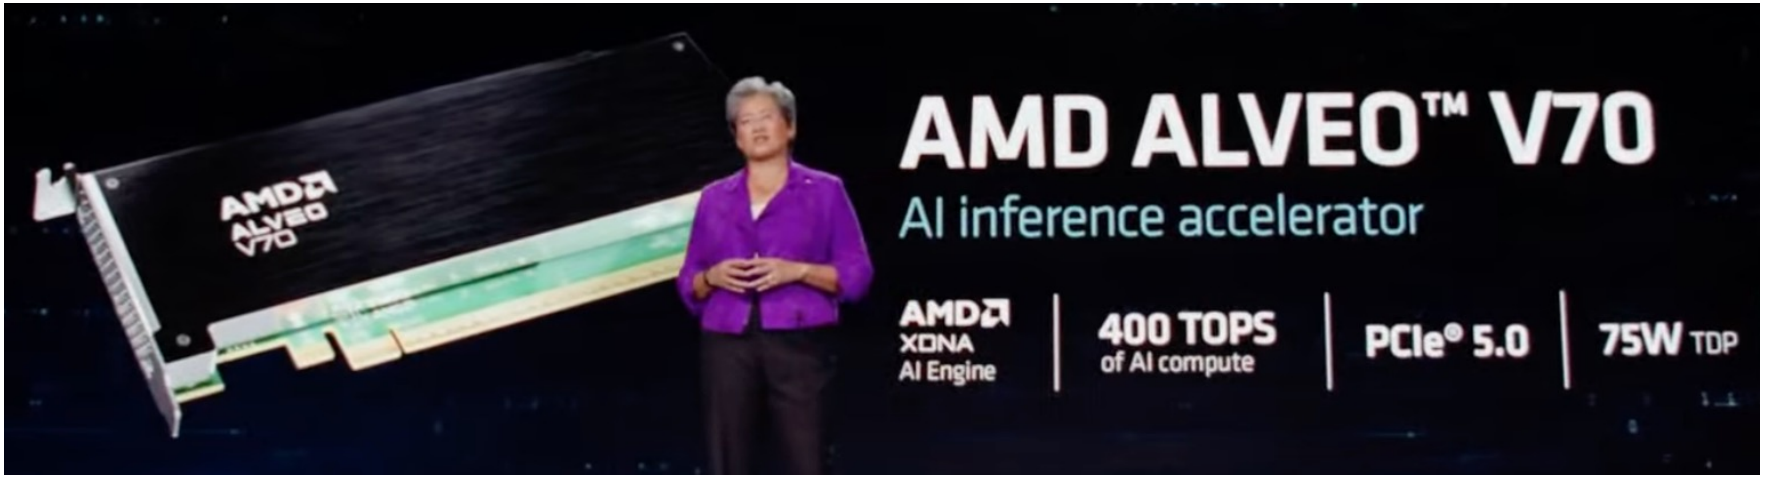

AMD CES 2023

Jan 2023

- **HPC** and **AI** System in a Module

Used in latest top **supercomputer** (Frontier)

AMD CES 2023

Jan 2023

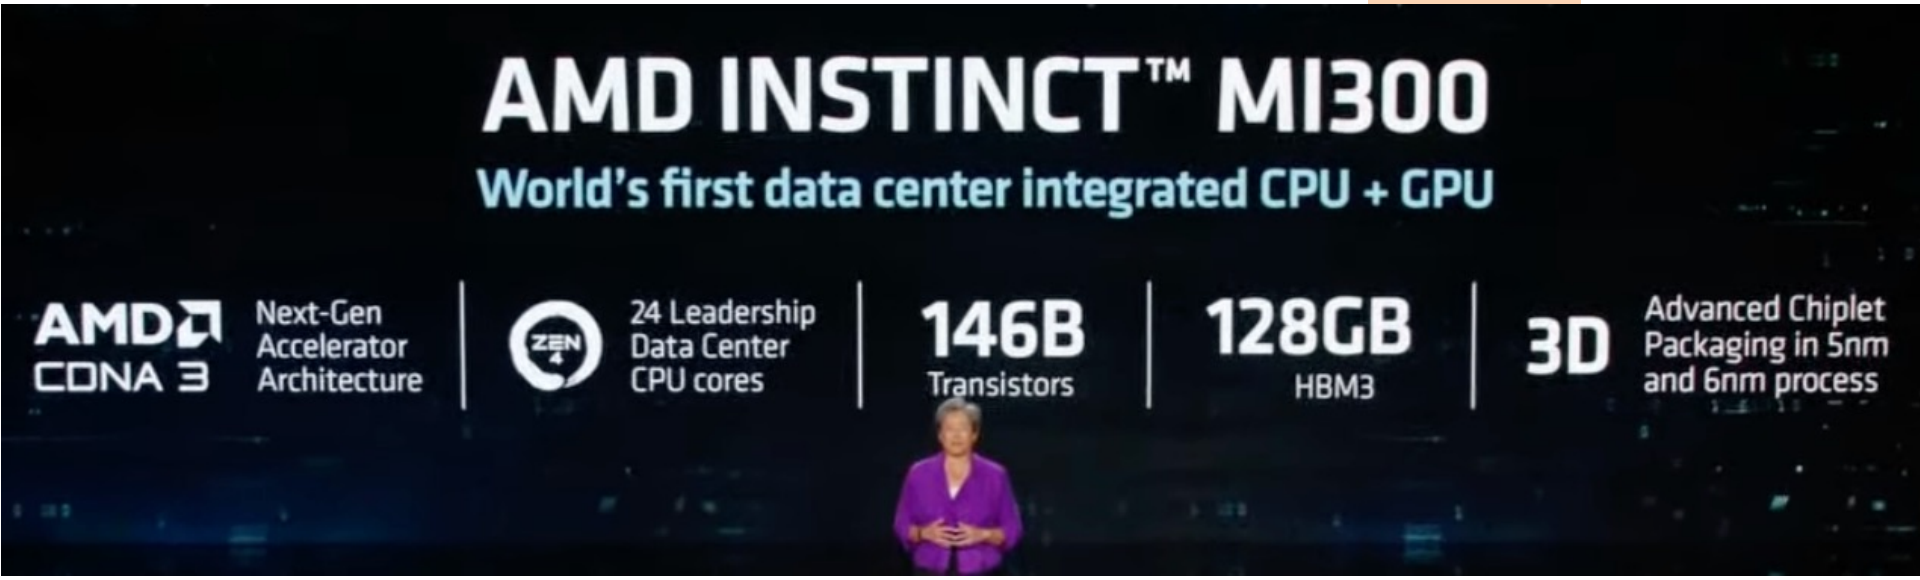

State of the Art

x86

Intel

Microprocessor Leaders: AMD vs Intel

COMP222

Intel Core i Family

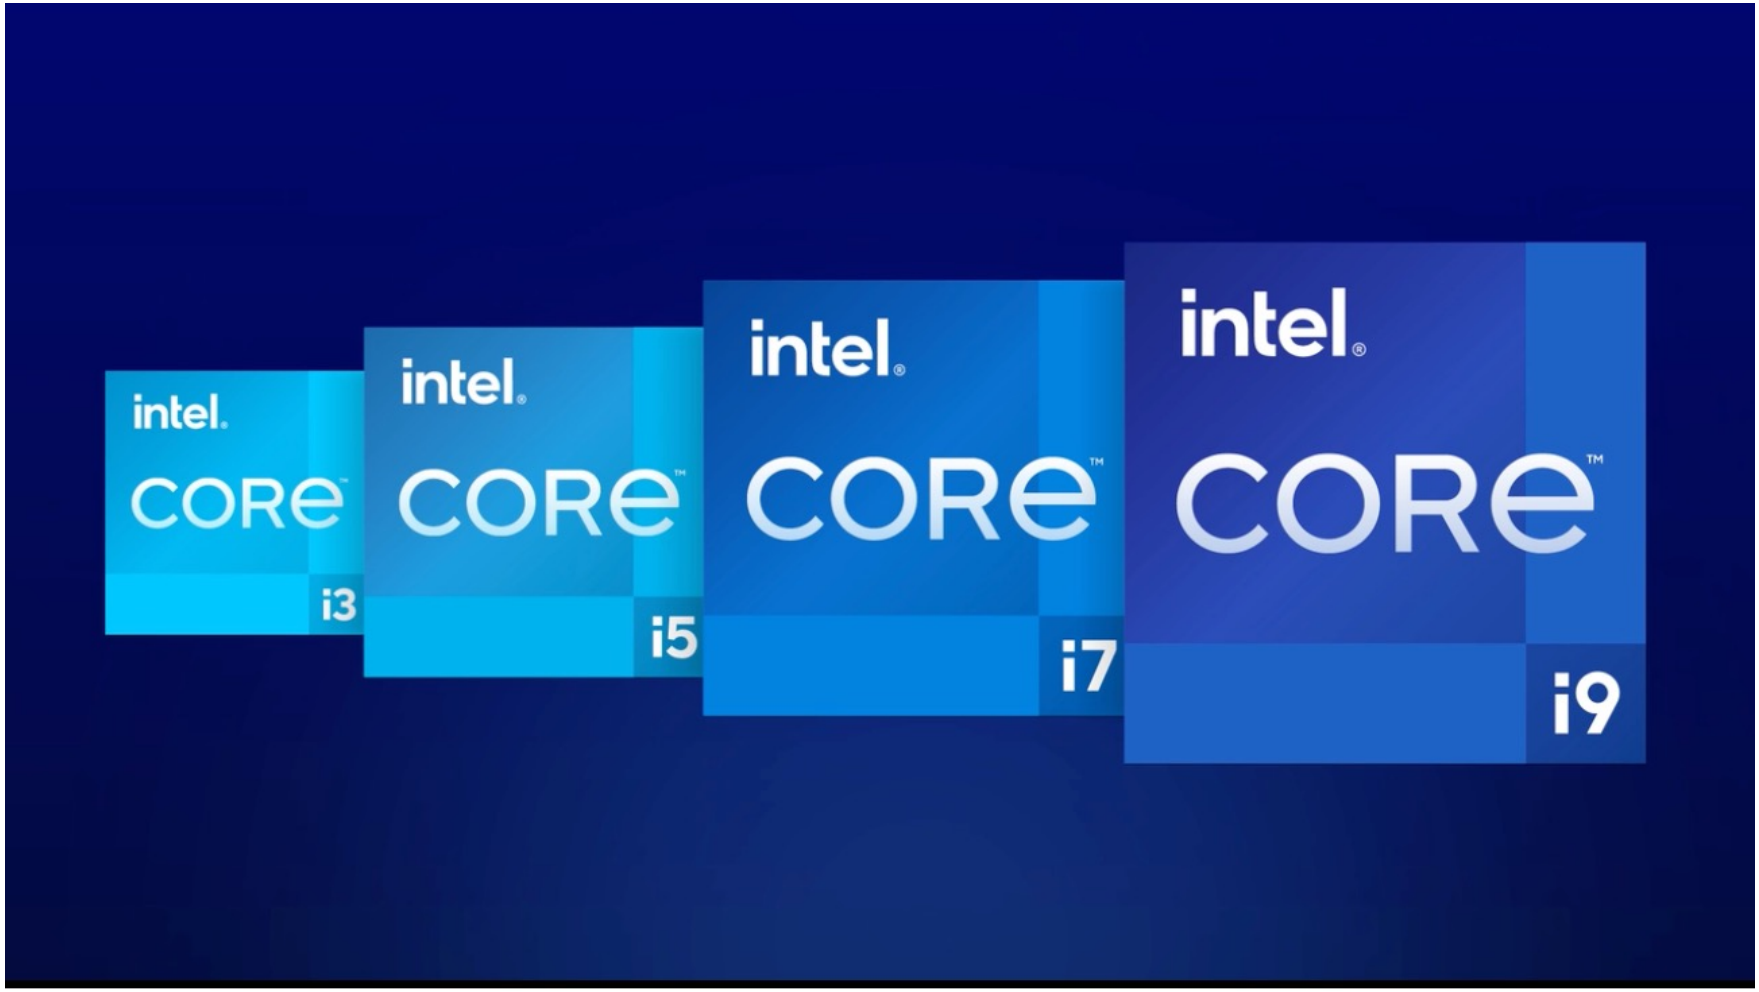

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Intel ‘s New 13th Gen

Oct 2022

13th Gen Intel® Core™ Unlock

Processor Number	Processor Cores (P+E)	Processor Threads	Intel® Smart Cache (L3)	Total L2 Cache	P-core Max Turbo Frequency (GHz)	E-core Max Turbo Frequency (GHz)	P-core Base Frequency (GHz)	E-core Base Frequency (GHz)
i9-13900K	24 (8+16)	32	36MB	32MB	Up to 5.8	Up to 4.3	3.0	2.2
i9-13900KF	24 (8+16)	32	36MB	32MB	Up to 5.8	Up to 4.3	3.0	2.2
i7-13700K	16 (8+8)	24	30MB	24MB	Up to 5.4	Up to 4.2	3.4	2.5
i7-13700KF	16 (8+8)	24	30MB	24MB	Up to 5.4	Up to 4.2	3.4	2.5
i5-13600K	14 (6+8)	20	24MB	20MB	Up to 5.1	Up to 3.9	3.5	2.6
i5-13600KF	14 (6+8)	20	24MB	20MB	Up to 5.1	Up to 3.9	3.5	2.6

Intel 's New 13th Gen

Oct 2022

Desktop Processors

Processor Graphics	Total CPU PCIe Lanes	Max Memory Speed (MT/S)	Memory Capacity	Processor Base Power (W)	Max Turbo Power (W)	RCP (USD)
Intel® UHD Graphics 770	20	DDR5 5600 DDR4 3200	128GB	125	253	$589
n/a	20	DDR5 5600 DDR4 3200	128GB	125	253	$564
Intel® UHD Graphics 770	20	DDR5 5600 DDR4 3200	128GB	125	253	$409
n/a	20	DDR5 5600 DDR4 3200	128GB	125	253	$384
Intel® UHD Graphics 770	20	DDR5 5600 DDR4 3200	128GB	125	181	$319
n/a	20	DDR5 5600 DDR4 3200	128GB	125	181	$294

Intel Xeon 3rd/4th Gen

Quora

Nov 2022

4th Gen Intel® Xeon® Scalable processors

AMX

4th Gen Intel® Xeon® Scalable processors will include a built-in accelerator dedicated to the matrix multiplication at the heart of deep learning workloads.

Intel® Advanced Matrix Extensions (Intel® AMX) combines a new instruction set that turns large matrices into a single operation with two-dimensional register files that store larger chunks of data for each core.

AI acceleration on 3rd Gen Intel® Xeon® Scalable processors

- Up to 1.74x higher INT8 batch inference throughput on BERT-Lasge SQuAD with Intel® DL Boost on 3rd Gen Intel® Xeon® Scalable processors vs. prior generation.[1]
- Up to 1.59x higher INT8 real-time inference throughput with Intel® DL Boost on 3rd Gen Intel Xeon Scalable processors vs. prior generation.[2]
- Up to 4.5x more images per second at INT8[3] and up to 6x more images per second at BF16[4] object detection (SSD-ResNet-34) using Intel® AMX on upcoming 4th Gen Intel®

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Intel Xeon 4th Gen

Jan 2023

Intel® Xeon® Scalable Processors	Intel® Xeon® Scalable processor family delivers unparalleled scale and performance for compute, storage, network, security.
Intel® Xeon® Processors	Built for data centers and workstations to handle the heavy processing demands of cloud, big data, modeling, AI, and more.
Intel® Core™ Processors	Intel's highest-performance CPUs for laptops and desktops, delivering advanced responsiveness, connectivity and graphics.
Intel® Processor	Intel's entry level CPUs provide the performance that you need with the affordability you want to connect, learn, and play anywhere.
Intel Atom® Processor	These small, powerful CPUs are ideal for mobile and IoT devices as well as high-density, low-energy data center applications.
Processors for IoT and Embedded Applications	For swift deployment of edge applications, see Intel's portfolio of edge-ready compute and connectivity technologies.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Intel Xeon 4th Gen

Jan 2023

Intel® Xeon® Scalable Processors

4th Gen Intel® Xeon® Scalable processors feature built-in accelerators and advanced security technologies for the most in-demand workload requirements — all while offering the greatest cloud choice and application portability.

Intel® Xeon® Platinum Processors

Advanced 2, 4 & 8 socket performance, designed for the most demanding workloads & services from the edge to cloud.

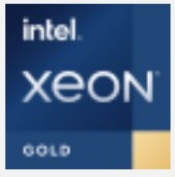

Intel® Xeon® Gold Processors

Up to 4 socket scalable performance, advanced reliability, and advanced security solutions.

Intel® Xeon® Silver Processors

Performance and power efficiency for entry compute, network and storage.

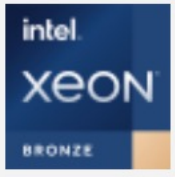

Intel® Xeon® Bronze Processors

Reliability and serviceability for small business and storage server solutions.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Intel Xeon 4th Gen

Jan 2023

Intel® Xeon® Platinum Processor

4th Gen Intel® Xeon® Scalable processors feature built-in accelerators and advanced security technologies designed over decades of innovation for the most in-demand workload requirements—all while offering the greatest cloud choice and application portability.

Built-in Intel Accelerator Engines for Performance and Security

AMX

Intel® Advanced Matrix Extensions

Intel® AMX is Intel's next-generation advancement for deep-learning training and inference on 4th Gen Intel® Xeon® Scalable processors.

Intel® Advanced Vector Extensions 512

AVX-512

Intel® AVX-512 is purpose-built to accelerate performance for the most demanding computational workloads in science, business and beyond.

Intel® Software Guard Extensions

Intel® SGX helps protect data in use via unique application-isolation technology.

Learn more

See All Accelerator Engines

Intel® Accelerator Engines are integrated accelerators on Intel® Xeon® Scalable processors that are purpose-built to deliver performance and power efficiency advantages across today's fastest-growing workloads.[1]

Intel Xeon AMX Benchmark

Jan 2023

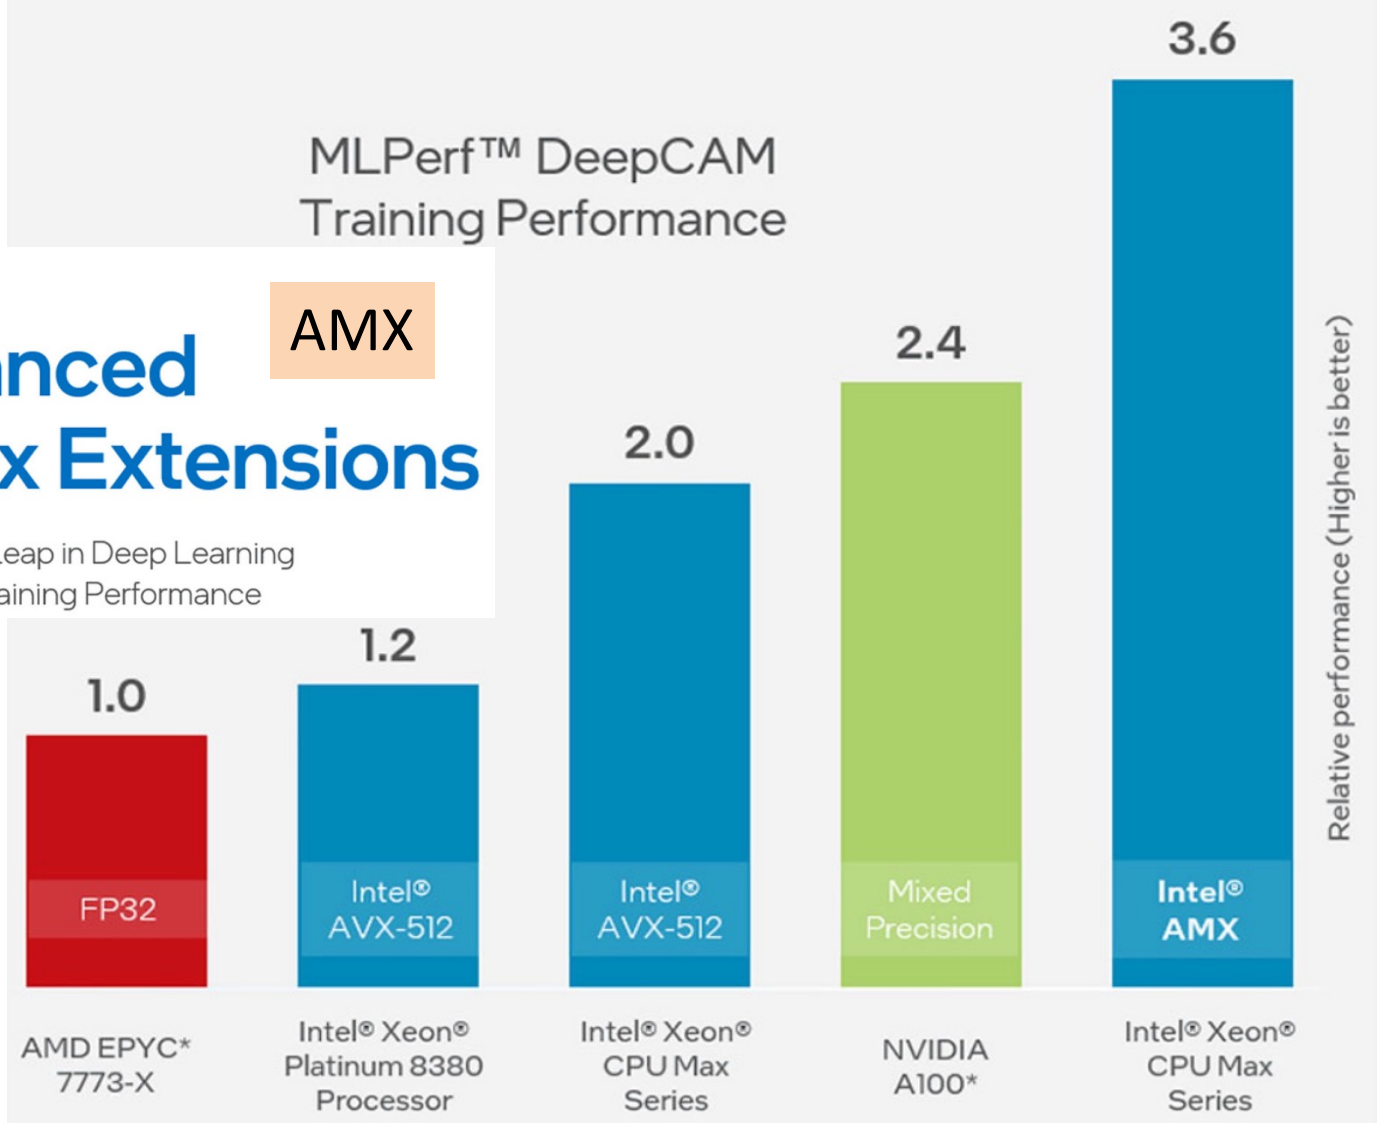

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Intel Xeon oneAPI

DR JEFF SOFTWARE
INDIE APP DEVELOPER

2020-23

Jan 2023

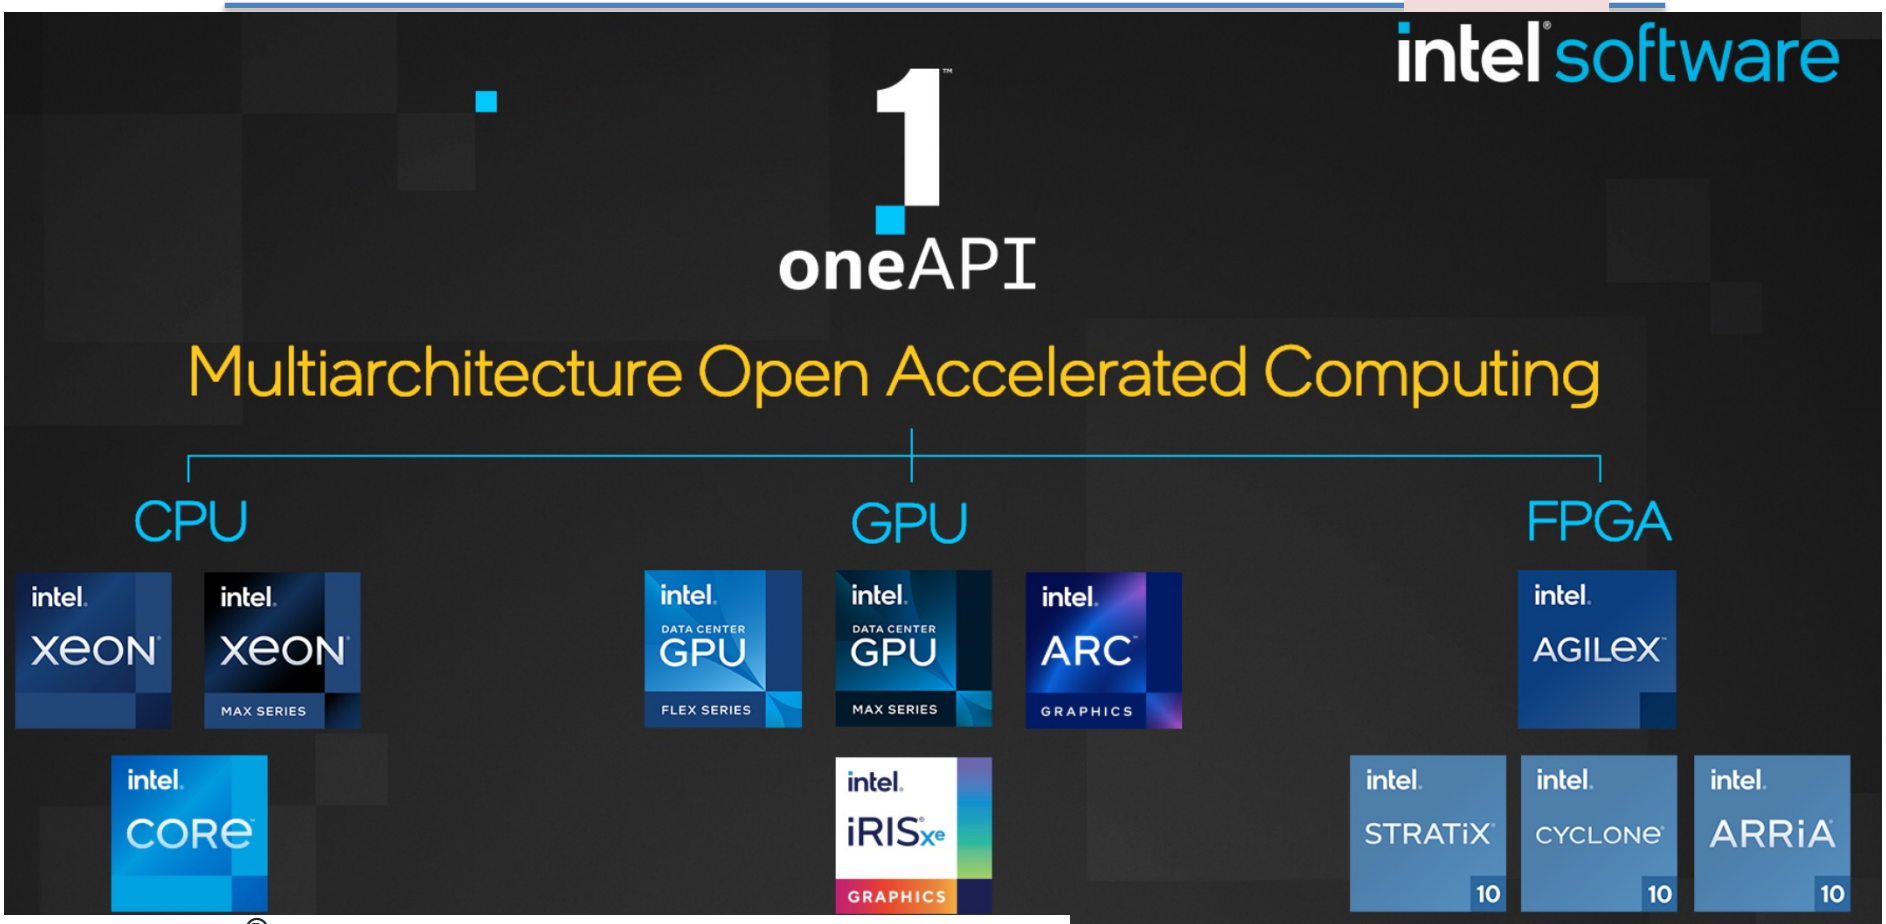

- Activate Intel® AMX support for int8 and bfloat16 data types using oneAPI performance libraries such as oneDNN, oneDAL, and oneCCL.
- Drive orders of magnitude for training and inference into TensorFlow and PyTorch AI frameworks which are powered by oneAPI and already optimized to enable Intel AMX.
- Activate the hardware's innovative features—Intel® X^e Matrix Extensions, vector engine, Intel® X^e Link, data type flexibility, and more—and realize maximum performance using oneAPI and AI Tools.
- Migrate CUDA code to SYCL for easy portability across multiple architectures—the new GPU as well as those from other vendors—with code migration tools to simplify the process.

Intel's Fabs

Intel Fabs

Quora

Can I say Intel has overcome TSMC now, because the 13th CPU is better than the AMD 7000 series?

Jeff Drobman
Works at Dr Jeff Software · Just now ·

NO! Both Intel and AMD are using the same TSMC 5nm process for their high end CPU's, and will move to TSMC 3e next. Intel is still behind TSMC by 1–2 generations, so they offer their inferior fabs to others in a foundry model.

Intel's Foundry Biz (IFS)

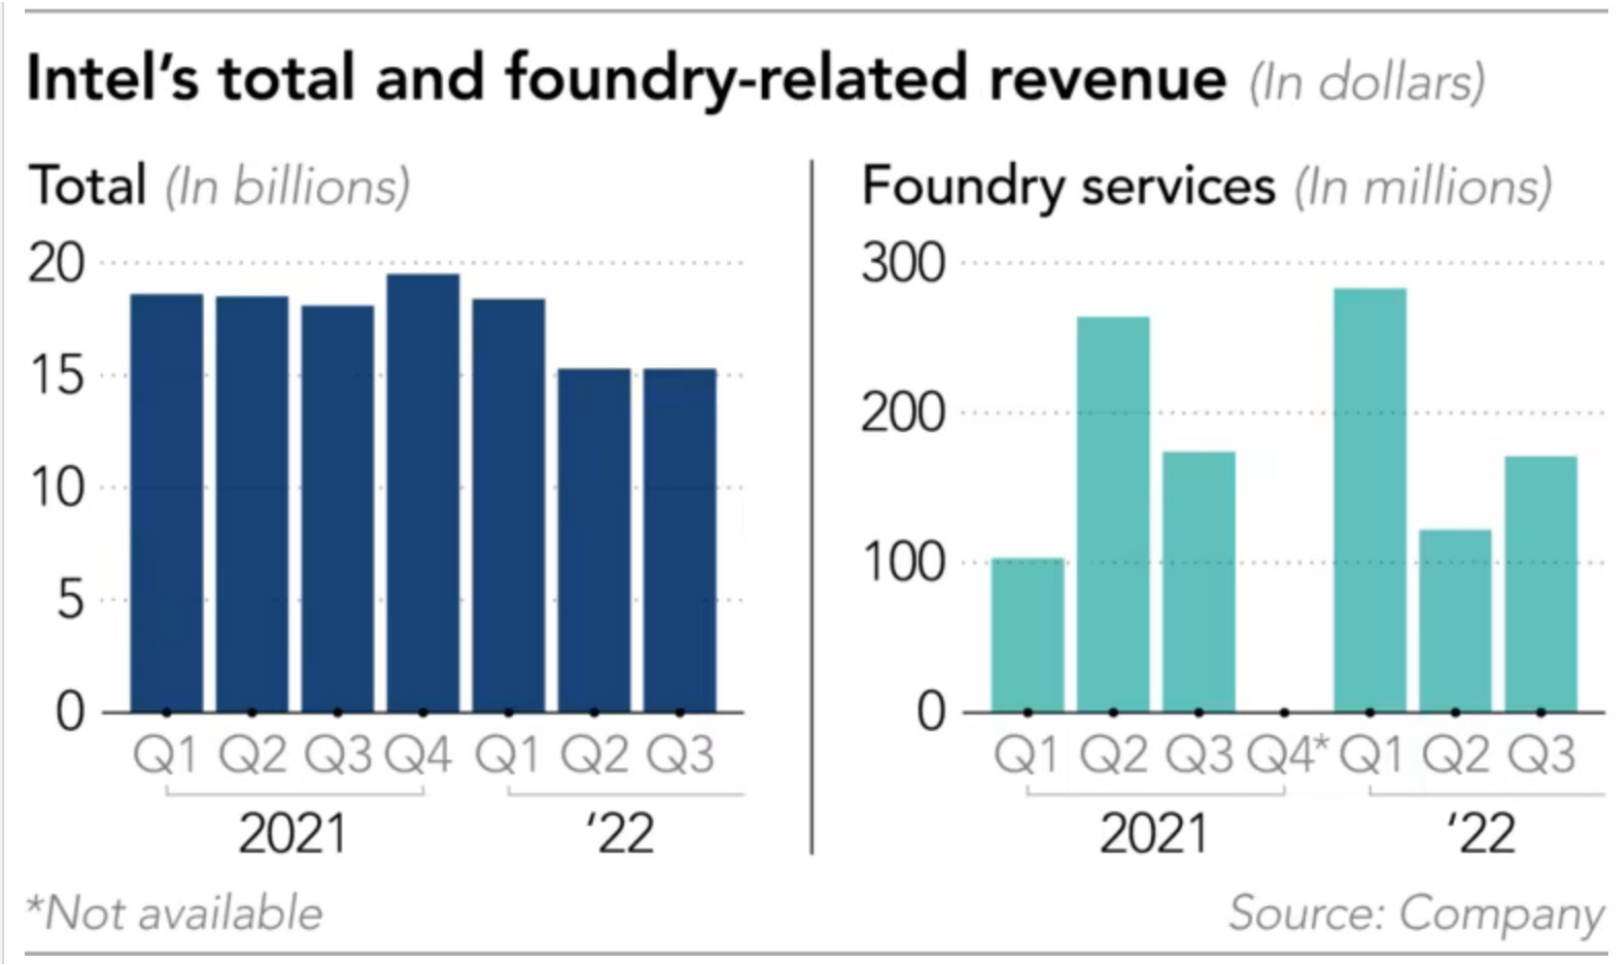

Intel's Foundry Biz (IFS)

The company's spending in this area includes $20 billion for a chip facility in Ohio and 17 billion euros ($16.8 billion) to build a plant in Germany, as well as $3.5 billion to expand its chip packaging facility in New Mexico, a $20 billion investment in Arizona fabs and a 17 billion euro expansion in Ireland. On top of that, Intel acquired Israeli foundry Tower Semiconductor for $5.4 billion in February.

Meanwhile, slowing global demand for chips has weighed on Intel's top line. The company reported a 20% year-over-year drop in third-quarter revenue last week, and lowered its 2022 full-year revenue outlook to between $63 billion and $64 billion, down as much as $4 billion from its previous guidance. Coupled with the heavy spending on its foundry business, Intel is now expecting to end 2022 with a negative $2 billion to $4 billion free cash flow, compared to the negative $1 billion to $2 billion it projected earlier this year.

Intel Revenue

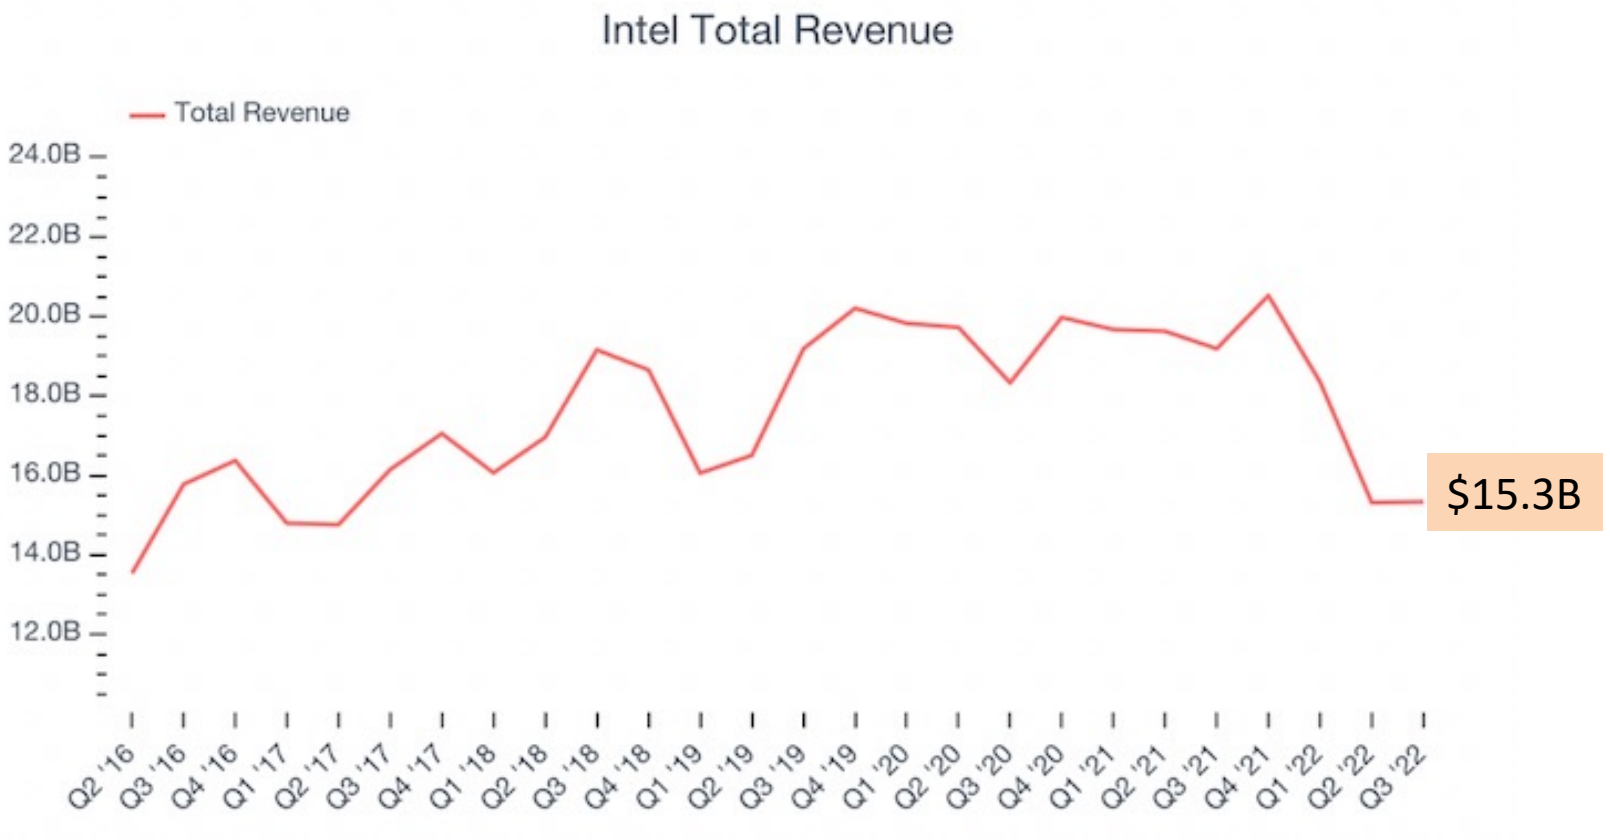

Reports 3Q22

3Q ($B)

- Tesla = $21.5
- TSMC = $20.23
- IBM = $14.1

- Amazon = $127.1
- Apple = $90.1
- Google = $69.1
- Microsoft = $50.1
- Intel = $15.3

- Qualcomm = $11.4
- Nvidia = $5.9B
- AMD = $5.6B

2Q ($B)

- Amazon = $121.2
- Apple = $83.0
- Google = $69.7
- Microsoft = $51.9
- Intel = $15.3
- Nvidia = $6.7
- AMD = $6.6

State of the Art

APPLE

Apple Event

9-7-22

- iPhone **14 Pro**

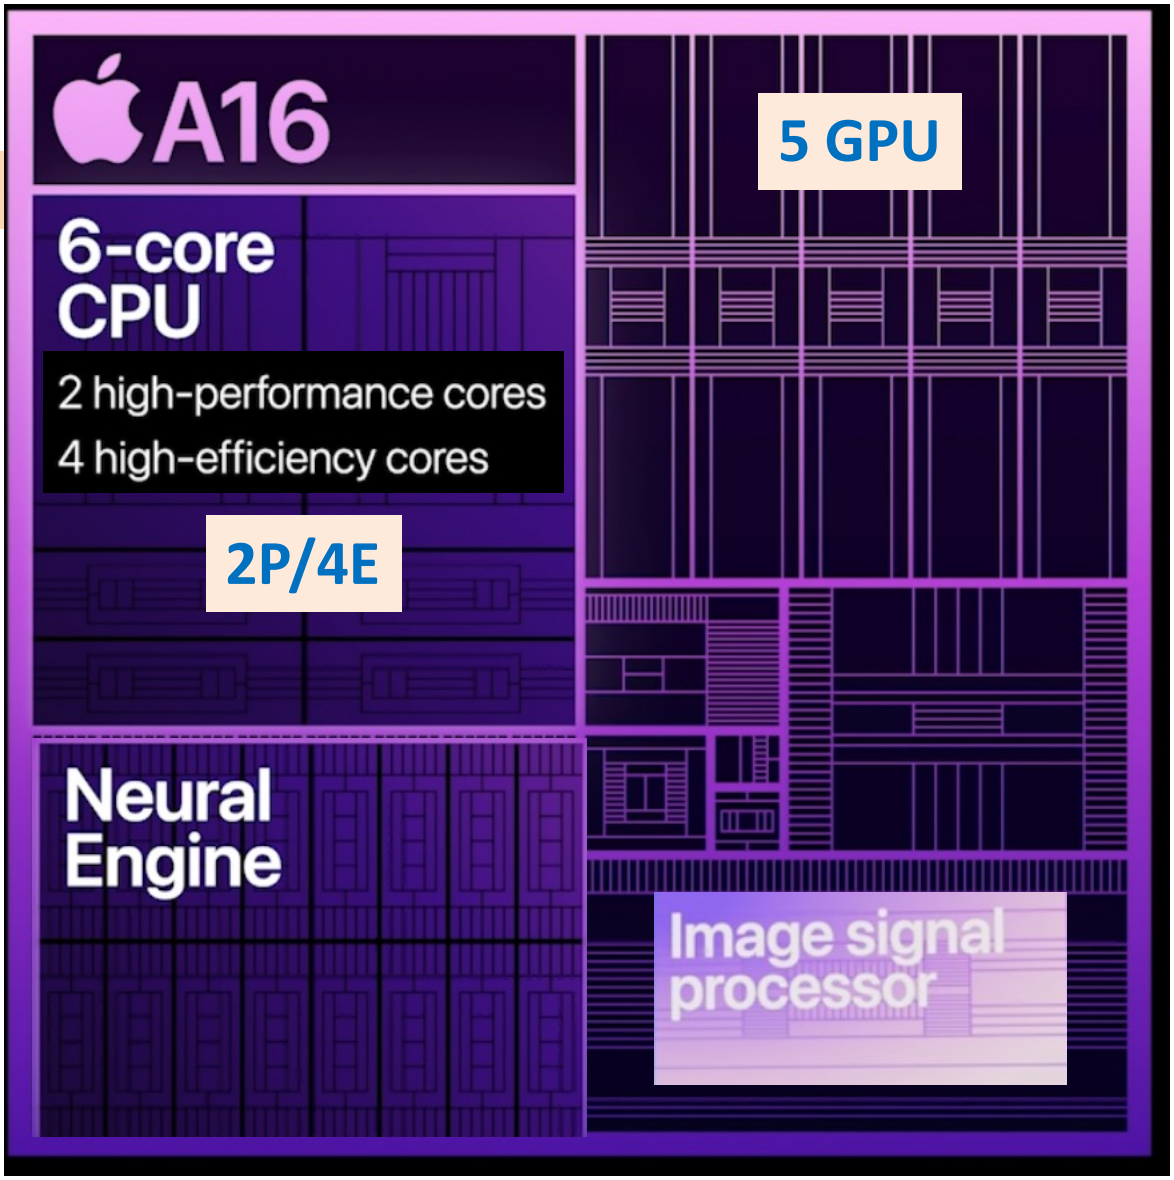

Apple Event

9-7-22

- iPhone **13/14**

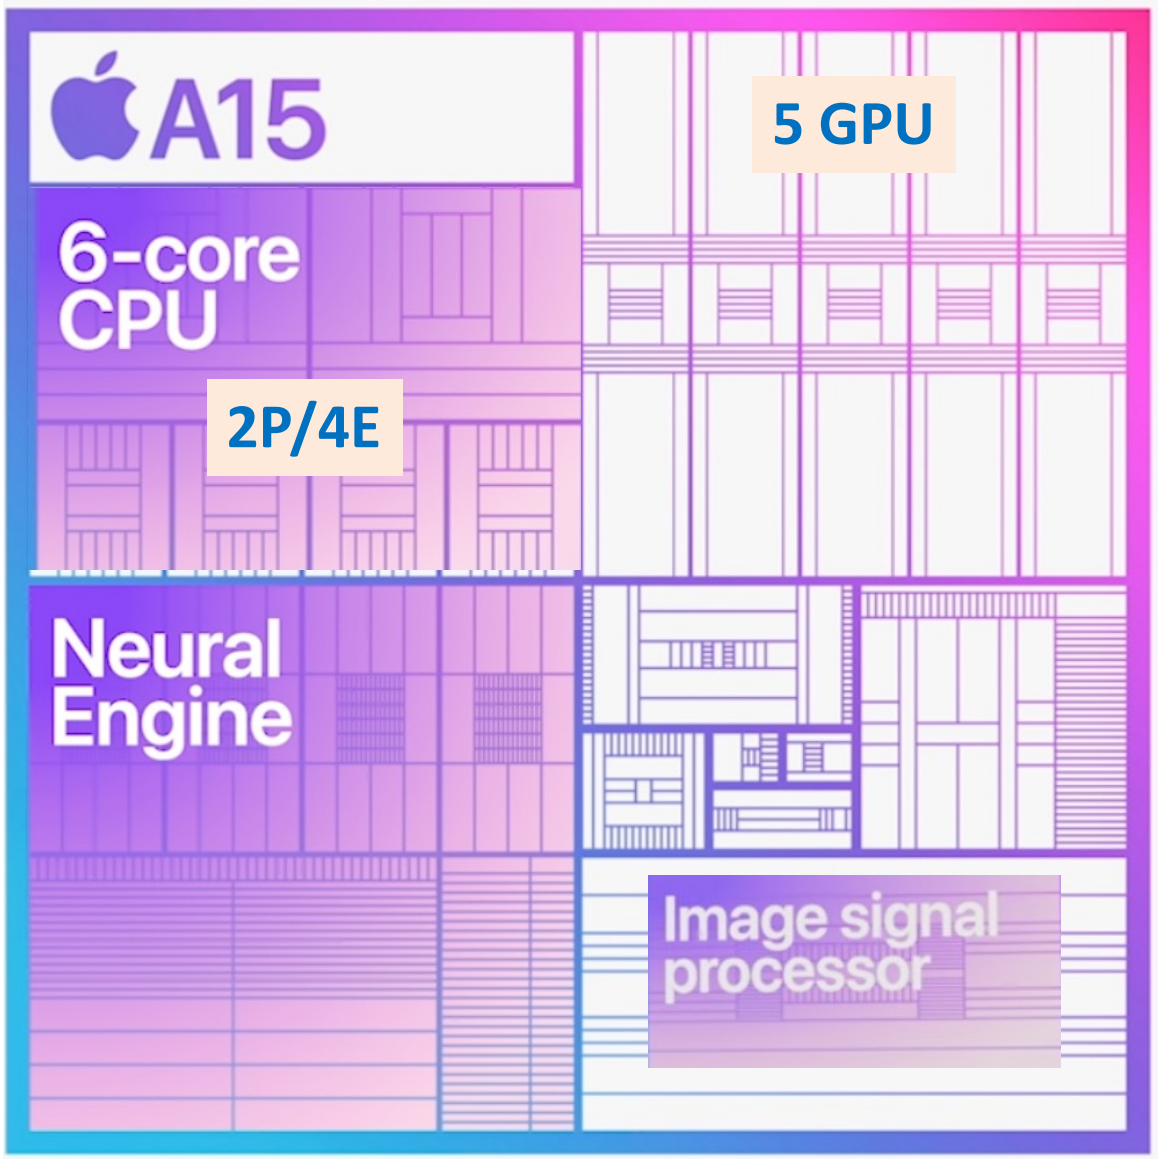

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Apple M1/2 for Mac

Quora

Why did Apple stop using Intel?

Jeff Drobman, Lecturer at California State University, Northridge (2016-present)

Answered just now

1. Intel had fallen behind in process nodes to the point where AMD was beating them by using 7nm at TSMC. Apple did not want to wait for Intel to catch up.
2. Apple wants control over their ISA's and architecture, and more control over their CPU ship supply.
3. Apple wants to unify their ISA's and architecture and thereby enable use of a single OS (MacOS will converge with iOS).

Mac M1/2 vs Core i7/9

Quora

What is the difference between the processor in a Mac and a PC with an equivalent CPU model?

Jeff Drobman
Lecturer at California State University, Northridge (2016–present) · Just now ·

Win PC's use x86 CPU chips from Intel (Core i7/i9) or AMD (Ryzen 5/7/9), with or without integrated graphics GPU cores. Macs used to also use these same chips, but now use Apple designed M1/2 SoC's with CPU and GPU cores (and NPU cores with ML accelerators). In a Mac, I suspect the AI support is mostly for its "Siri" assistant, but might be usable by some gaming software.

Apple ARM ISA

Quora

What is the latest version of the Advanced RISC Machine (ARM) architecture?

Jeff Drobman
Lecturer at California State University, Northridge (2016–present) · Just now ·

ARM has 3 families of architecture: A, R and M. they each have minute differences in their “micro” architecture (including multi-threading). but there are now only 2 basic ISA: 32-bit v7 and 64-bit v8. Licensees such as Apple have created their own ISA extensions, which are called v8.1, v8.2 etc.

Apple M1

November 10, 2020

11 trillion
Operations per second

11 Tera FLOPS

- Cores
  - 8 CPU
  - 8 GPU
  - 16 NPU

- CPU cores
  - 4 Hi Perf (20W **P**)
  - 4 Hi Efficiency (1.3W low power **E**)

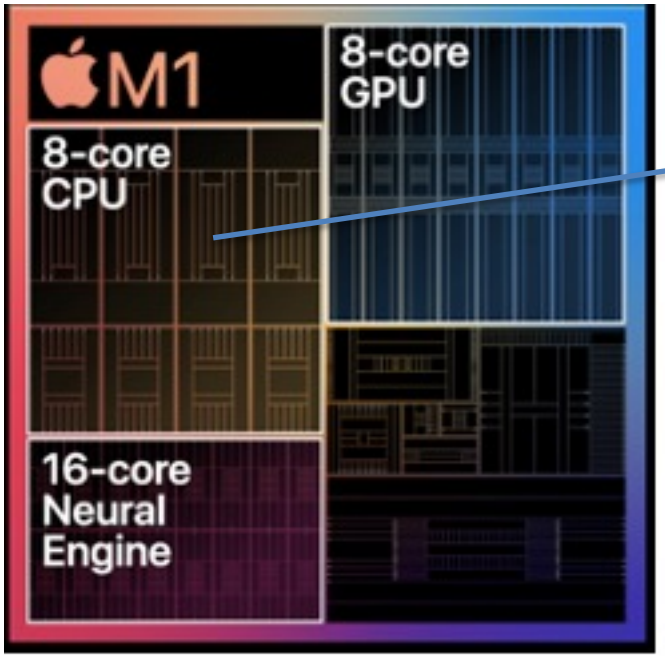

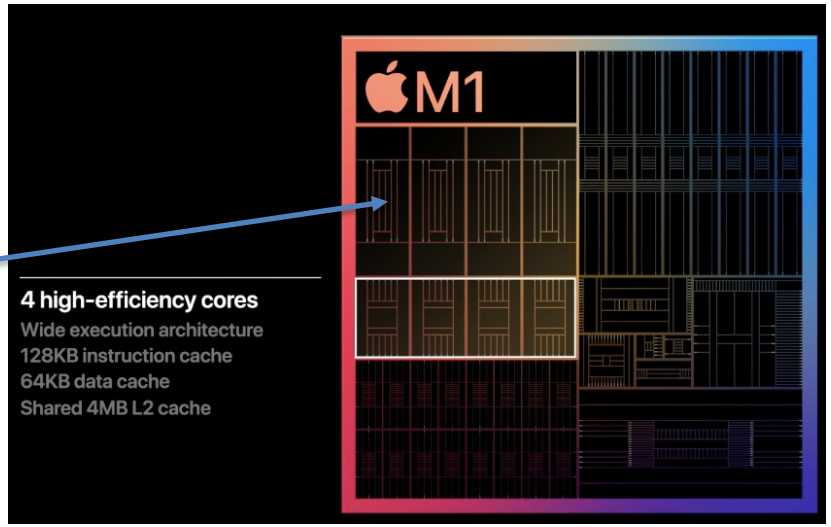

Apple M1 Module

5-nanometer process

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

16 billion transistors

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

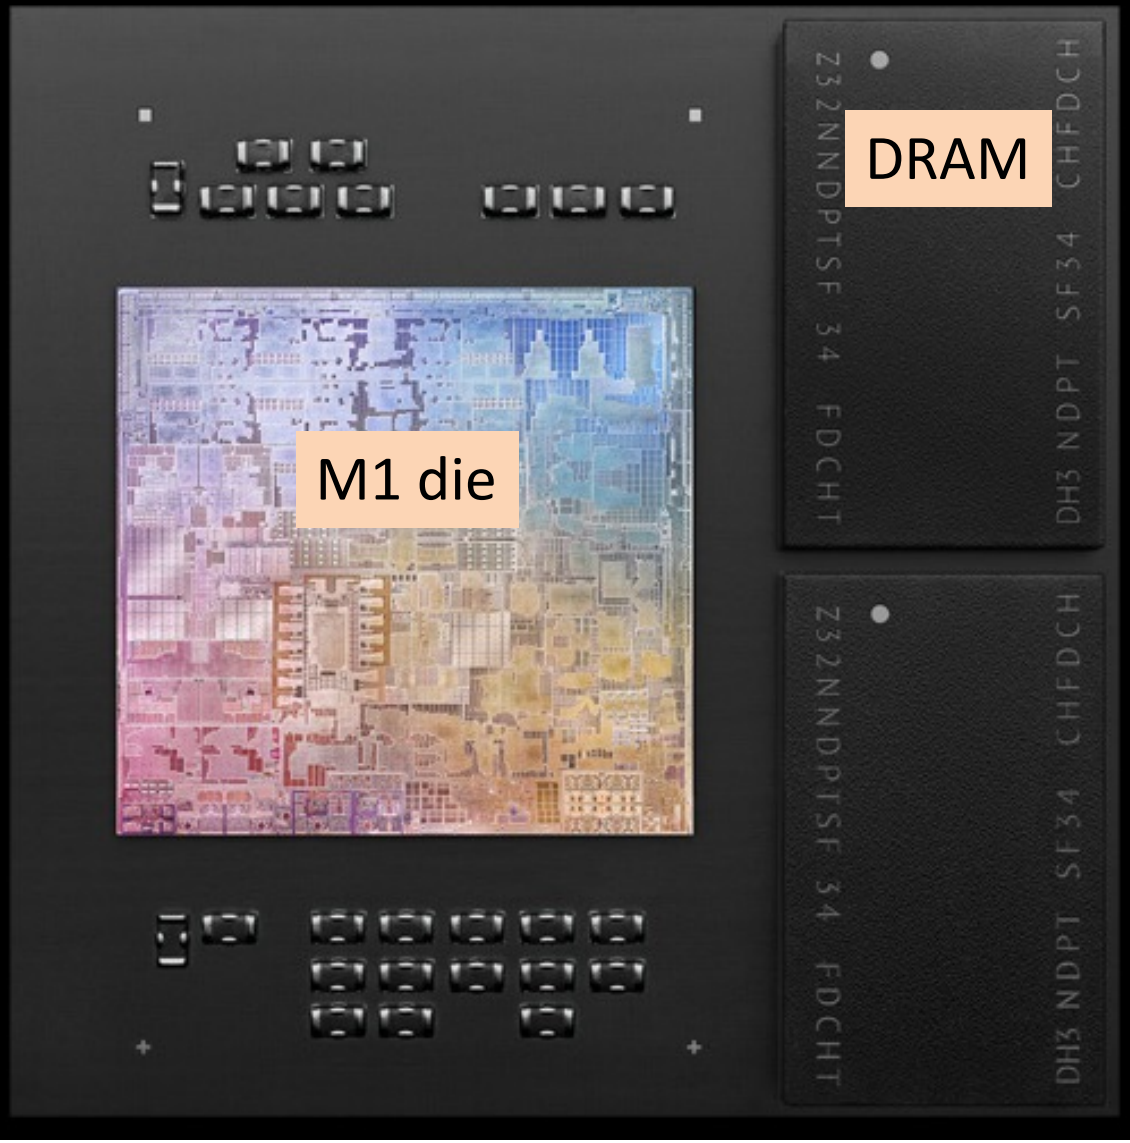

Apple M1 Max Die

Annotated

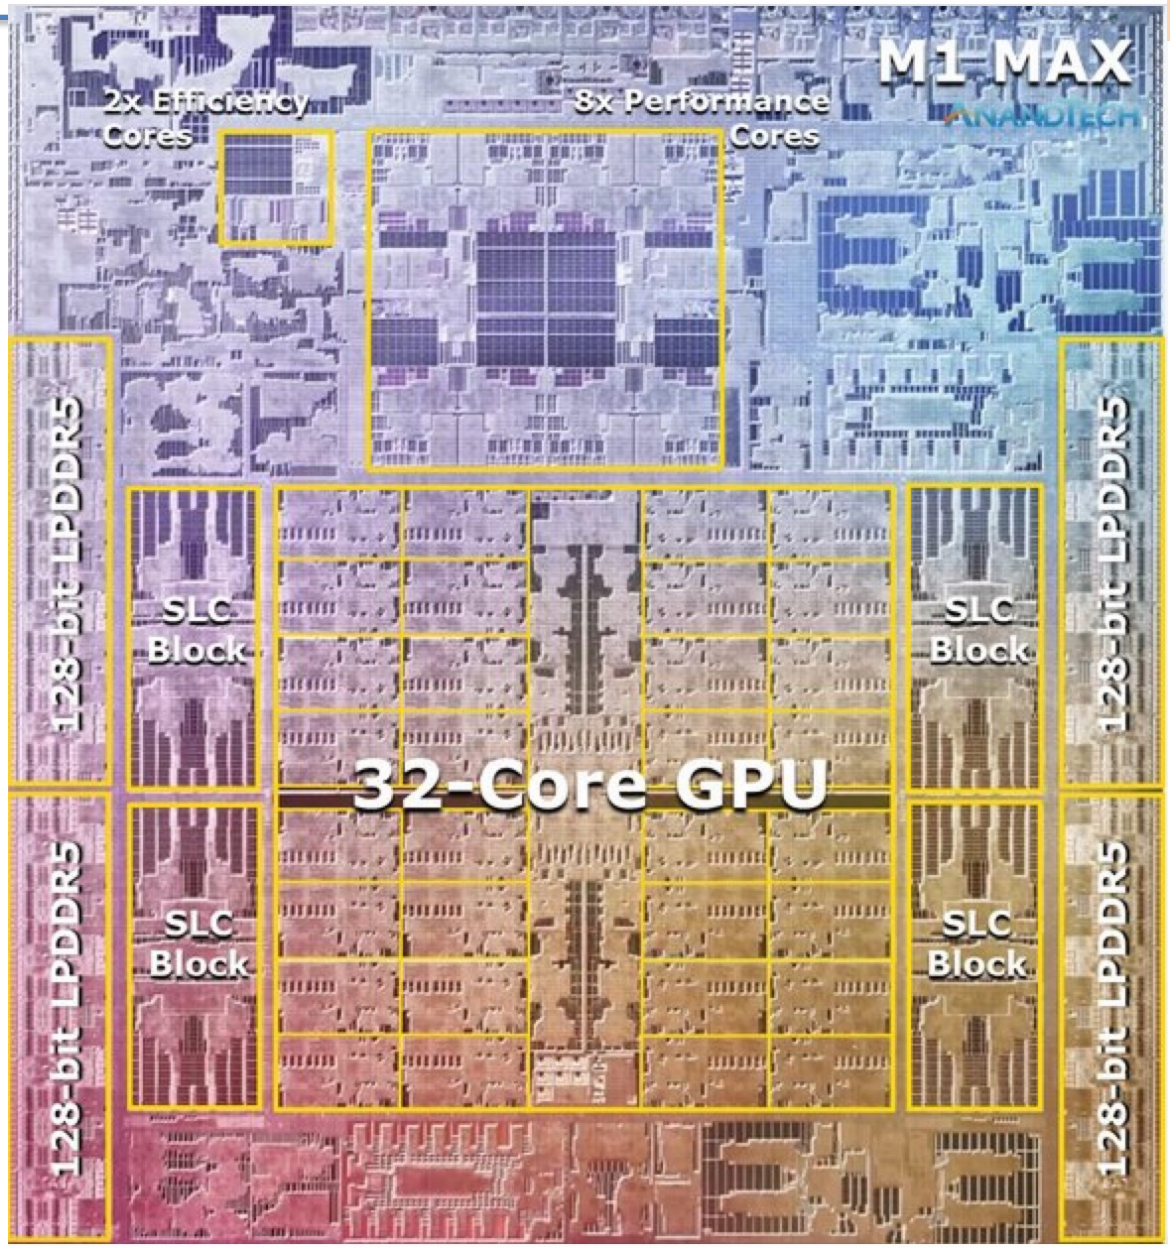

Apple M1 Models

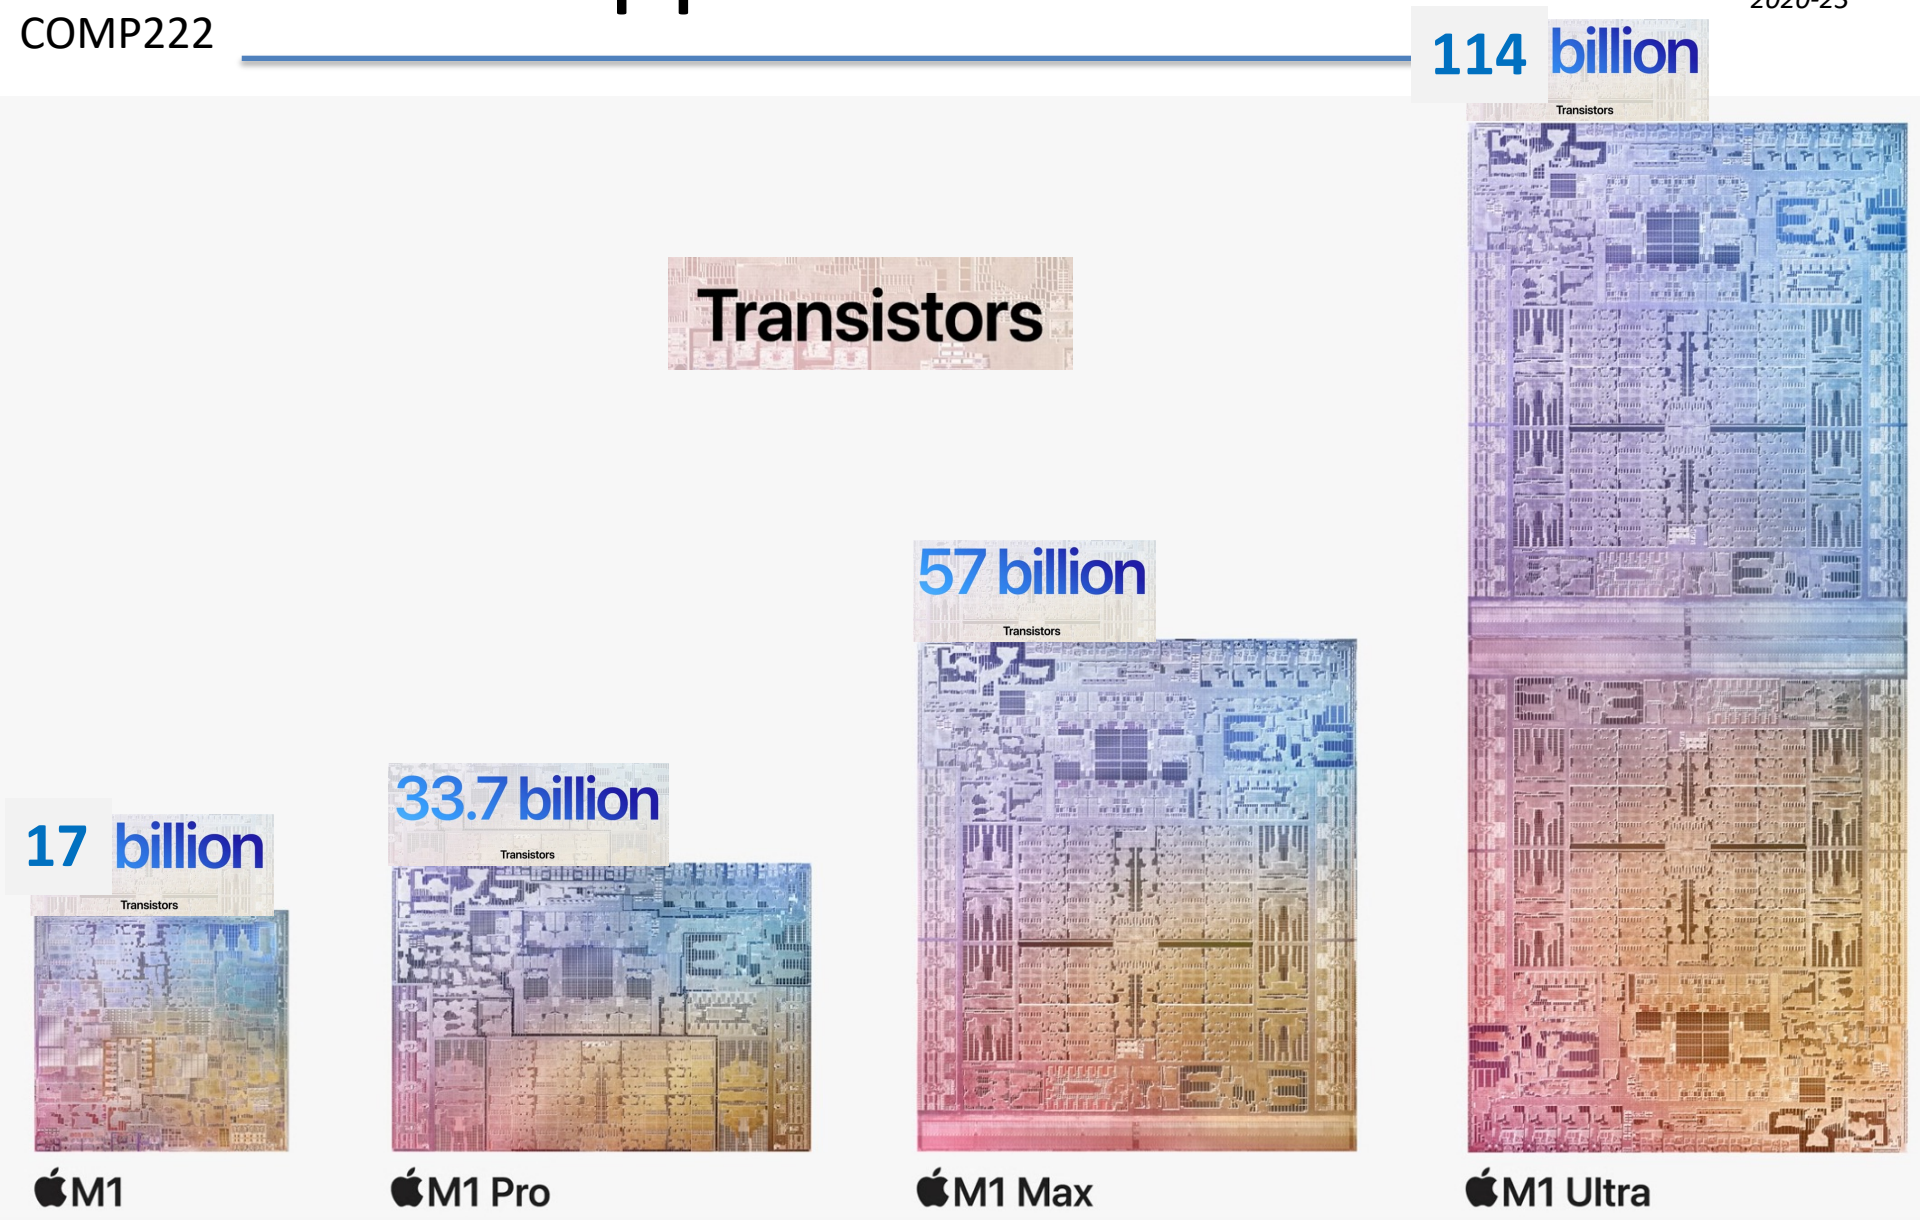

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Max/Pro

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Jan 17, 2023

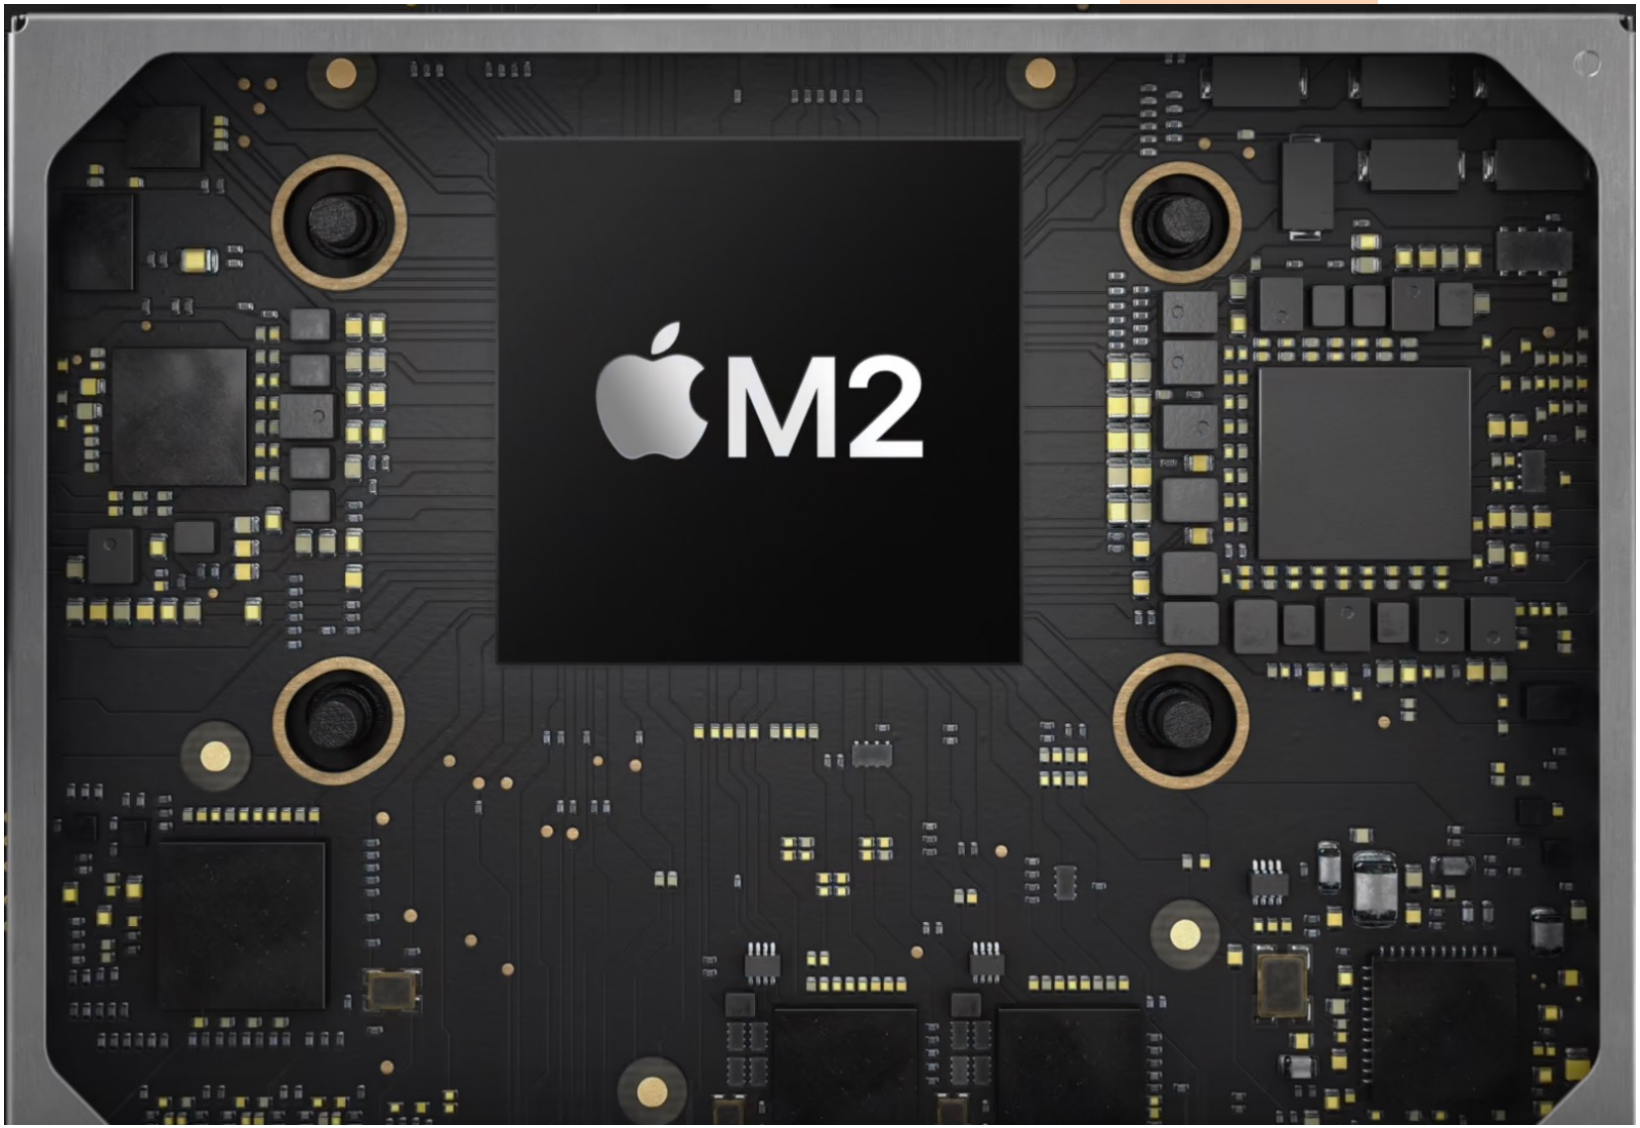

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Die

DR JEFF SOFTWARE
INDIE APP DEVELOPER

June 6, 2022

Second-generation
5 nanometer

17 billion
Transistors

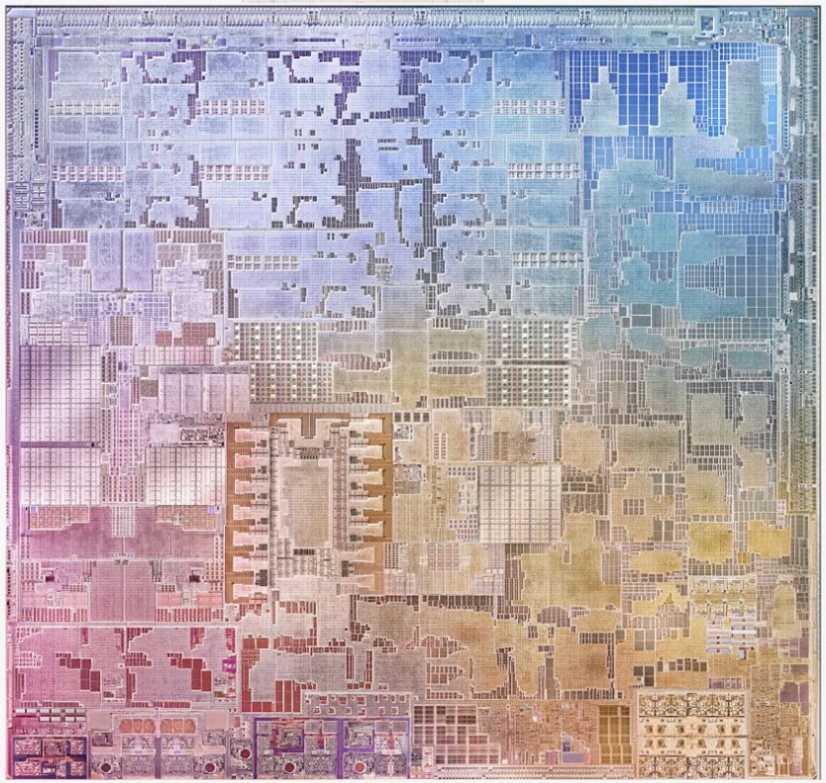

M1

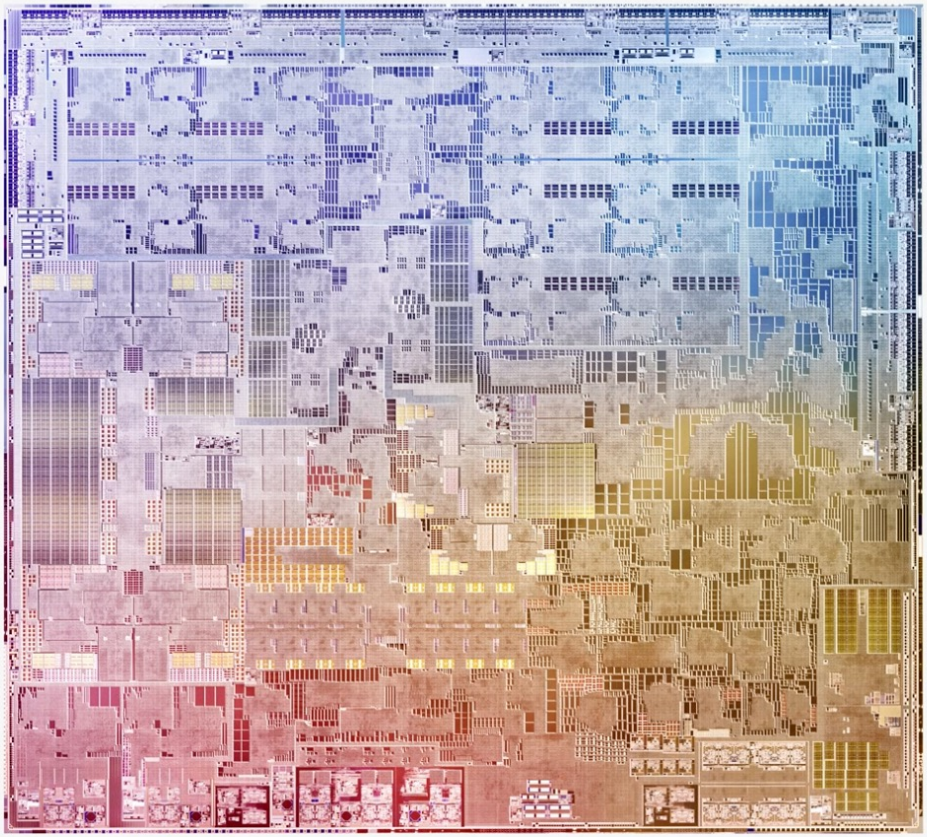

M2

20 billion
Transistors

Apple M2

8-core CPU

P

4 high-performance cores

Ultrawide microarchitecture

192KB instruction cache

128KB data cache

Shared 16MB cache

E

4 high-efficiency cores

Wide microarchitecture

128KB instruction cache

64KB data cache

Shared 4MB cache

10 GPU cores

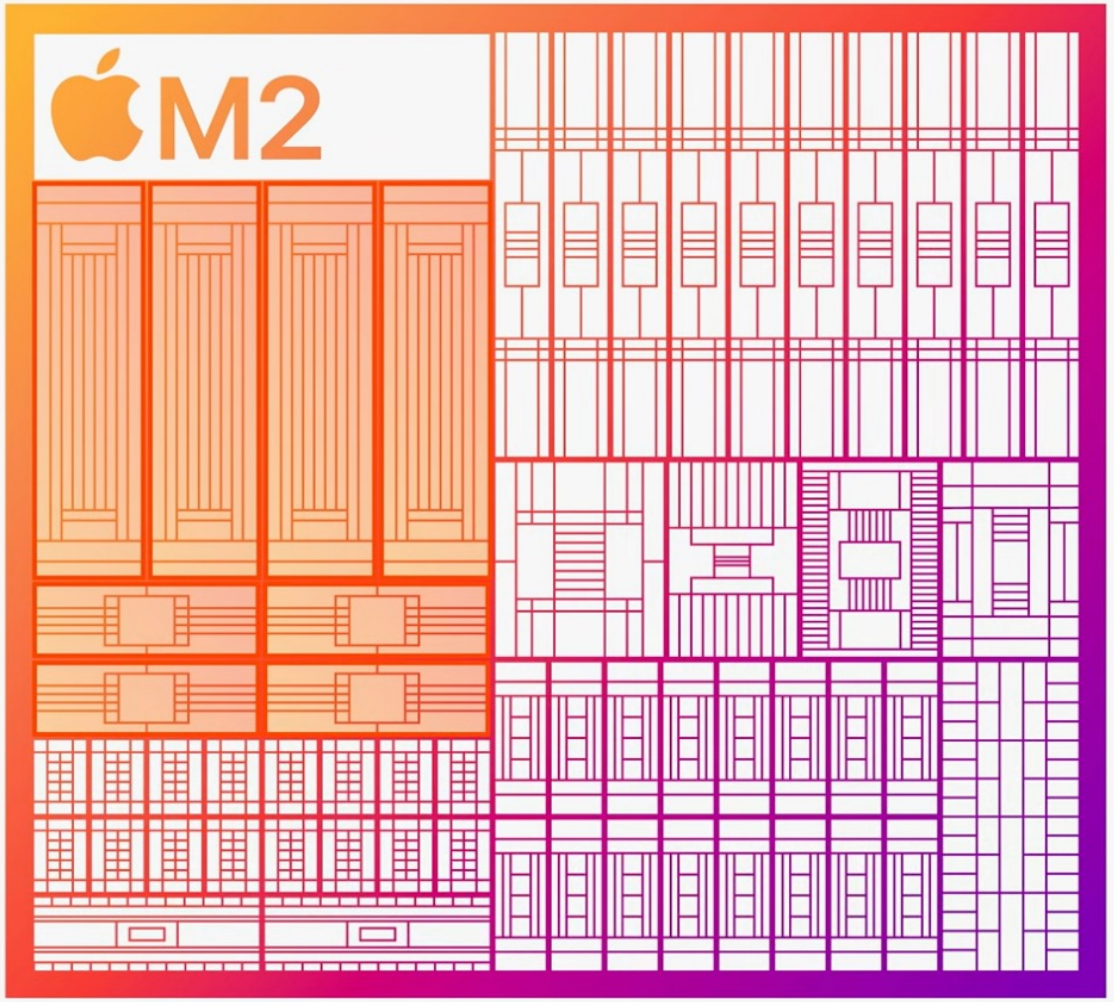

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Apple M2

Jan 17, 2023

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Apple M2 Mac Pro

Jan 17, 2023

CSUN
CALIFORNIA
STATE UNIVERSITY
NORTHRIDGE

DR JEFF
SOFTWARE
INDIE APP DEVELOPER

CO

Apple M2 Pro

Jan 17, 2023

Transistors

20 billion
Transistors

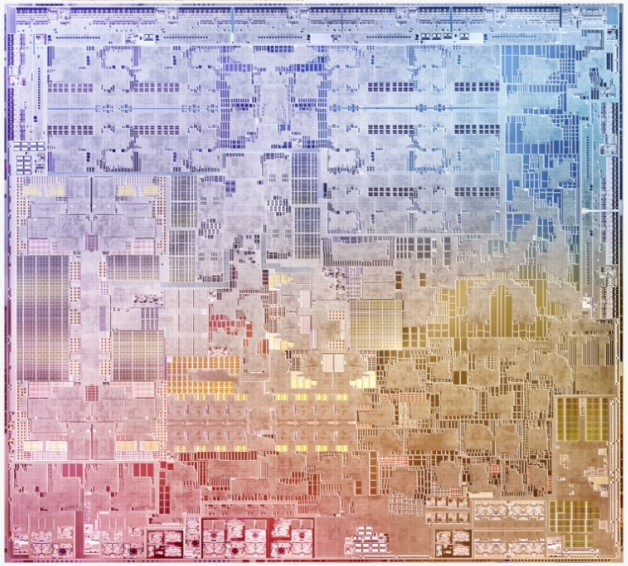

40 billion
Transistors

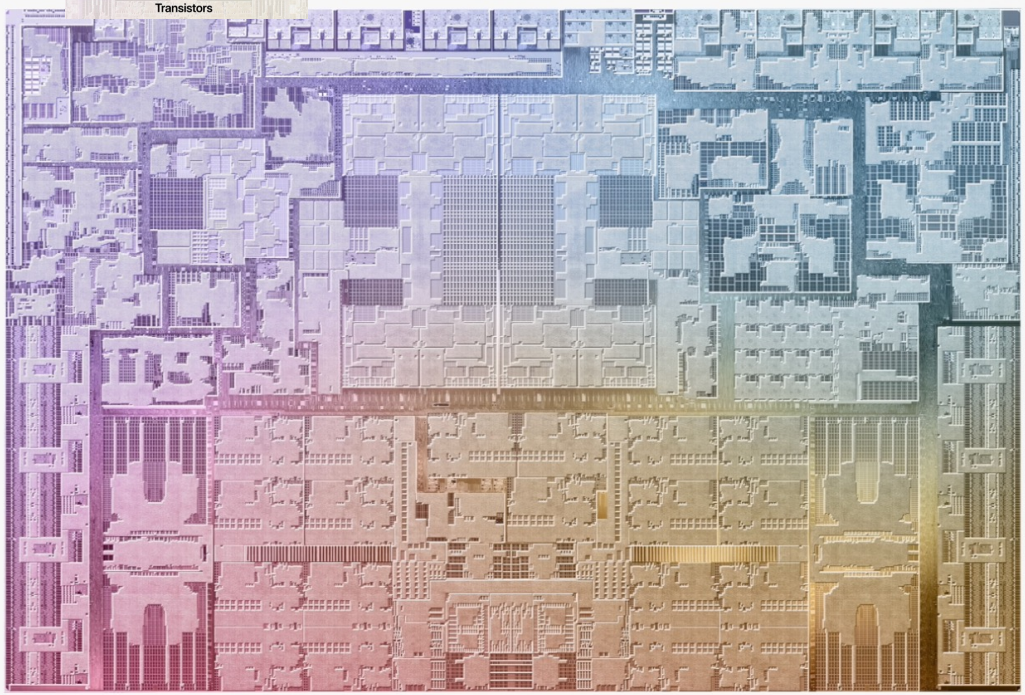

M2

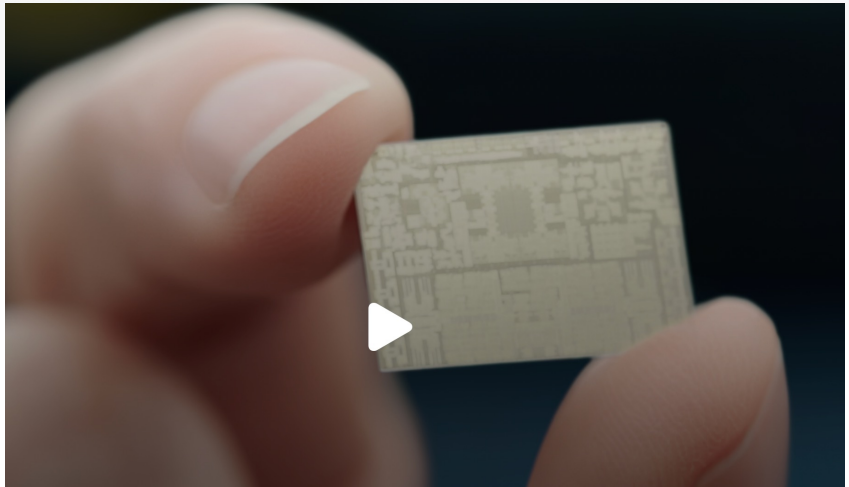

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Max

Jan 17, 2023

DR JEFF SOFTWARE
INDIE APP DEVELOPER

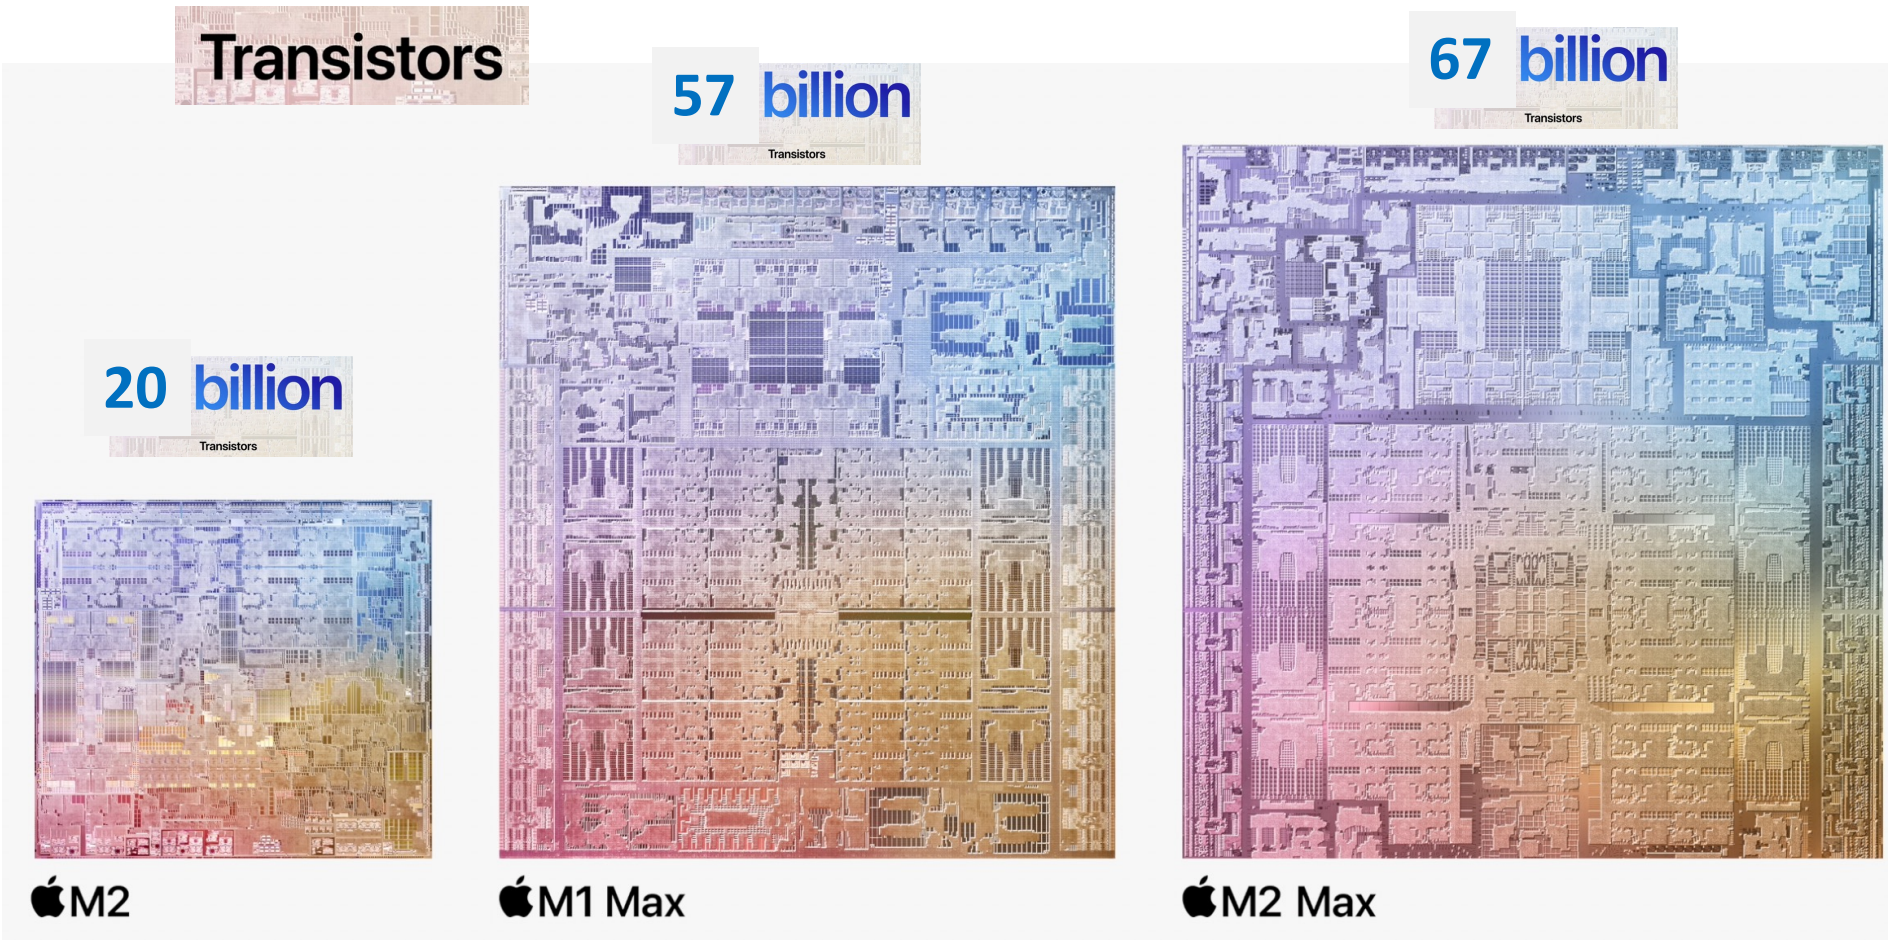

Apple M2 Pro

Jan 17, 2023

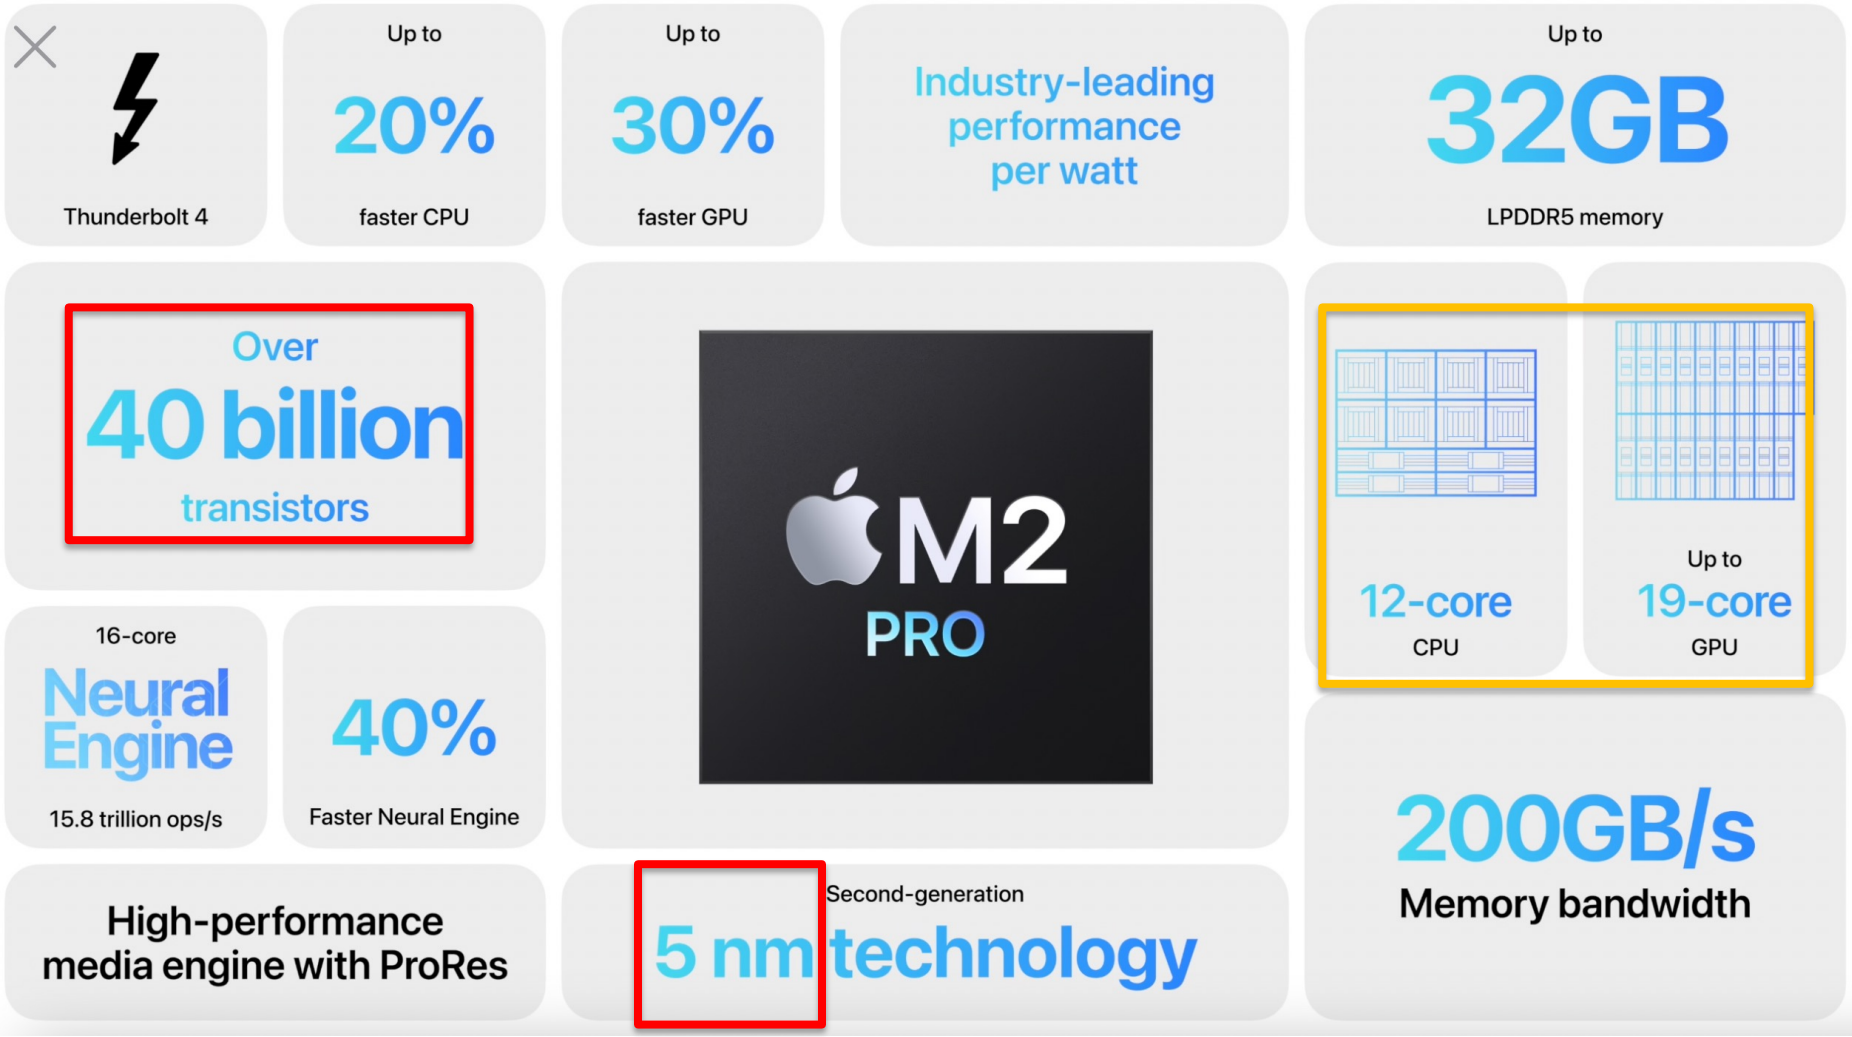

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Max

Jan 17, 2023

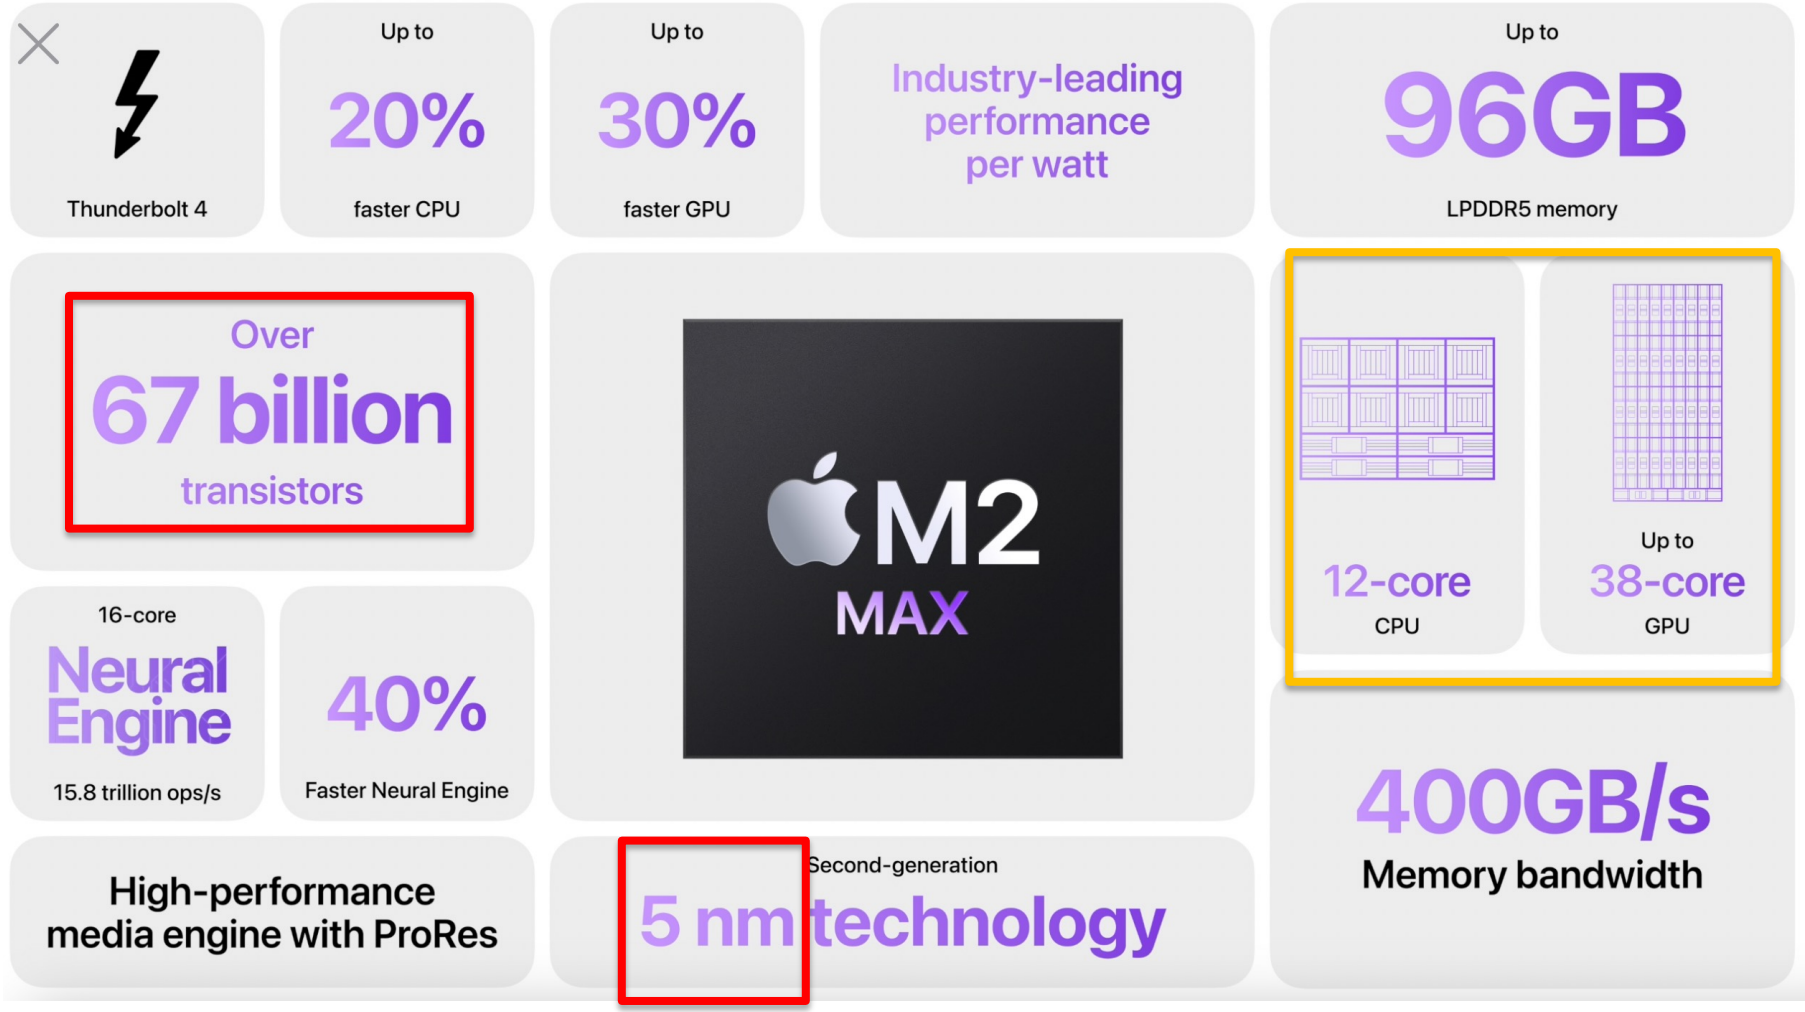

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Pro

Jan 17, 2023

DR JEFF SOFTWARE
INDIE APP DEVELOPER

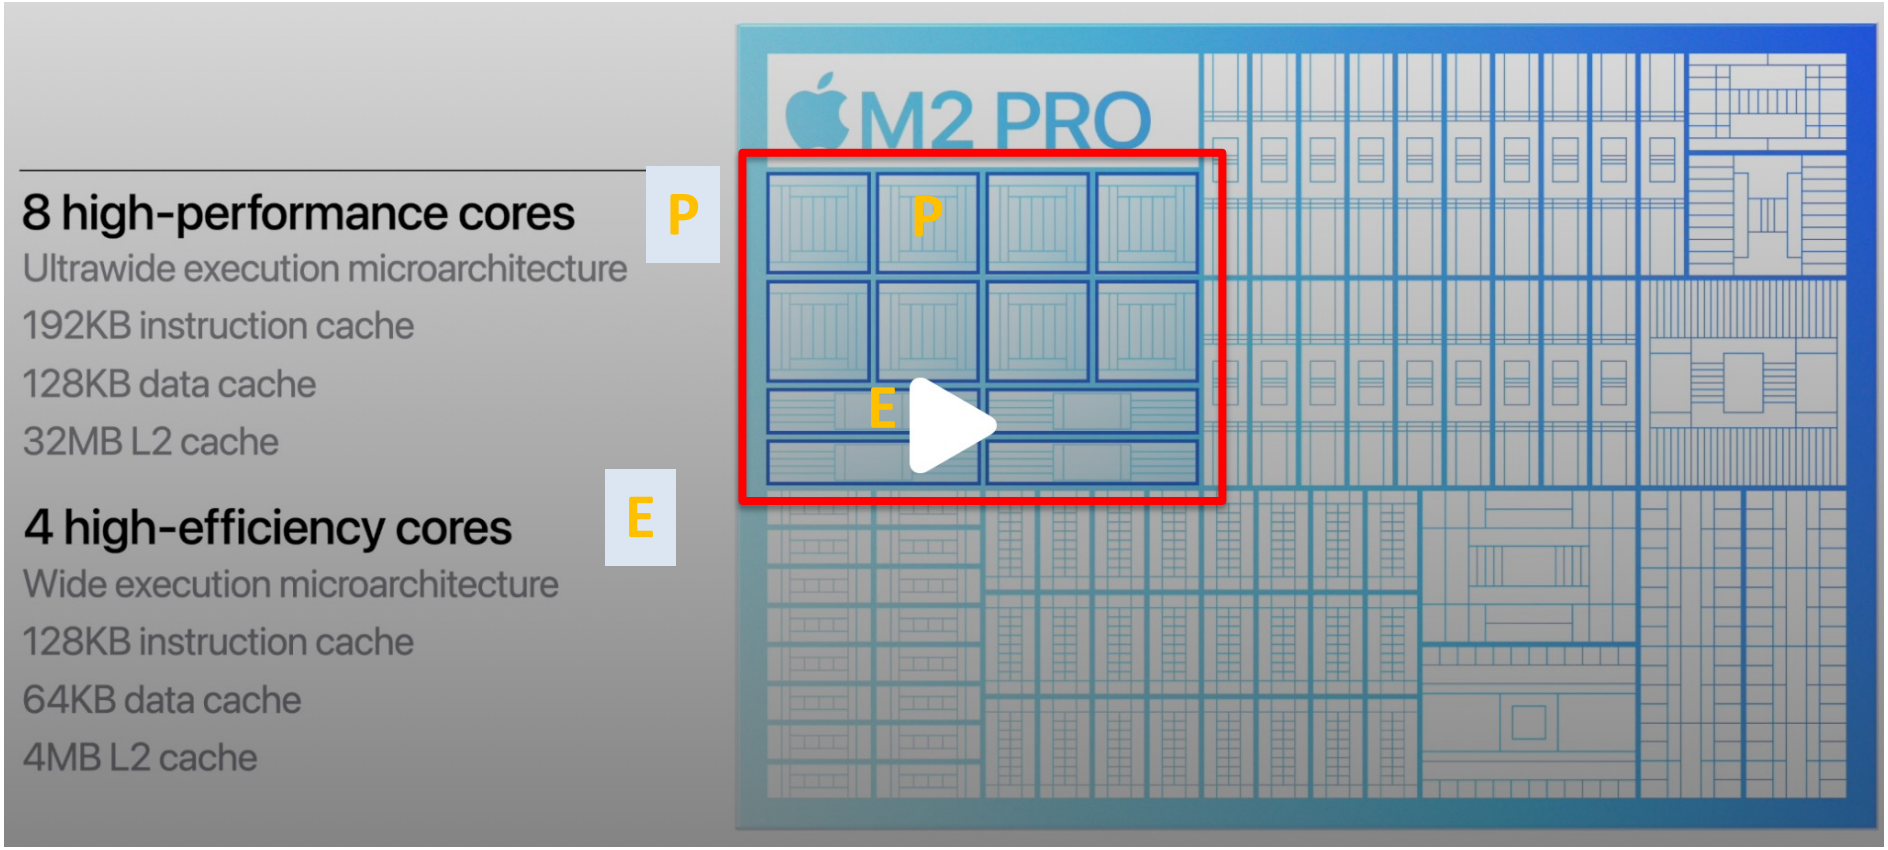

Apple M2 Max

Jan 17, 2023

(same)

8 high-performance cores P

Ultrawide execution microarchitecture

192KB instruction cache

128KB data cache

32MB L2 cache

4 high-efficiency cores E

Wide execution microarchitecture

128KB instruction cache

64KB data cache

4MB L2 cache

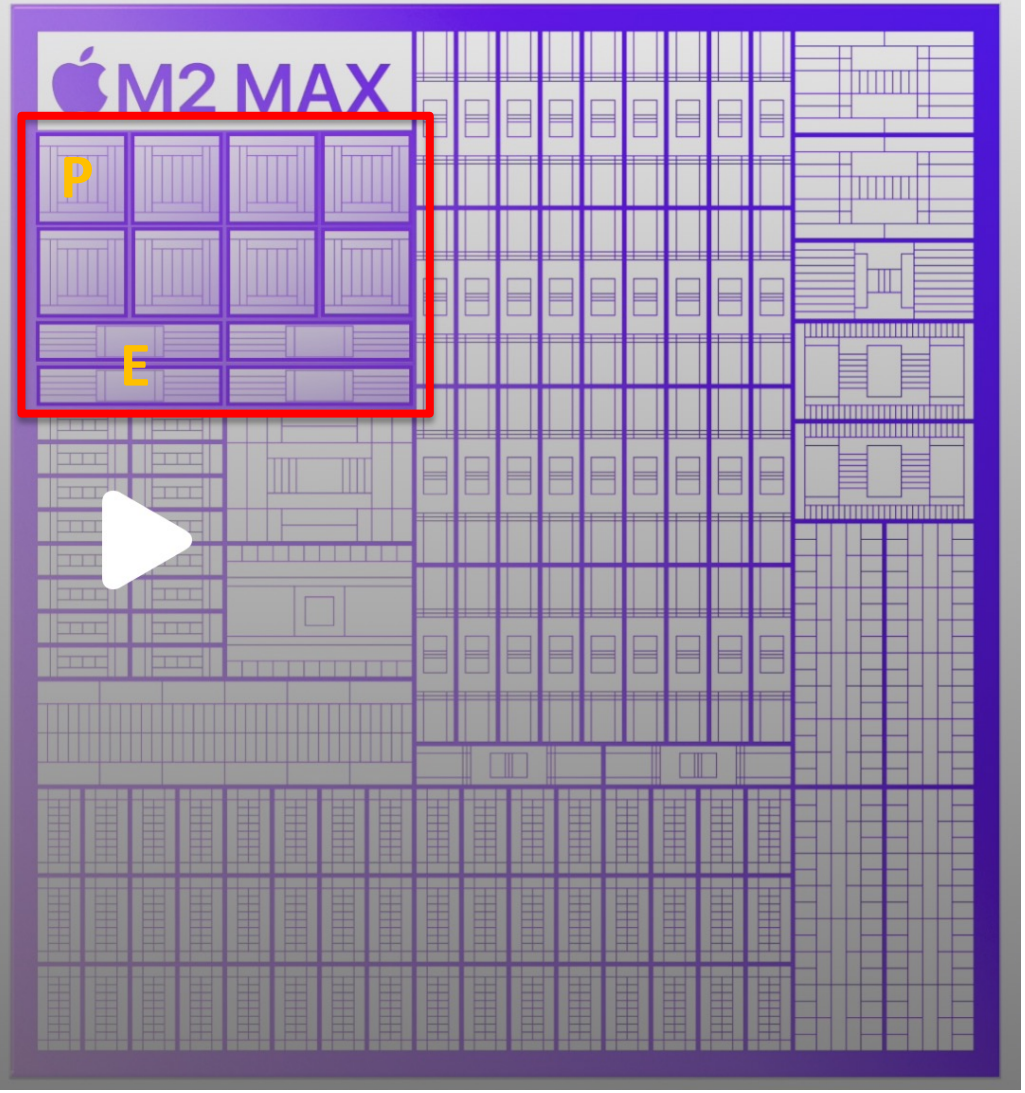

Apple M2 Pro

Jan 17, 2023

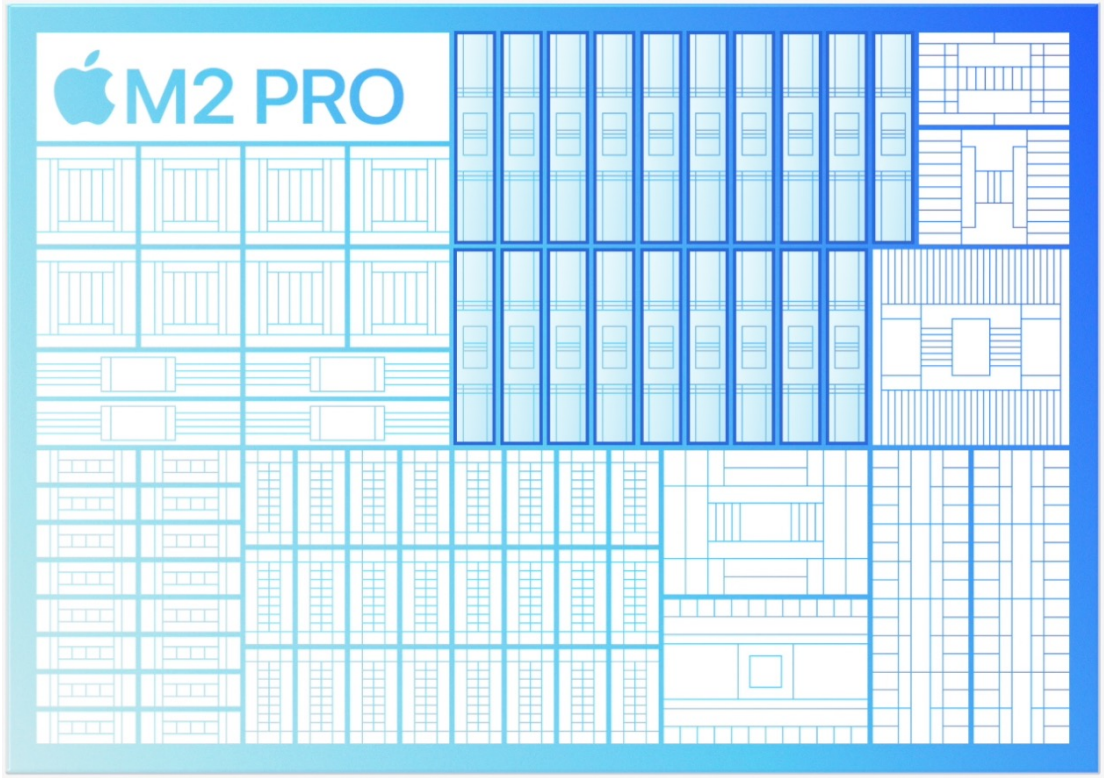

19-core GPU
2,432 execution units
6.8 teraflops
212 gigatexels/second
106 gigapixels/second

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Max

Jan 17, 2023

DR JEFF SOFTWARE
INDIE APP DEVELOPER

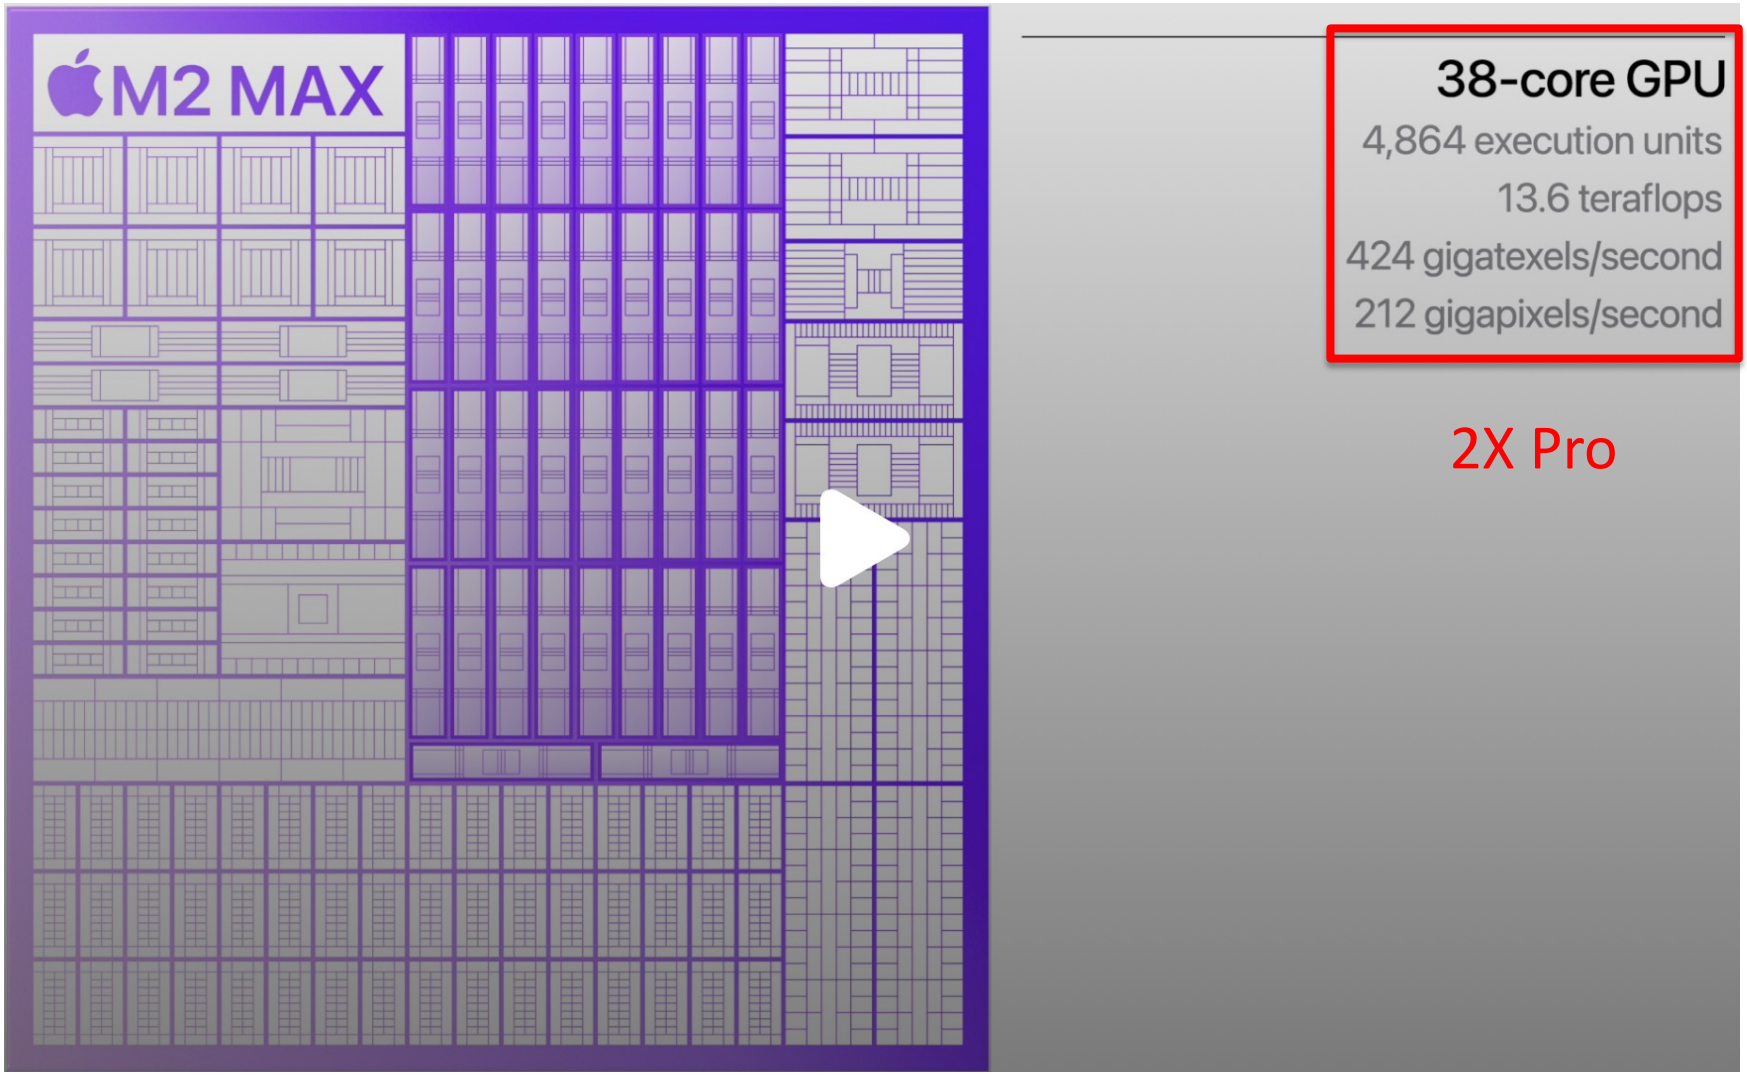

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Pro Module

Jan 17, 2023

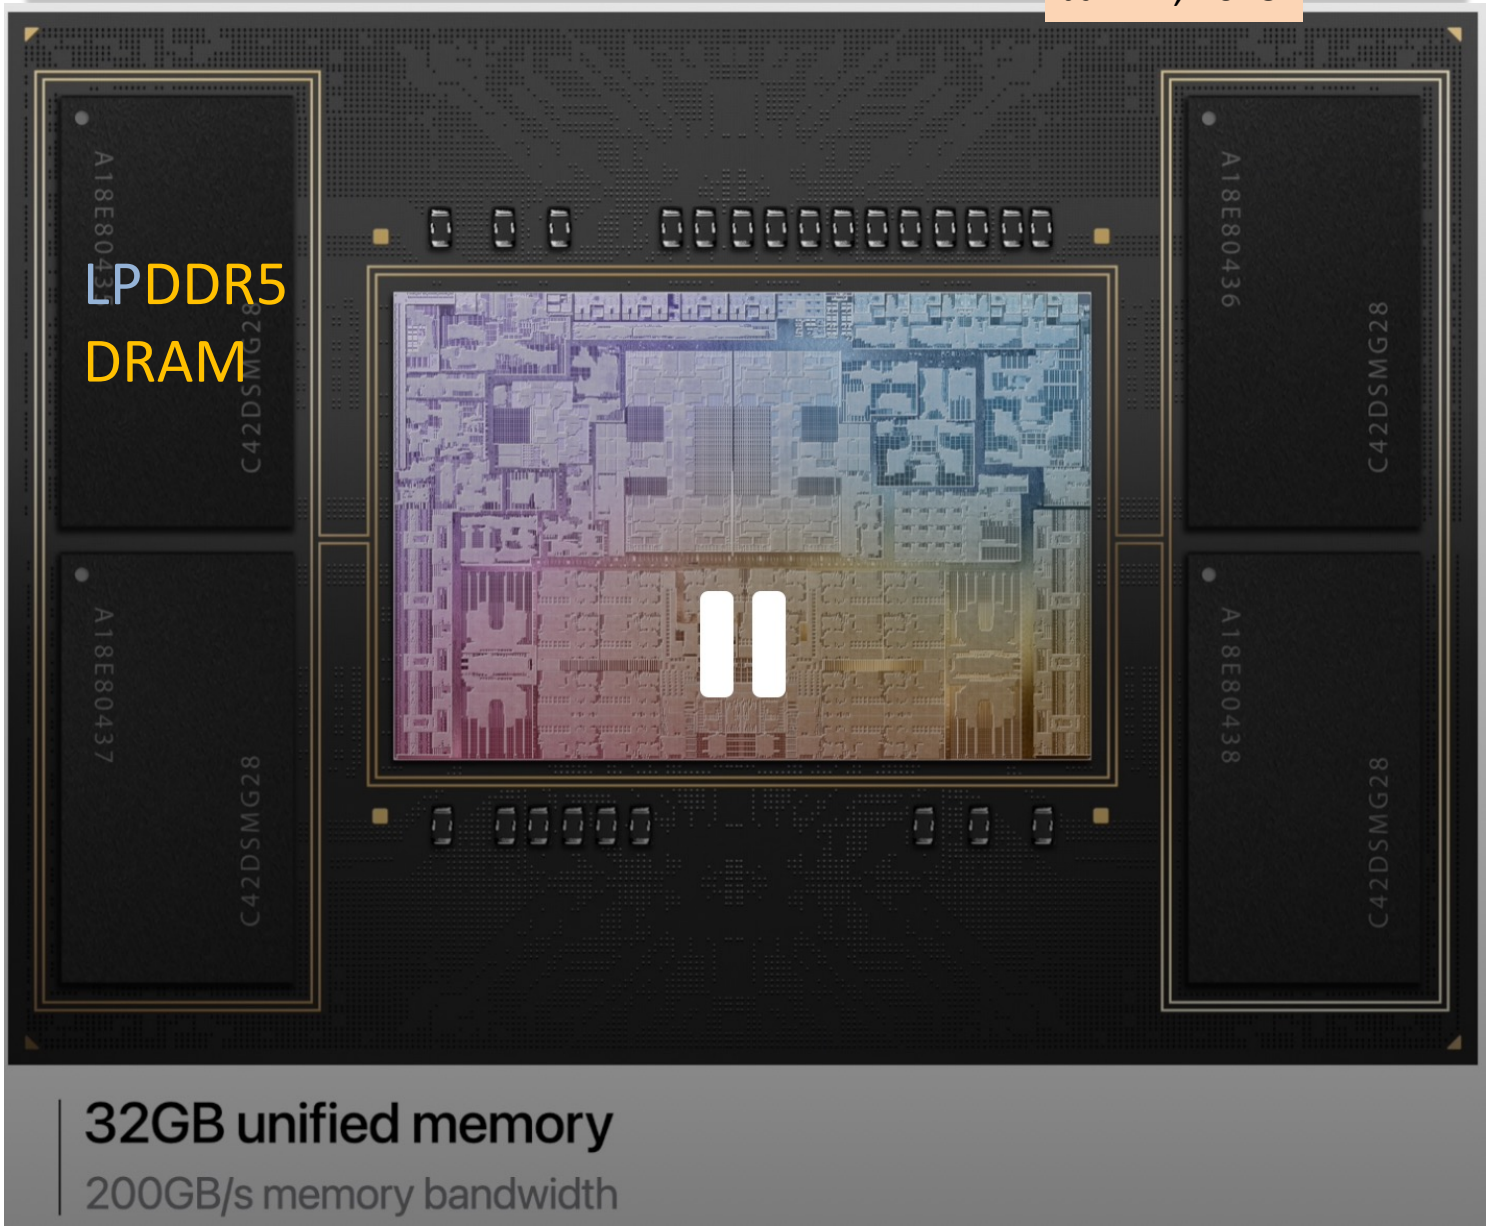

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple M2 Max Module

Jan 17, 2023

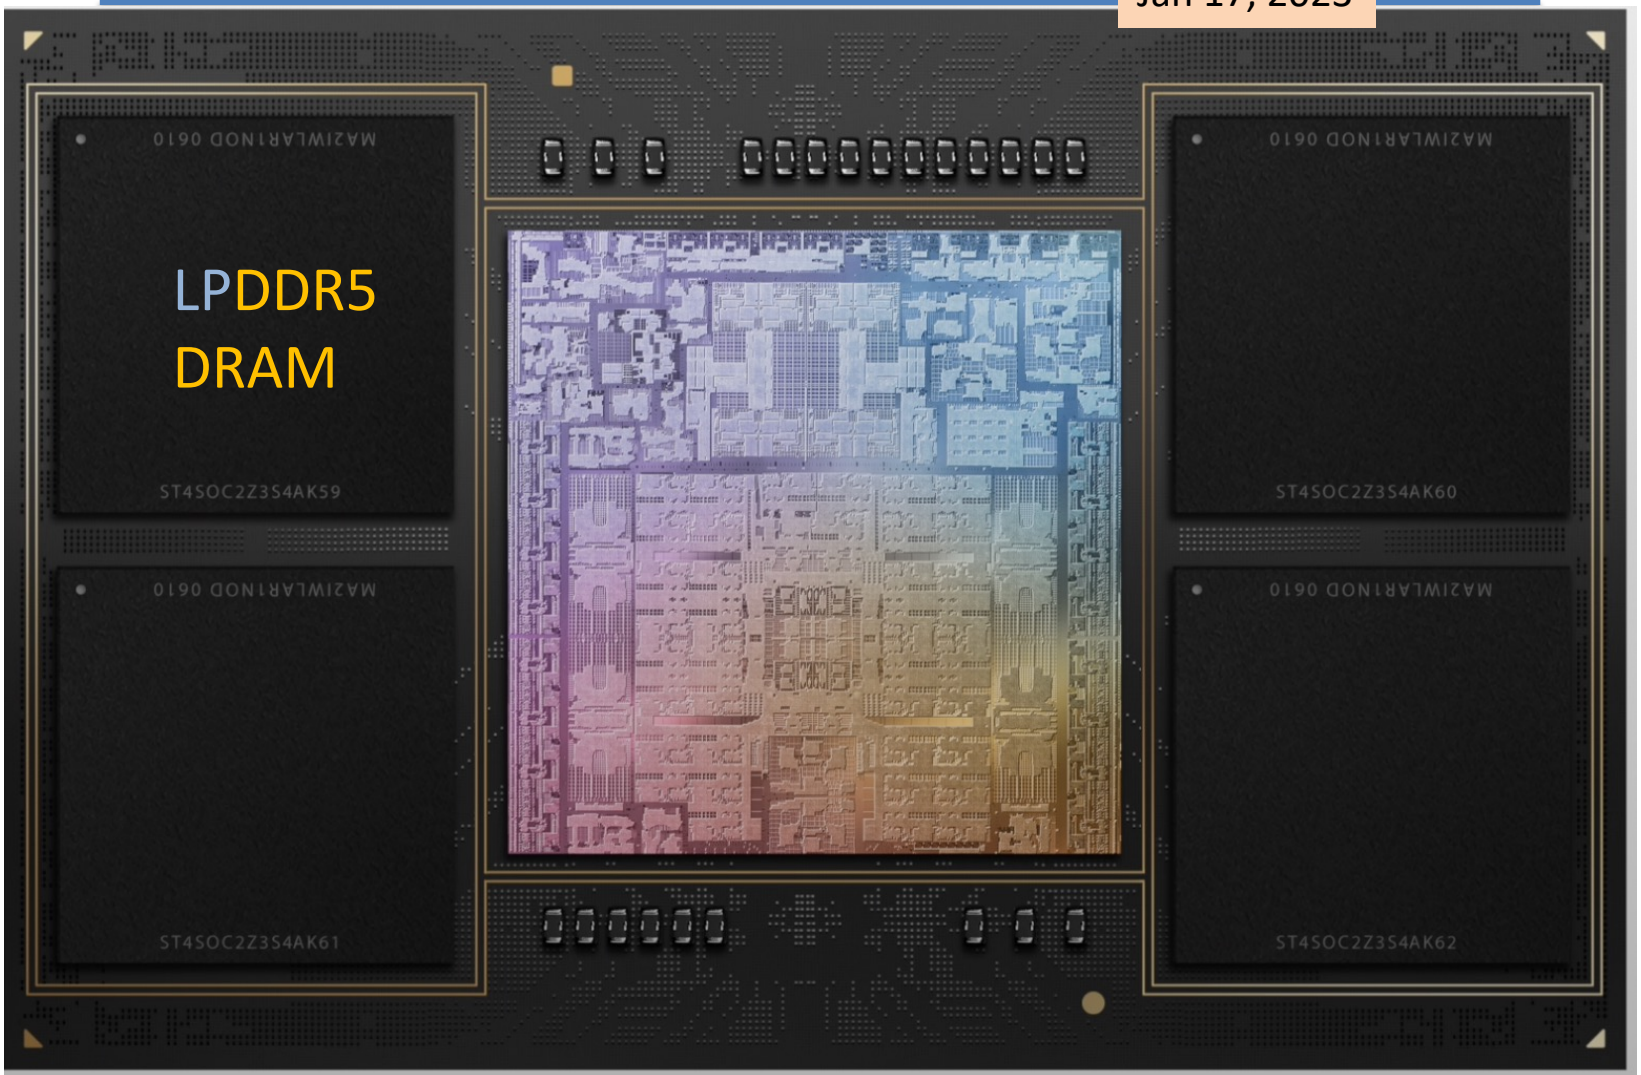

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

Apple Xcode/Swift

Jan 17, 2023

DR JEFF SOFTWARE
INDIE APP DEVELOPER


```
    renderer.c_axiom_solver.setSourceDensity(Float32(sender.value));
}
@IBAction func sliderTemperature(_ sender: UISlider) {
    renderer.c_axiom_solver.setSourceTemperature(Float32(sender.value));
}
@IBAction func sliderBuoyancy(_ sender: UISlider) {
    renderer.c_axiom_solver.setBuoyancy(Float32(sender.value))
}
@IBAction func sliderBuoyancy(_ sender: UISlider) {
    renderer.c_axiom_solver.setBuoyancy(Float32(sender.value))
}
override func touchesBegan(_ touches: Set<UITouch>, with event: UIEvent?) {
    for touch in touches {
        allTouches.append(touch)
```

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Apple Xcode/Swift

Jan 17, 2023

DR JEFF SOFTWARE
INDIE APP DEVELOPER


```
textureDescriptor_screenSelection_depth.usage = MTLTextureUsage.unknown
textureDescriptor_screenSelection_depth.pixelFormat = MTLPixelFormat.depth32Float_stencil8
textureDescriptor_screenSelection_depth.storageMode = .private

textureDescriptor_screenSelection_depth.width = 256
textureDescriptor_screenSelection_depth.height = 256

self.c_axiom_texture_screenSelection_depth = self.device.makeTexture(descriptor:
    textureDescriptor_screenSelection_depth)!

self.c_axiom_buffer_screenSelection = self.device.makeBuffer(length: MemoryLayout<Float>.size *
    textureDescriptor_screenSelection_depth.width * textureDescriptor_screenSelection_depth.height, options:
    [.storageModeShared])

self.c_axiom_renderPassDescriptor_screenSelection = MTLRenderPassDescriptor()

self.c_axiom_renderPassDescriptor_screenSelection.colorAttachments[0].texture = self.c_axiom_texture_screenSelection

self.c_axiom_renderPassDescriptor_screenSelection.depthAttachment.texture = self.c_axiom_texture_screenSelection_depth
self.c_axiom_renderPassDescriptor_screenSelection.stencilAttachment.texture = self.c_axiom_texture_screenSelection_depth

self.c_axiom_renderPassDescriptor_screenSelection.colorAttachments[0].loadAction = .clear;
self.c_axiom_renderPassDescriptor_screenSelection.colorAttachments[0].clearColor = MTLClearColorMake(0,0,0,0);
self.c_axiom_renderPassDescriptor_screenSelection.colorAttachments[0].storeAction = .store;

let textureDescriptor = MTLTextureDescriptor.textureBufferDescriptor(with: MTLPixelFormat.rgba16Float, width: 1,
    resourceOptions: [], usage: MTLTextureUsage.shaderWrite)
```

State of the Art

Qualcomm

- Snapdragon

Apple M1 SoC

Quora

Could Qualcomm have developed an M1 competitor without ex-Apple engineers?

Jeff Drobman
Lecturer at California State University, Northridge (2016–present) · Just now

anything is possible. but Apple has been a partner with ARM since about 1990, and has always had the ability and license to customize the CPU and GPU cores. base ARM cores were single-threaded. I suspect Apple extended them to SMT2 (like AMD) or even hyper-threading (like Intel). then Apple probably also modified ARM Mali GPU cores into an improved or extended design.

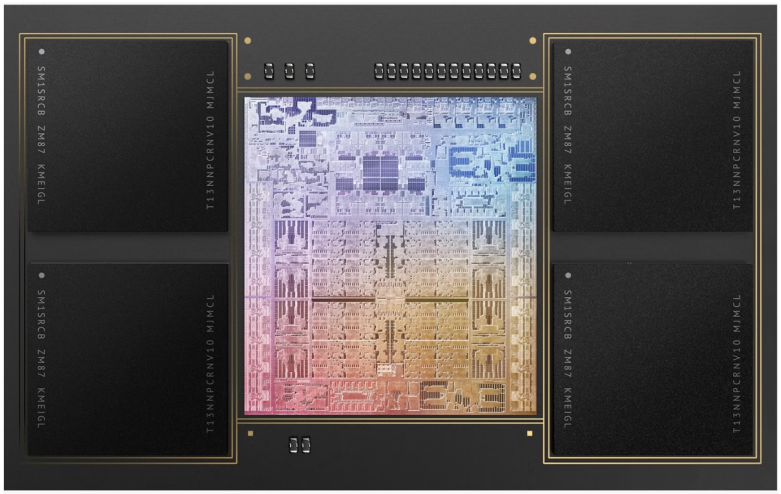

64GB unified memory
High bandwidth, low latency
512-bit LPDDR5 interface
Apple-designed custom package

Qualcomm

Qualcomm will continue using **multiple foundries** to manufacture its chips, one of its executives said on Tuesday.

The US chip giant was continuing to collaborate with **Samsung** Foundry and will use multiple foundries, from the South Korean company, **TSMC** to **GlobalFoundries**, depending on their technological maturity going forward, Qualcomm senior vice president Don McGuire said during Qualcomm's Snapdragon Summit 2022 in Hawaii in a meeting with South Korean press.

Qualcomm has currently given all its workload for 4-nanometer (nm) and 3nm chips to **TSMC**, the world's largest foundry. McGuire's comments indicate that the US firm could give orders to **Samsung** Foundry again for follow-up nodes such as gate-all-around (**GAA**).

At the summit, Qualcomm unveiled its latest application processor the **Snapdragon 8** Gen 2. The chip boasts 4.35 times increased AI performance and 25% faster processing speed compared to its predecessor.

Snapdragon 8 Gen 2 also uses Qualcomm's new **GPU** and has a CPU with 40% increase in performance. The chip will be made using **TSMC's 4nm** node. For Snapdragon 8 Gen 1, **Samsung** Foundry was the initial contract manufacturer but Qualcomm gave the order to **TSMC** during the second half of last year.

Qualcomm

Sources had said **Samsung's** low yield rate for 4nm was the cause and the US chip firm will be giving **TSMC** the order for 3nm chips because of this.

McGuire said Qualcomm's orders were too large for it to use a single foundry and using multiple foundries is not only advantageous in supply but also price and scale.

The US firm also needed multiple foundries to expand in other business areas besides smartphones, he added.

Meanwhile, **Samsung**, while it facing difficulty with yield with 4nm, was the first to start production of 3nm GAA chips.

TSMC has also started 3nm chip production but for these uses a FinFET structure rather than GAA. The Taiwanese giant is reportedly planning to apply GAA structure starting with 2nm.

State of the Art

Samsung

- Exynos

Quora Samsung Exynos SoC

Drazen Zoric · Follow
Lives in Cork, Ireland ·

Samsung is huge enterprise made of many smaller and separate Samsung companies. For electronics two are most important:

- Samsung Electronics - has fabs for IC manufacturing
- Samsung Semiconductor - designs CPU, SoC and everything else

Samsung Semi most known product is Exynos SoC used in Samsung Galaxy phones. Well, US, S Korea use Qualcomm Snapdragon SoC while rest of the World uses Exynos. Is Exynos good? Well, it is not so good as Snapdragon or high end MediaTek (Dimensity). But later on this.
For example latest Exynos 2200 uses AMD Radeon GPU and chip was total disaster. So big disaster that Samsung canceled big presentation show 1 day before show. Chip was very slow and ran very hot. In meantime they fixed some problems but still no show.

Quora Samsung Exynos SoC

Drazen Zoric · Follow
Lives in Cork, Ireland ·

Samsung Elec has fabs for IC manufacturing. They produce Exynos as well as for others, eg NVidia, Qualcomm, etc. Samsung Elec was first one to commercially switch from FinFET to GAA in 2022. TSMC and Intel plan this switch in 2 or so years.
But Samsung Elec has problems which cost them NVidia and Qualcomm moving to TSMC. For example Snapdragon 8 Gen 1 was produced by Samsung while Snapdragon 8+ Gen 1 by TSMC and it is some 10% faster and runs quite cooler (lower power consumption).
There are several articles how Samsung Elec has lots of internal problems with management what is visible in above mentioned problem.

State of the Art

TSMC

TSMC

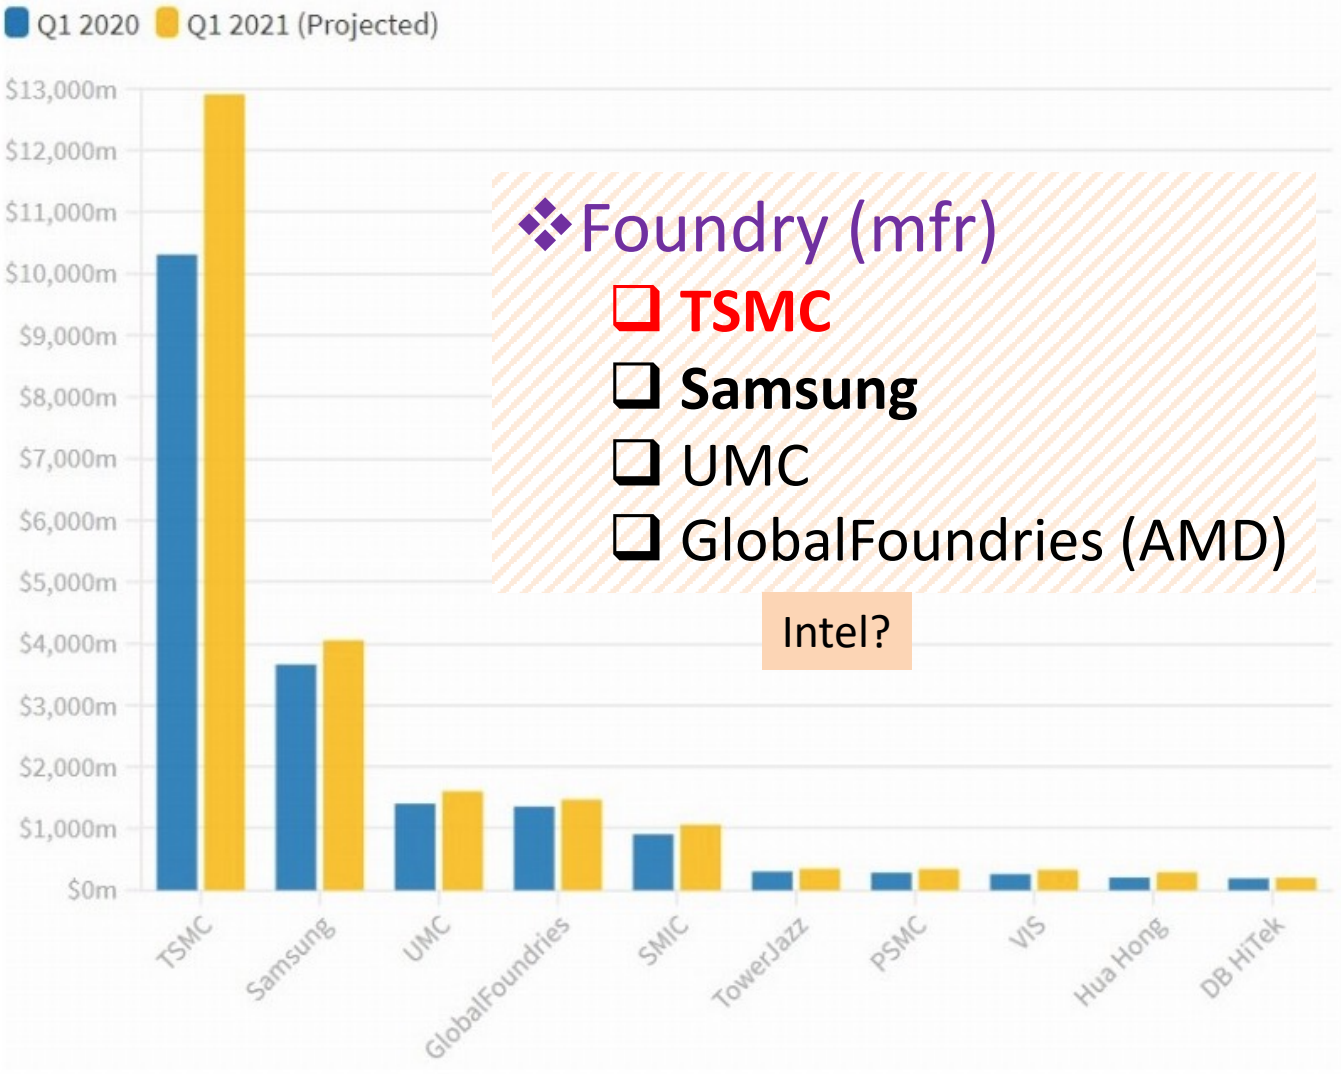

- Foundry (mfr)
  - **TSMC**
  - **Samsung**
  - UMC
  - GlobalFoundries (AMD)

Intel?

Apple

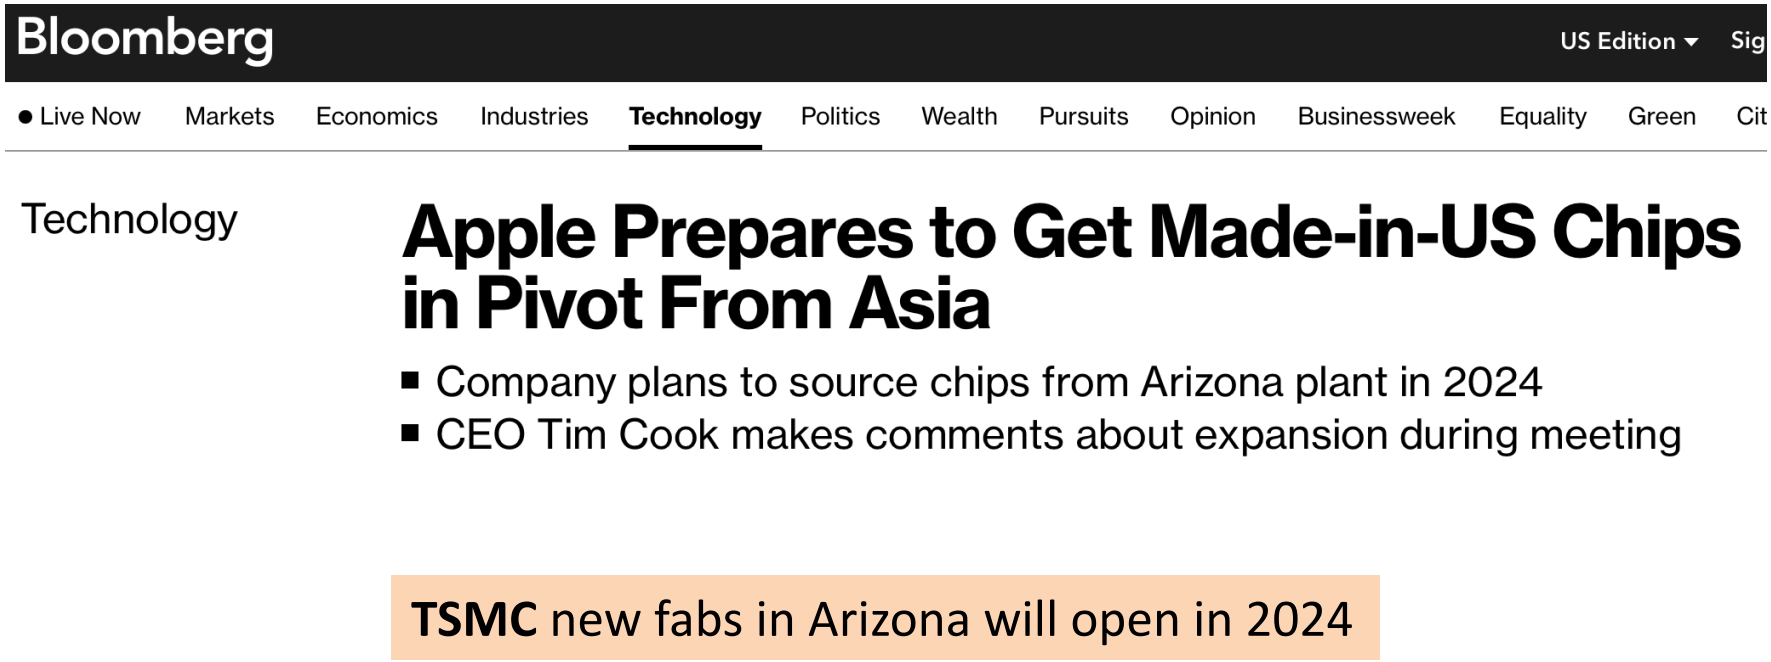

Bloomberg

US Edition ▾ Sig

● Live Now Markets Economics Industries **Technology** Politics Wealth Pursuits Opinion Businessweek Equality Green Cit

Technology

Apple Prepares to Get Made-in-US Chips in Pivot From Asia

- Company plans to source chips from Arizona plant in 2024
- CEO Tim Cook makes comments about expansion during meeting

TSMC new fabs in Arizona will open in 2024

TSMC in US

12-6-22

Apple, Nvidia CEOs to join President Biden at Taiwan Semi chip plant in Arizona

Dec 06, 2022 10:43 AM ET | **Apple Inc. (AAPL)** | Rex

Taiwan Semi (TSM), which is the world's biggest chip foundry, said it will build a second chipmaking plant in Arizona, which will boost its investments in the state to $40B after initially saying it will put $12B into its first plant in the state.

Among those heading to Taiwan Semi's (TSM) plant near Phoenix are some of the heavyweight of the semiconductor market, including Apple (NASDAQ:AAPL) Chief Executive Tim Cook, Nvidia (NASDAQ:NVDA) CEO and founder Jensen Huang and Micron Technology (NASDAQ:MU) CEO Sanjay Mehrotra. The three will be joined by Taiwan Semi (TSM) founder Morris Chang.

The White House said President Biden and the tech executives will hold a "tool-in" ceremony that involves a ceremonial moving of equipment to the production area of the Taiwan Semi (TSM) plant/

TSMC 3Q22

- Revenue = $20.23B
- EPS = $1.79
- Wafer revenue shares
  - 5nm = 28%
  - 7nm = 26%

Samsung **7nm** did pretty good but Samsung **5nm** and **4nm** had serious PDK/yield problems and Samsung **3nm** is not really competitive against TSMC N3 and it requires new design considerations for **GAA**.

Foundry Shares 2Q22

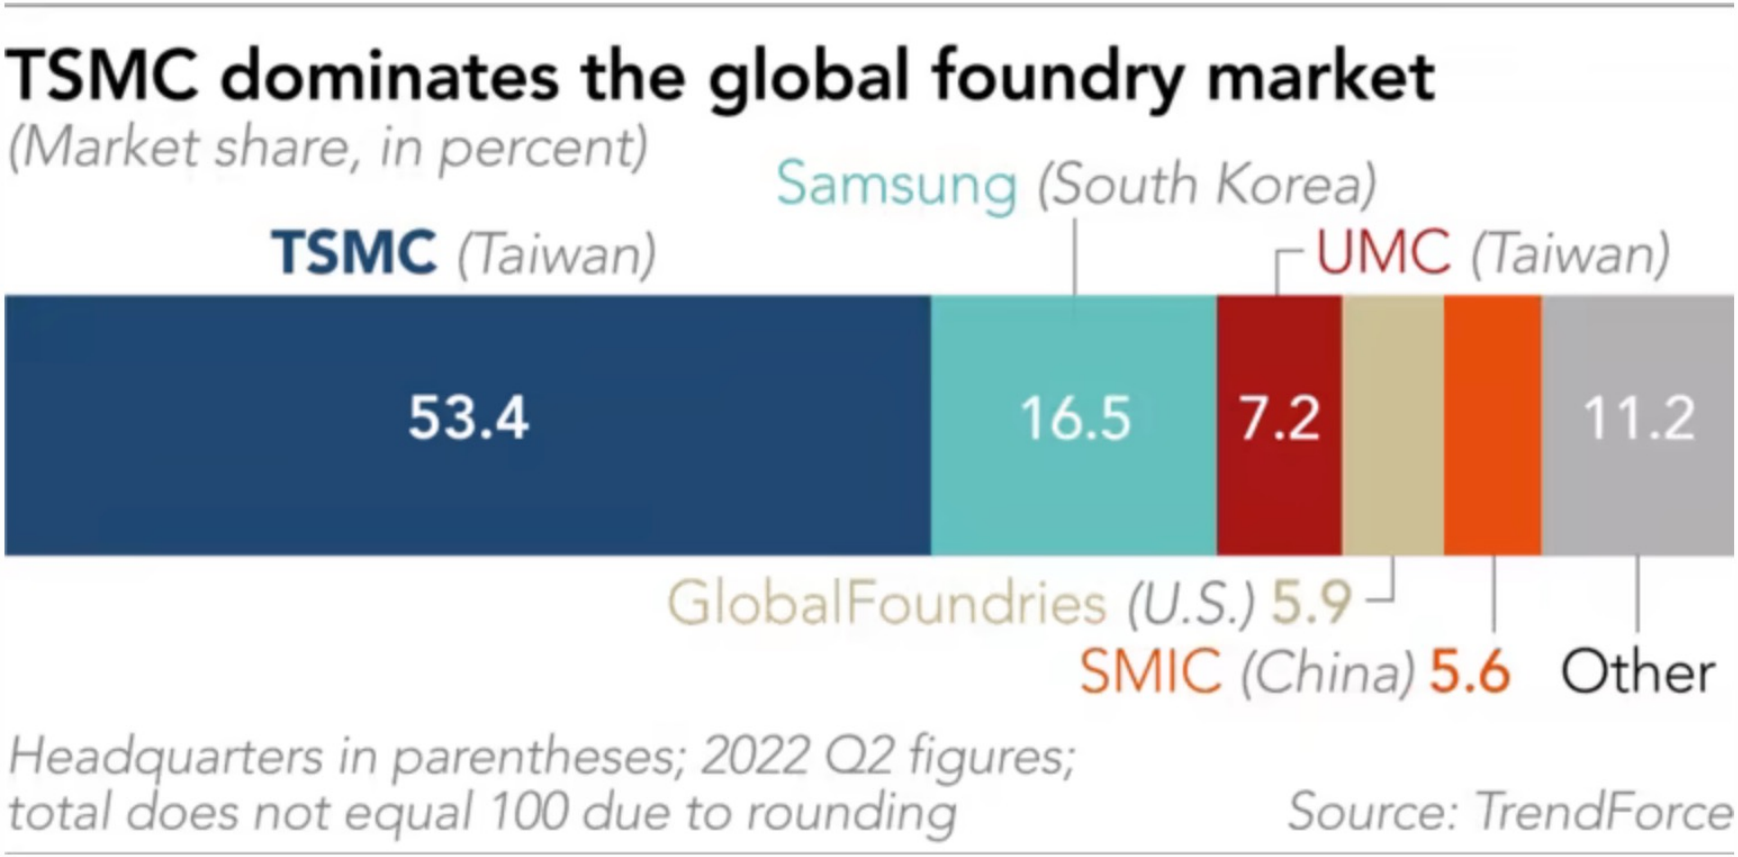

TSMC Node Segments

Semi Wiki

Revenue by Technology

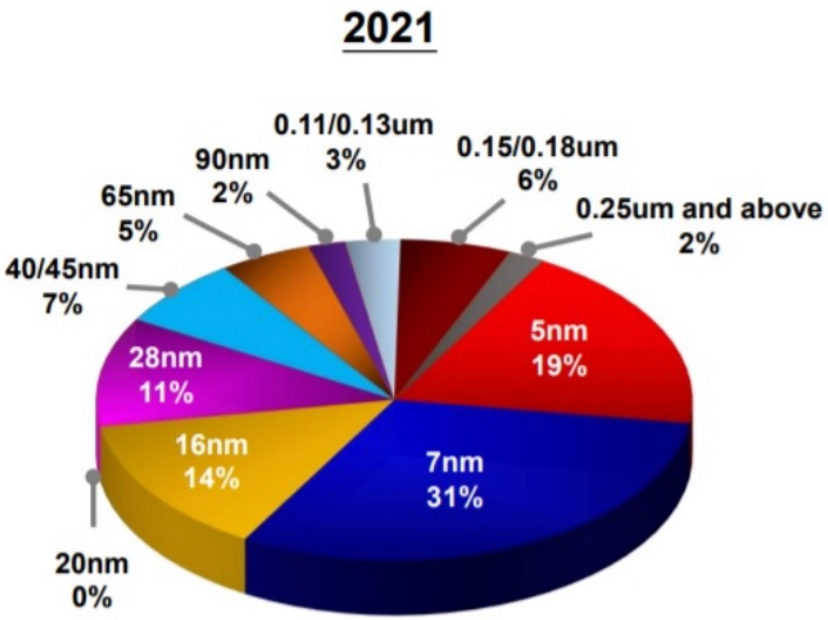

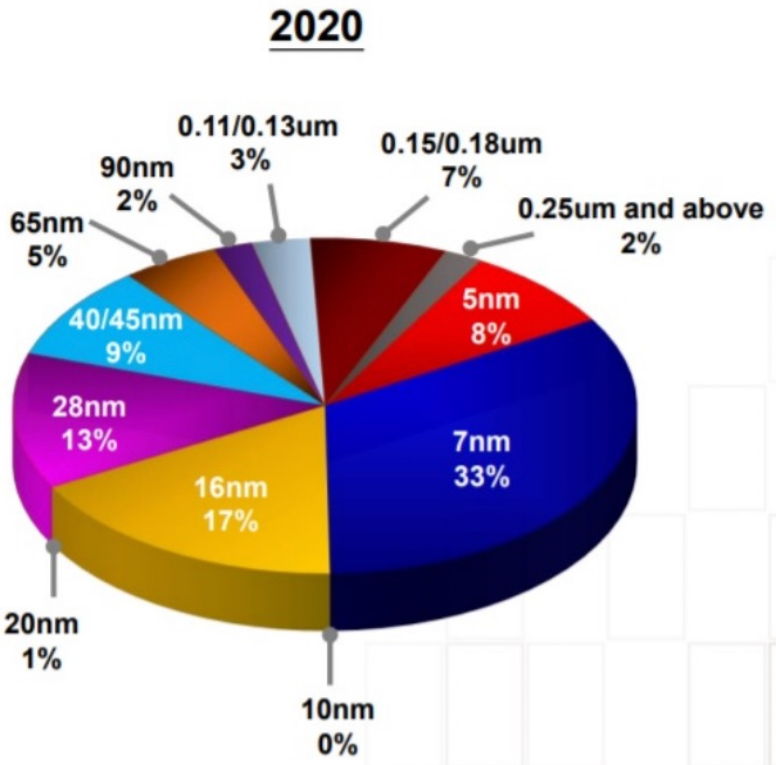

TSMC Revenue Segments

Semi Wiki

2021 Revenue by Platform

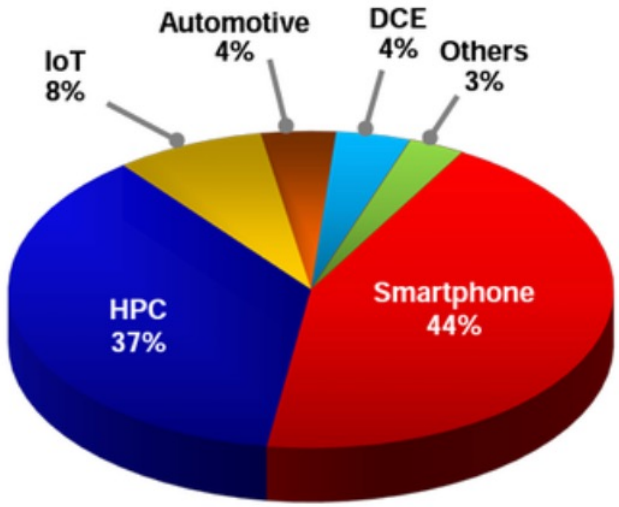

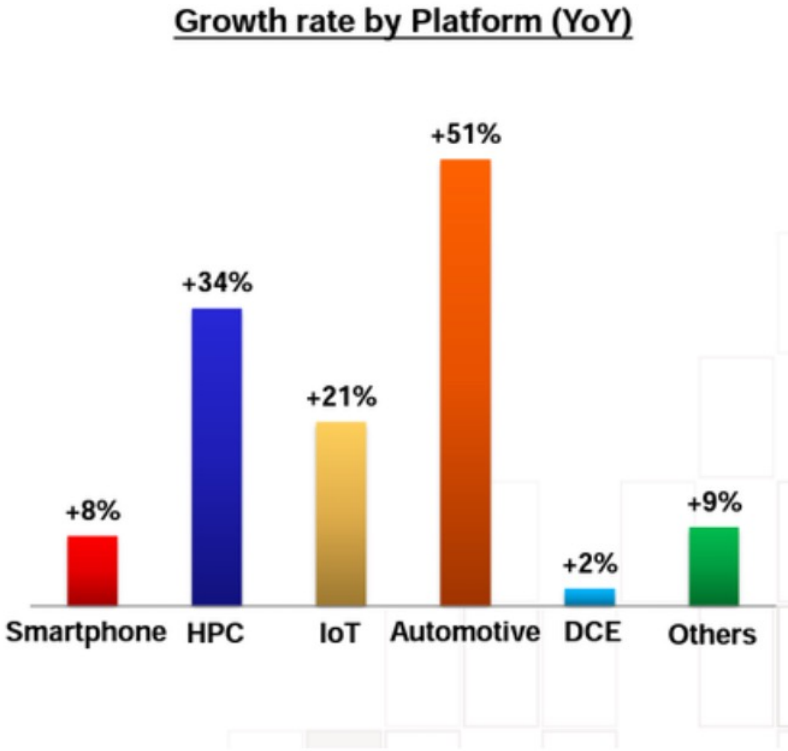

TSMC's current business mix

TSMC vs Intel Nodes

Intel Versus TSMC Nodes and Timing

	2017	2018	2019	2020	2021	2022	2023	2024	2025	2026	2027
TSMC											
- Node	7	7/7+	5	5	3	3	2	2	1.5	1.5	
- Status	Risk	Full	Risk	Full	Risk	Full	Risk	Full	Risk	Full	
Intel											
- Node	14	14	10	10SF	10SF	7	7	5	5	3	3
- EN	13.8	13.8	7.1	7.1	7.1	4.1	4.1	2.4	2.4	1.3	1.3
- Status	Full	Full	Full	Full	Full	Ramp	Full	Ramp	Full	Ramp	Full

- Risk = risk starts, Ramp = production ramp, Full = full production, EN = TSMC equivalent node.
- TSMC is assumed to stay on a new node every two years with shrink similar to the 5nm and 3nm announced shrinks.
- Intel 5nm and 3nm are assumed to be on two year intervals and to be 2x density improvements consistent with the announced 7nm density shrink.

IC KNOWLEDGE LLC

Intel Process Nodes

Slower Node Transitions Versus Foundries

IC KNOWLEDGE LLC

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

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

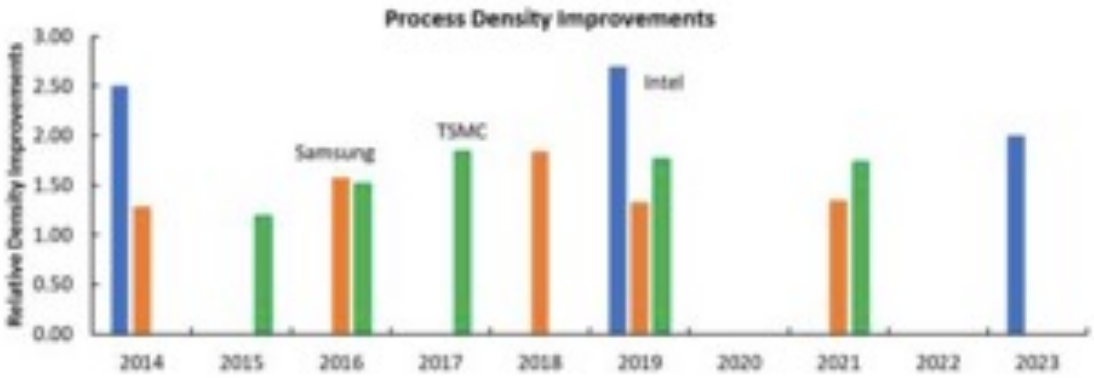

3/24/2021 4

Figure 4. Node Introductions.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

TSMC 3nm (N3)

TECHSPOT

TSMC hoping to move 3nm production to US, now starting work towards 1nm

Just in time for the next set of iPhones

TECHSPOT

TSMC 3nm (N3)

Currently, TSMC only produces 3nm nodes within its home country of Taiwan. While this doesn't necessarily cause any significant issues or hang-ups during Apple's development process, there are ways to streamline it. The two companies have one idea: move all of TSMC's 3nm production to the United States.

TSMC 3nm (N3)

TECHSPOT

In 2020, TSMC began plans to build a processing and development plant within the US. Initial estimates put the construction's completion in 2021. However, after two delays, the deadline is now Q1 2023. Assuming they meet this timeframe, iPhone 15 models will feature a brand new 3nm processor made in the United States.

TECHSPOT

TSMC 3nm (N3)

The big picture: It's been an eventful few months in the processing industry, with tech companies releasing new products left and right. The influx of devices isn't stopping, as TSMC has plans to relocate 3nm development to the United States, and that's not the only news.

Taiwan Semiconductor Manufacturing Company (TSMC) has been in the news frequently lately, with many companies releasing products that include the chip fabricator's semiconductors and process nodes. The Taiwanese chip producer has reliably supplied AMD with its 5nm processors since 2020 and even provided new 4nm nodes for Nvidia's recent Ada Lovelace graphics cards.

The iPhone 14 also features TSMC's 4nm process node. However, as its largest customer, Apple often gets the latest and greatest developments from the supplier. As such, Cupertino plans to move to TSMC's 3nm process for the iPhone 15.

TECHSPOT

TSMC 1nm (N1)

To this end, TSMC partnered with the Massachusetts Institute of Technology (MIT) and the National Taiwan University (NTU) to research and develop new methods. After loads of engineering and testing, they discovered that combining "2D materials" and "semi-metallic bismuth" results in extremely low resistance, which may overcome the most challenging aspect of producing 1nm nodes.

The research teams did confirm that 1nm nodes are still years away from being produced and sold in consumer products. We will not see current plans for 2nm nodes until late 2024. So it could easily be at least five years for such a breakthrough. Who knows, maybe we will see picometer-sized (1000 pm = 1nm) nodes before the end of the decade if these breakthroughs continue.

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

SemiWiki
The open forum for semiconductor professionals

TSMC Update

TSMC 2022 OIP Ecosystem Forum
10/26 North America 11/8 Europe 12/6 China

Nov 10, 2022

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Starting with N3 there's something new called FinFlex that used Design Technology Co-Optimization (DTCO), promising an improved Power, Performance and Area (PPA) for segments like energy-efficient and high-performance. With the FinFlex approach a designer can choose from three transistor configurations, based on their design goals:

TSMC Update

TSMC 2022 OIP Ecosystem Forum

Nov 10, 2022

There are eight choices of packaging in 3DFabric:

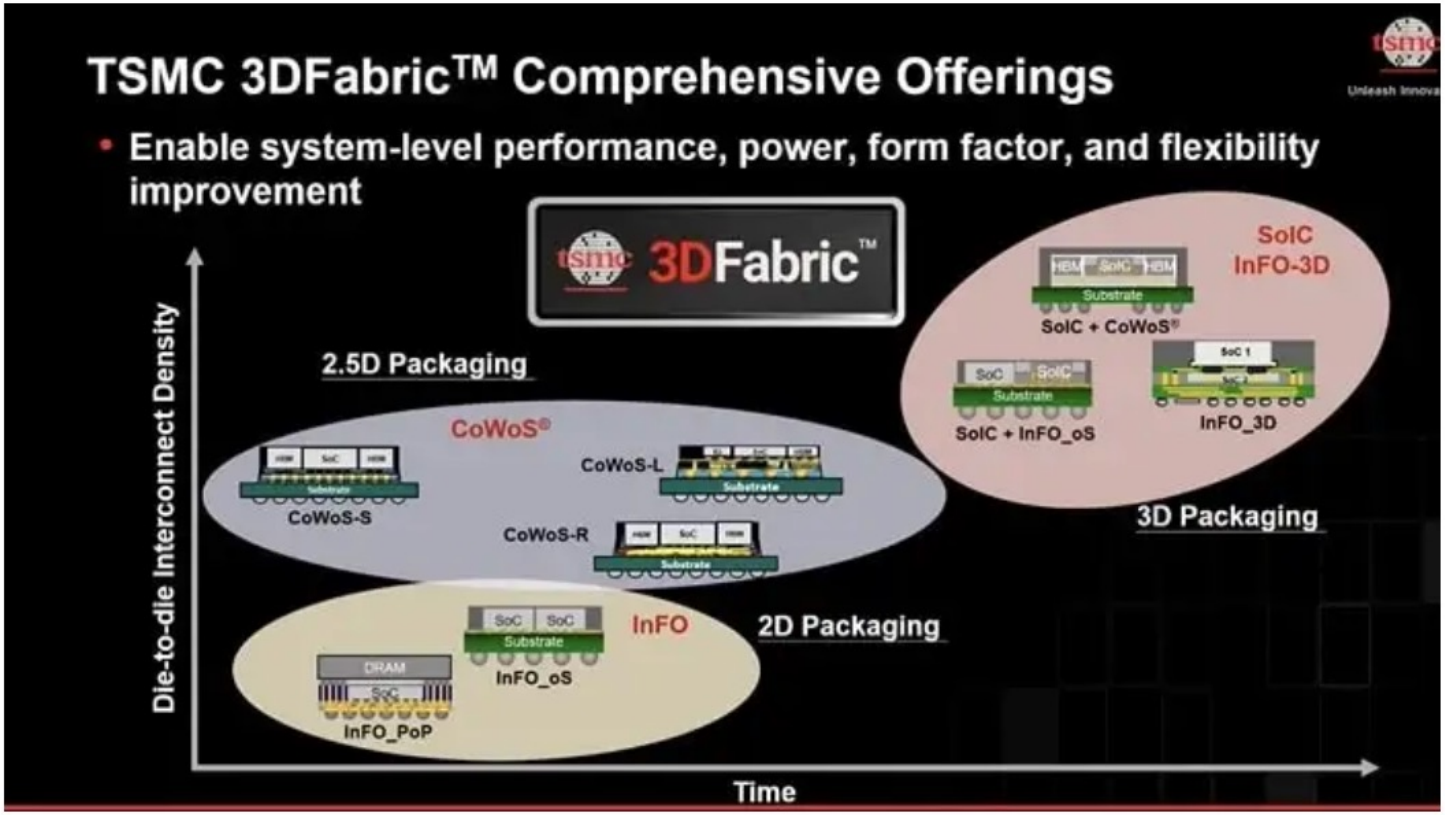

A recent example using SoIC packaging was the AMD EPYC Processor, a data center CPU, which showed a 200X interconnect density improvement over 2D packaging, a 15X density improvement over traditional 3D stacking, producing a 50-80% better CPU performance.

TSMC Update

SemiWiki
The open forum for semiconductor professionals

TSMC 2022 OIP Ecosystem Forum
10/26 North America 11/8 Europe 12/6 China

Nov 10, 2022

- 3-2 fin blocks, for high-performance
- 2-2 fin, for efficient performance
- 2-1 fin, for lowest-power, best density

3D Gate designs

FinFlex

The history of fin block choices used in process nodes N16 to N3 are shown below:

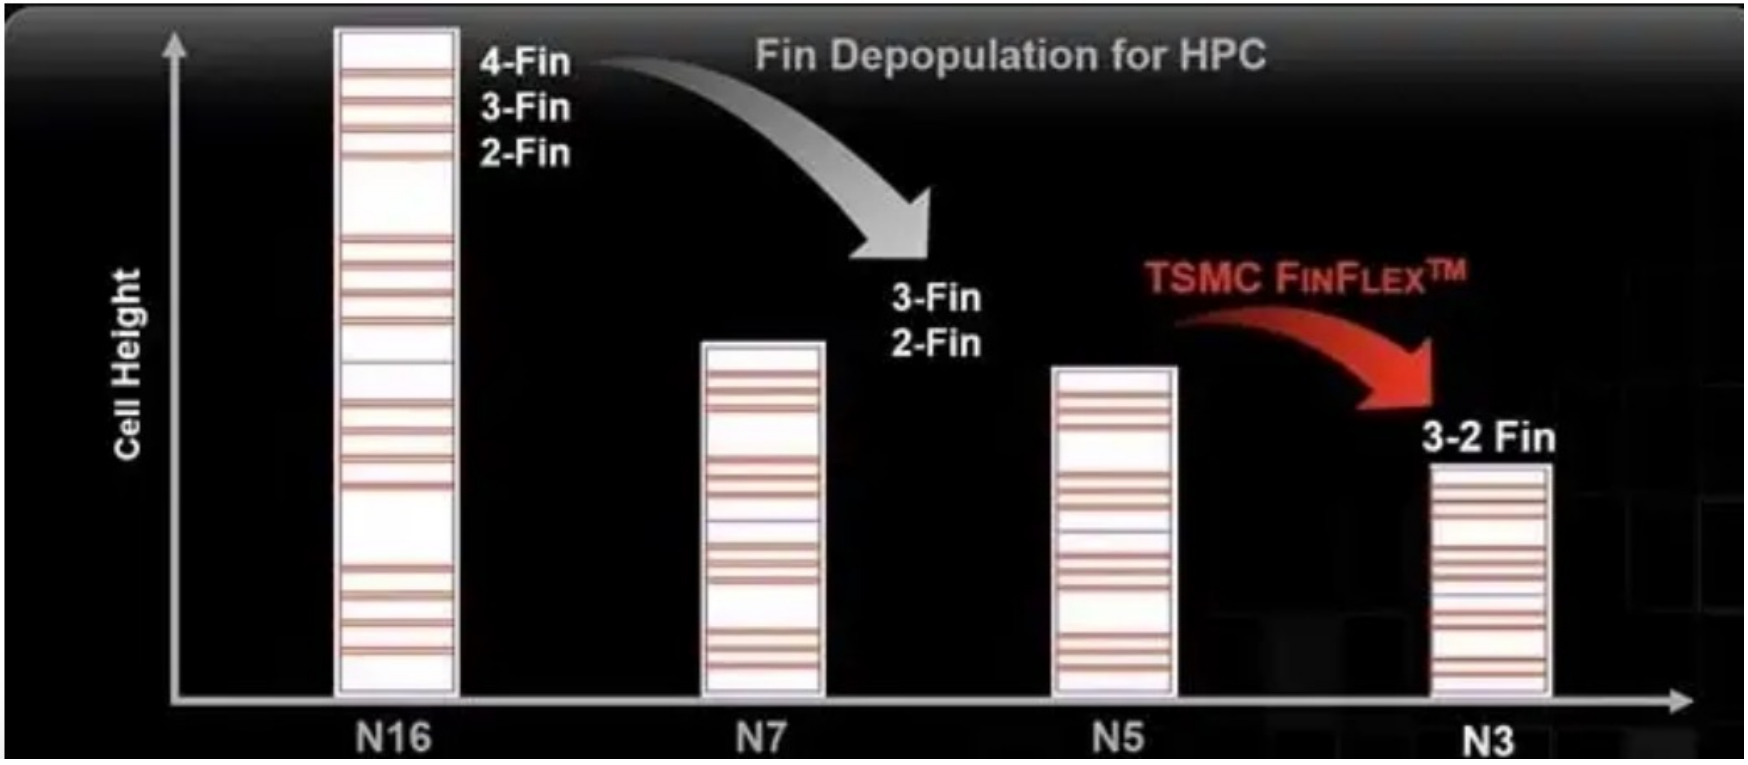

CSUN CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

SemiWiki The open forum for semiconductor professionals

TSMC Update

TSMC 2022 OIP Ecosystem Forum
10/26 North America 11/8 Europe 12/6 China

Nov 10, 2022

N3E Foundation IPs are Ready for Design Start

IP Category	IP List	Mobile	High Performance Computing	Vendors
Foundation Library/IP	Standard Cell	🟢	🟢	TSMC
	GPIO/ESD	🟢	🟢	TSMC, Synopsys
	PLL	🟡	🟡	SiCr, ABI, Synopsys
	SRAM Compiler	🟢	🟢	TSMC
	ROM Compiler	🟢	🟢	TSMC
Non Volatile Memory	Electrical Fuse	🟡	🟡	TSMC
	OTP	🔵	🔵	TSMC, Synopsys, eMemory
Interface IP	DDR4/5		🟡	Synopsys, Cadence
	LPDDR4/4X/5/5X	🟡		Synopsys, Cadence
	GDDR6/7		🔵	Cadence
	PCIe G2/3/4/5/6	🔵	🟡	Synopsys, Cadence, Alphawave
	MIPI G2/3	🟡		Synopsys
	HBM		🔵	Synopsys, GUC
	PAM4 SerDes		🟡	Synopsys, Cadence, AlphaWave
	eUSB/USB 2/3.x/4	🟡	🟡	Synopsys
Chiplet	HBI, XSR, BoW		🔵	Synopsys, Cadence, Alphawave
	UCIe		🔵	Synopsys, Cadence, Alphawave, Credo

🟢 Silicon Report Ready 🟡 Pre-Silicon DK Ready 🔵 In Development/ Planning

State of the Art

Samsung

Samsung Processes

Logic

IC Knowledge

Samsung Keynote at IEDM
by Scotten Jones on 01-27-2022 at 6:00 am

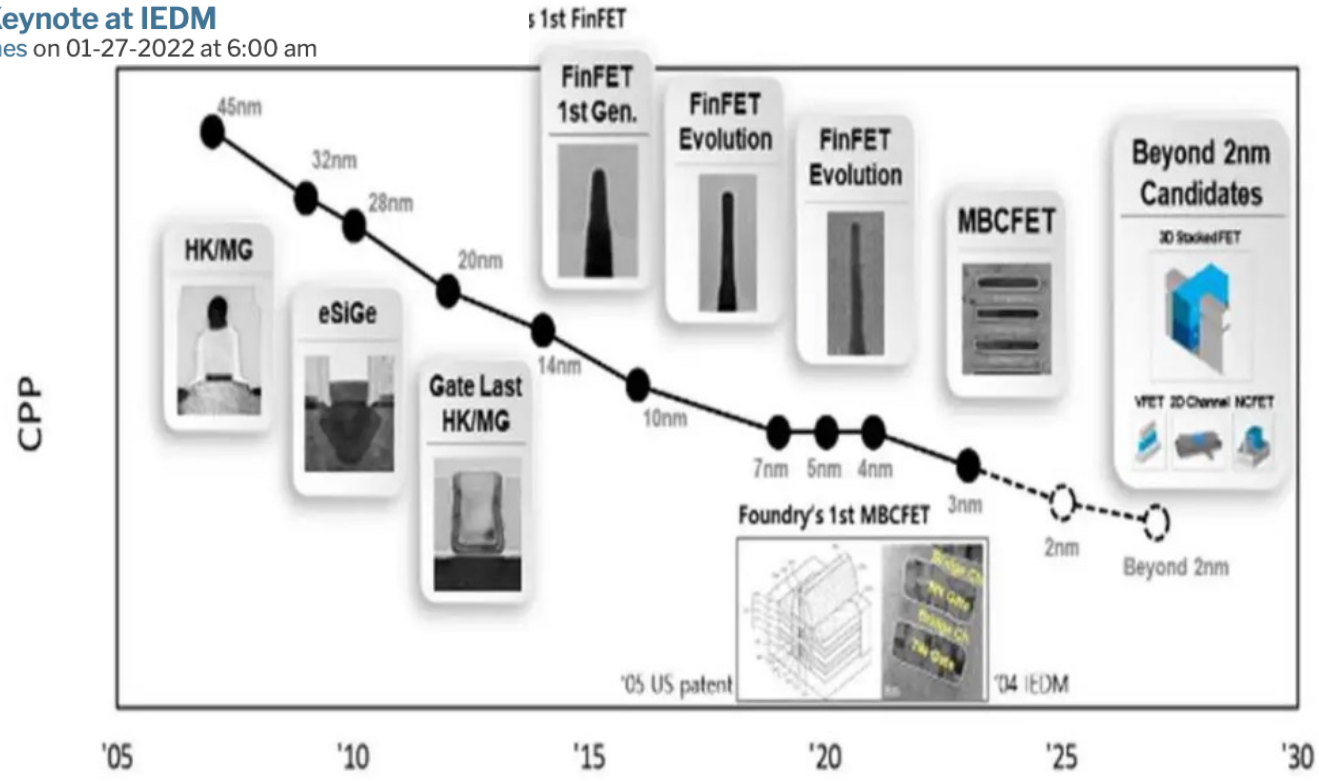

Figure 1. Logic Roadmap.

n figure 1 we can see how the contacted poly pitch (CPP) of logic processes has caled over time. In the planar era we saw high-k metal gate (HKMG) introduced by ntel at 45nm and by the foundries at 28nm as well as innovations like embedded

Samsung Processes

IC Knowledge

DRAM

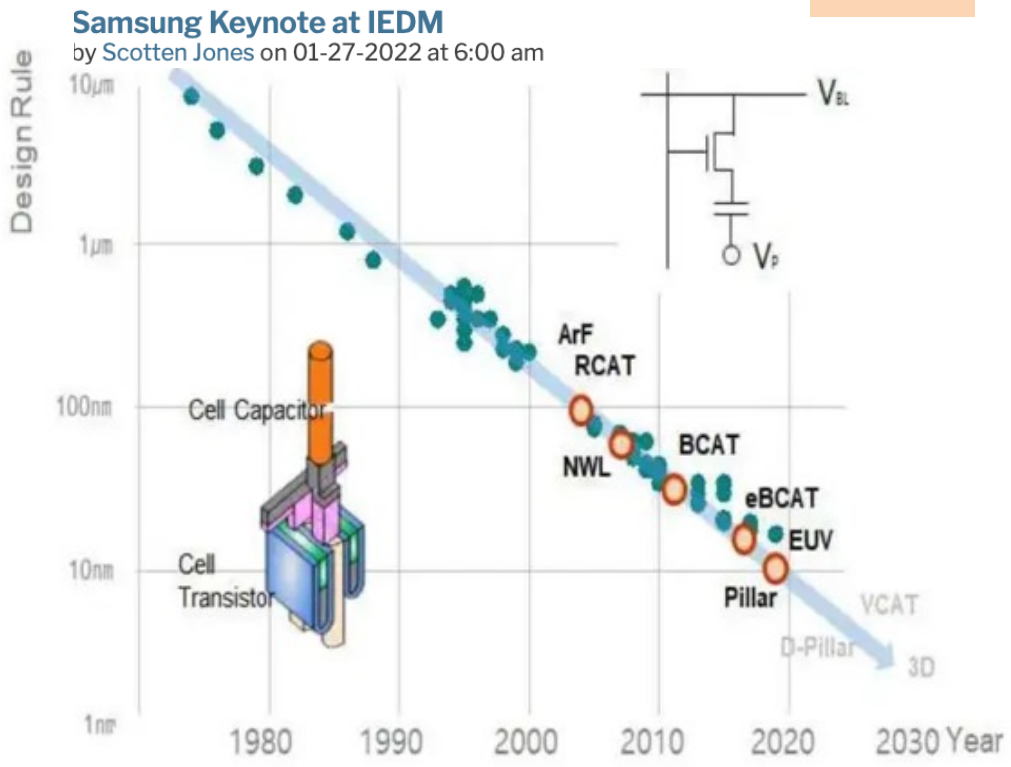

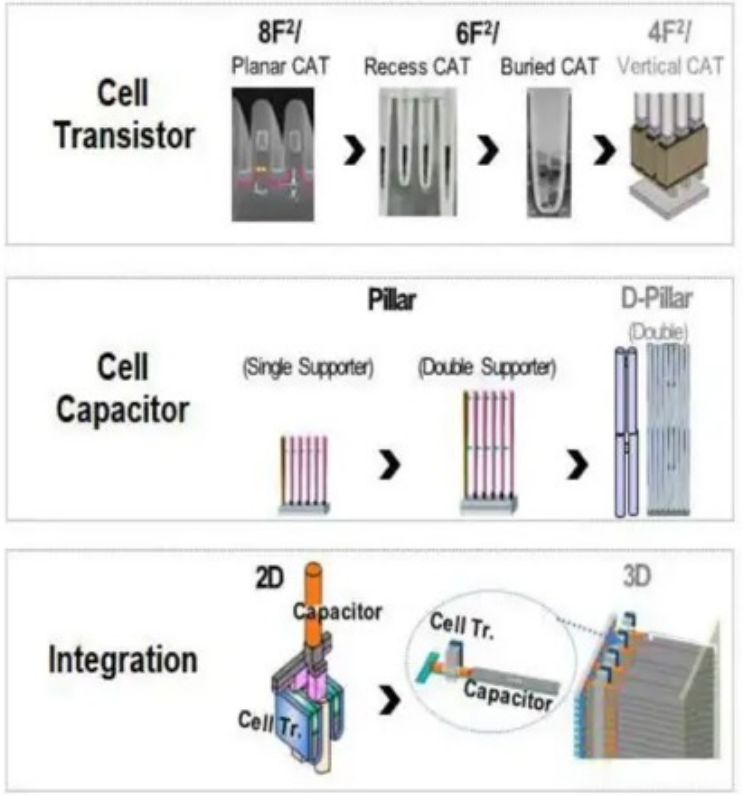

Figure 2 DRAM Roadmap

With EUV already ramping up in DRAM, the next challenges are shrinking the memory cell. Samsung is anticipating staking two layers of capacitors soon. A switch

State of the Art

Memory

SRAM 6T Cell

Quora

Drazen Zoric · Follow

This is 6T SRAM:

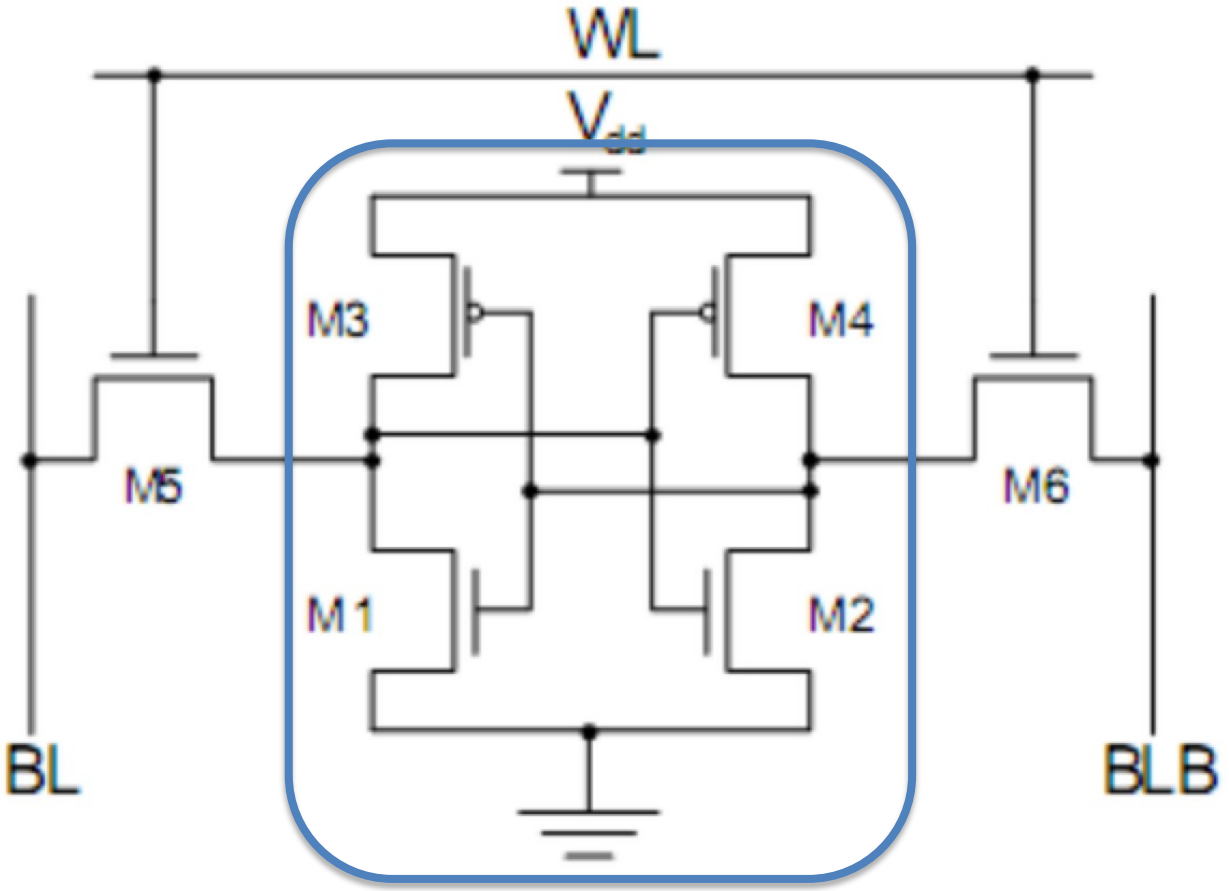

Two extra N MOSFETs are used as switches selected with WL (Word Line). Depending what is on BL (Bit Line) or BLB (Bit Line Bar or /BL) cell works in read or write mode (WL must be high to get cell out of hold mode). Driving BLs with stronger driver will write bit into cell when WL is high - M1 - M4 transistors are not strong, have higher resistance and BL drivers can properly change voltage.

SRAM Cells

I am guessing it's either 0.25 / 0.35 / 0.5 micron.

SRAM cell width/height seems regular enough that I could figure out critical dimensions.

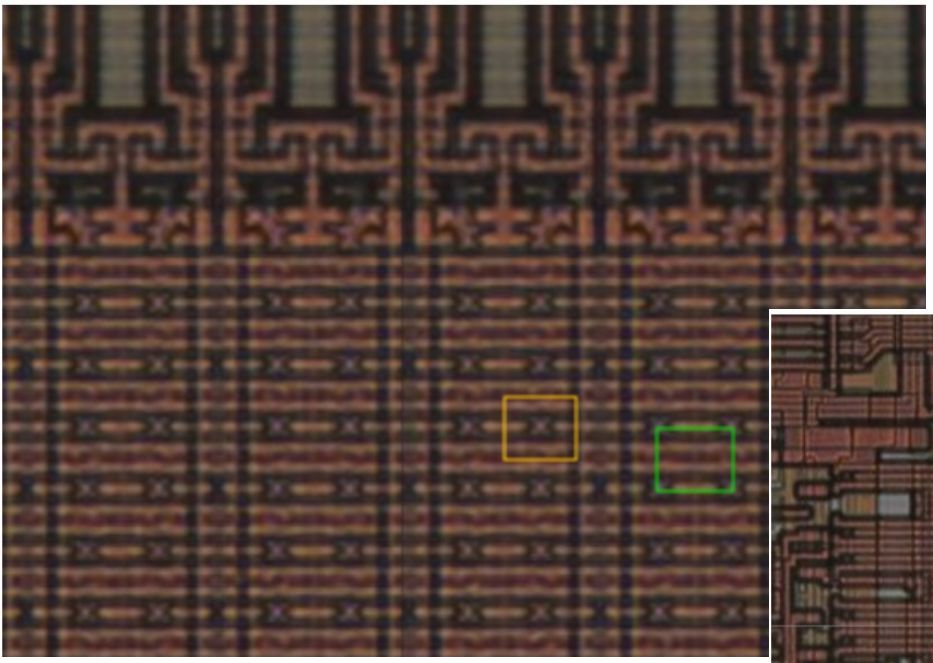

SRAM or ROM?

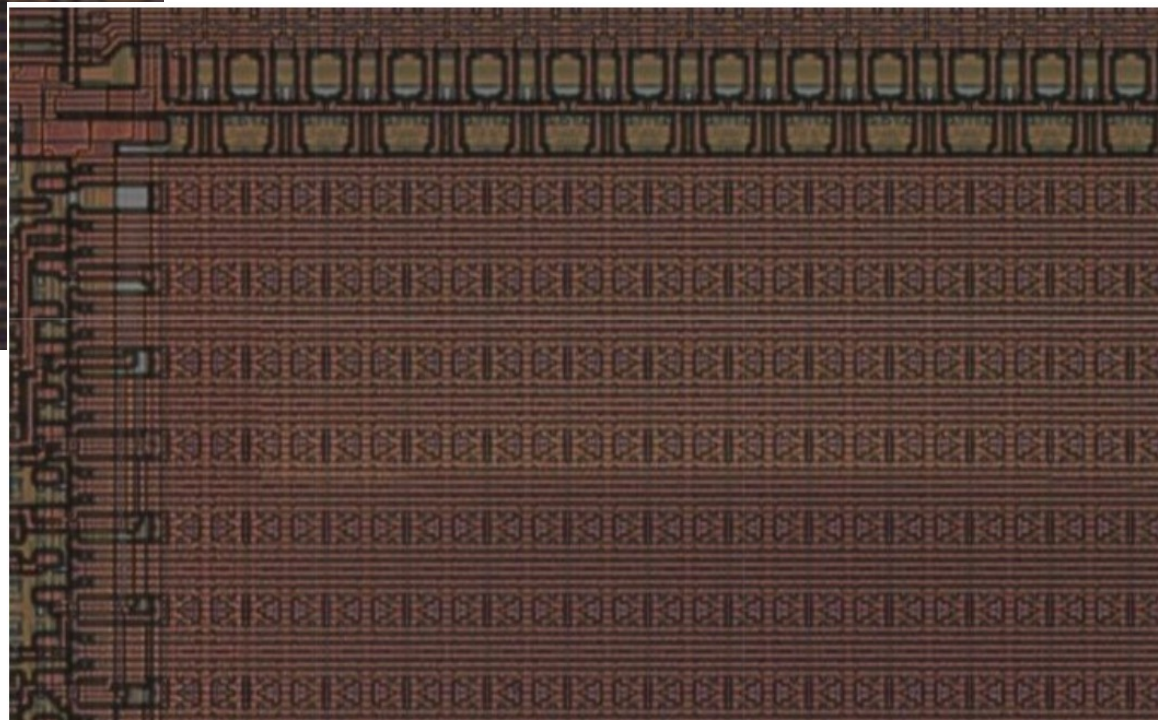

Am1101A 256x1 SRAM

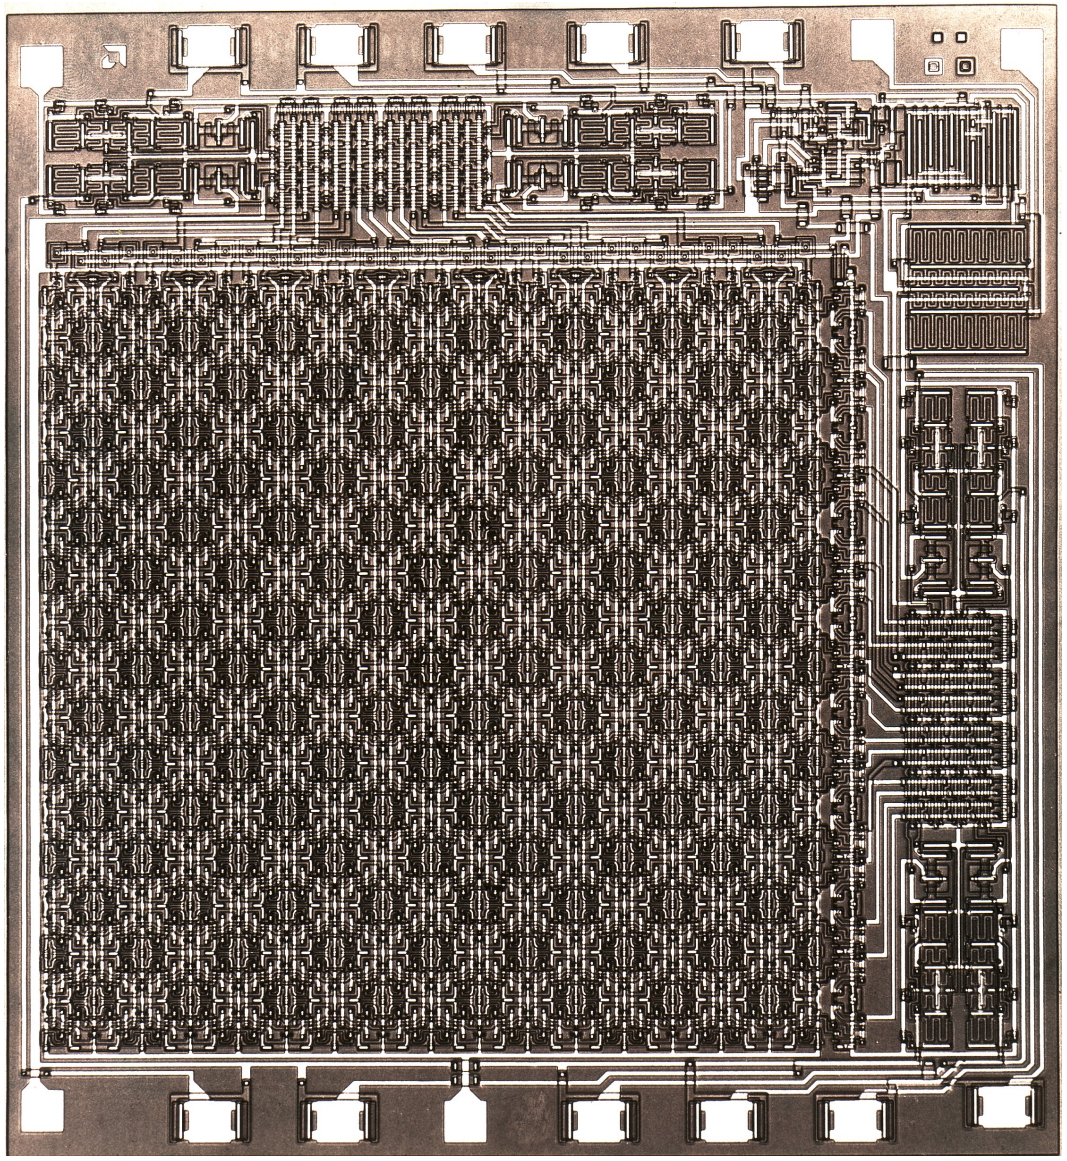

DRAM (AMD 1Kx1)

Am1103A

1971

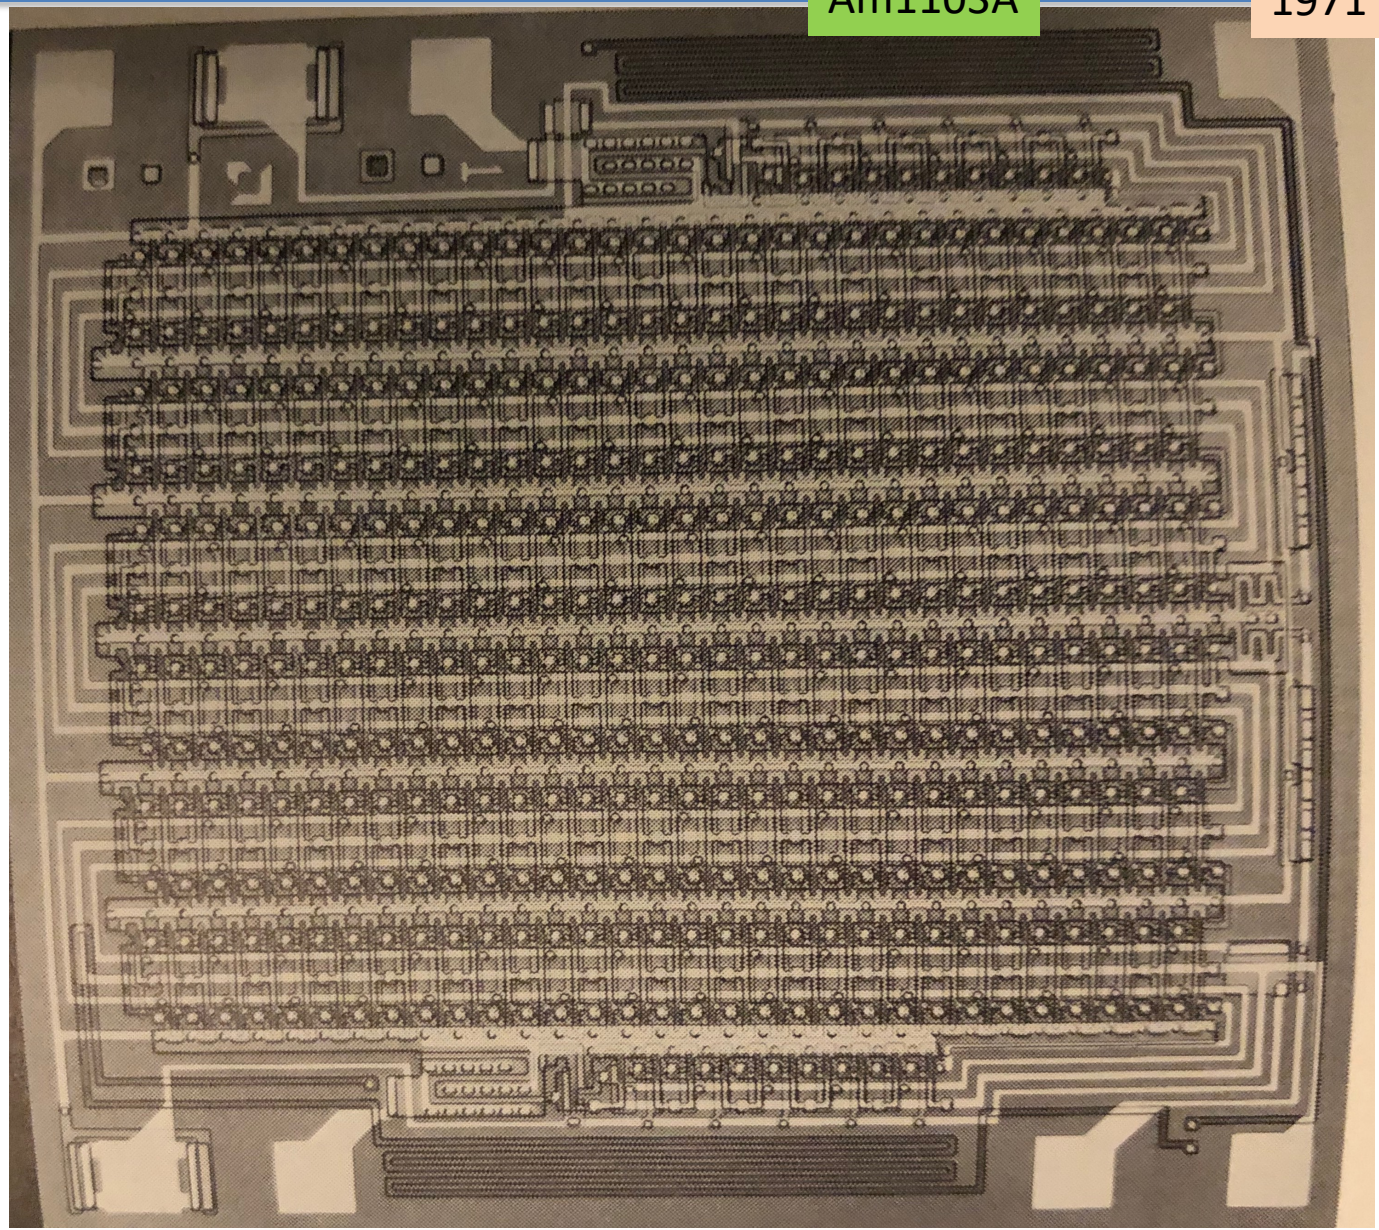

Sequential Logic: Flip-Flops

[from Wikipedia]

Clock	D	Q_{next}
Rising edge	0	0
Rising edge	1	1
Non-Rising	X	Q

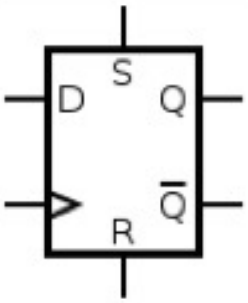

D flip-flop symbol

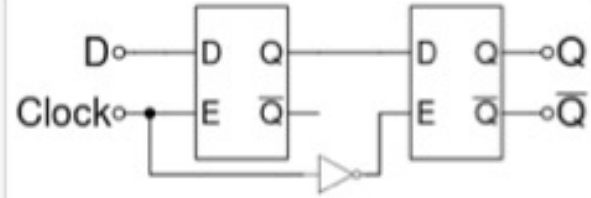

A master–slave D flip-flop. It responds on the falling edge of the *enable* input (usually a clock)

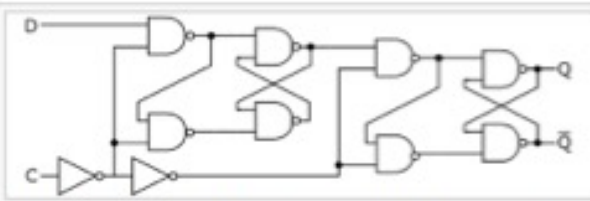

An implementation of a master–slave D flip-flop that is triggered on the rising edge of the clock

"D" FF

Timing: setup, hold, delay

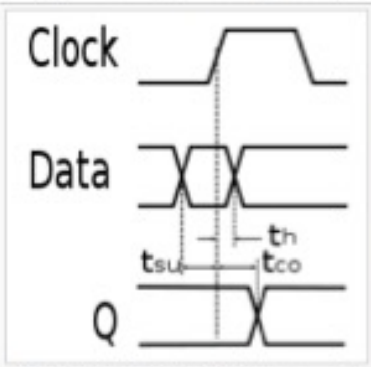

Flip-flop setup, hold and clock-to-output timing parameters

Other FFs

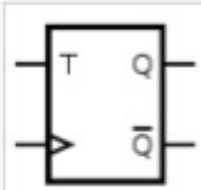

A circuit symbol for a T-type flip-flop

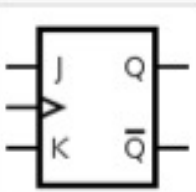

A circuit symbol for a positive-edge-triggered JK flip-flop

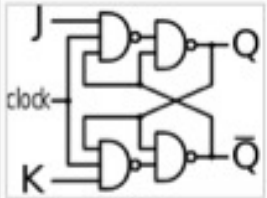

A JK flip-flop made of NAND gates

Micron Memory

Micron had better hope that Samsung shows restraint and also cuts spend so the flood of memory chips doesn't turn into a Tsunami. In the past Samsung and other aggressive memory manufacturers have cranked production in weak conditions to try to push smaller players out which is in large part why the US went from 7 memory makers to just one remaining, Micron.

CSUN CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

SemiWiki The open forum for semiconductor professionals

Micron Memory

Micron- "The first cut isn't the deepest"- Chops production & forecast further

by Robert Maire on 11-23-2022 at 6:00 am
Categories: Semiconductor Advisors, Semiconductor Services
2 Comments

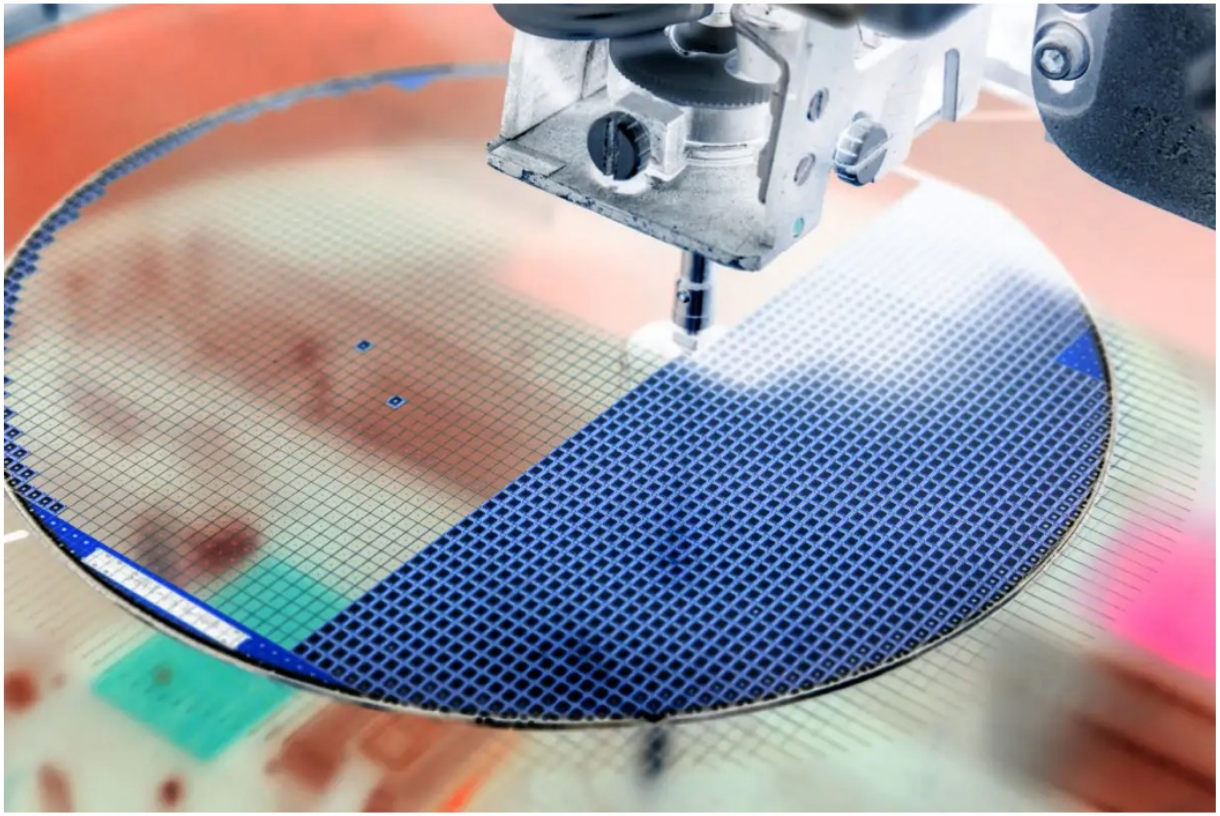

-Micron announces more production cuts & lower forecast
-DRAM will be negative- NAND sounds barely flat (for now)
-Capital spending to be cut to near zero- essential only
-Will the rest of the industry follow suit?

Virtual Memory

Quora

Why do we use virtual memory addresses instead of real physical addresses? ...

Jeff Drobman
Lecturer at California State University, Northridge (2016–present) · Just now ·

we DO use physical addresses — for “embedded” designs. for general computing with an operating system, the OS needs to be able to locate and relocate code and data anywhere in the physical address space. since a programmer cannot know where their code will be located, they use a “virtual” address which gets translated into a physical from the OS’ “page table” and its cache called “TLB”.

VM – Page

Quora

What is the definition of a page in computer science?

Jeff Drobman
Lecturer at California State University, Northridge (2016–present) · Just now

a “page” is a unit of main memory segmentation anywhere in size from 1KB to 1MB, but typically 64KiB (e.g., MIPS). there can be both *virtual* and *physical* pages in a virtual memory, with a mapping managed by the OS. in a virtual memory, the upper part of a virtual address is the page number, while the lower portion is the untranslated address within the page.

AMD64 Virtual Memory

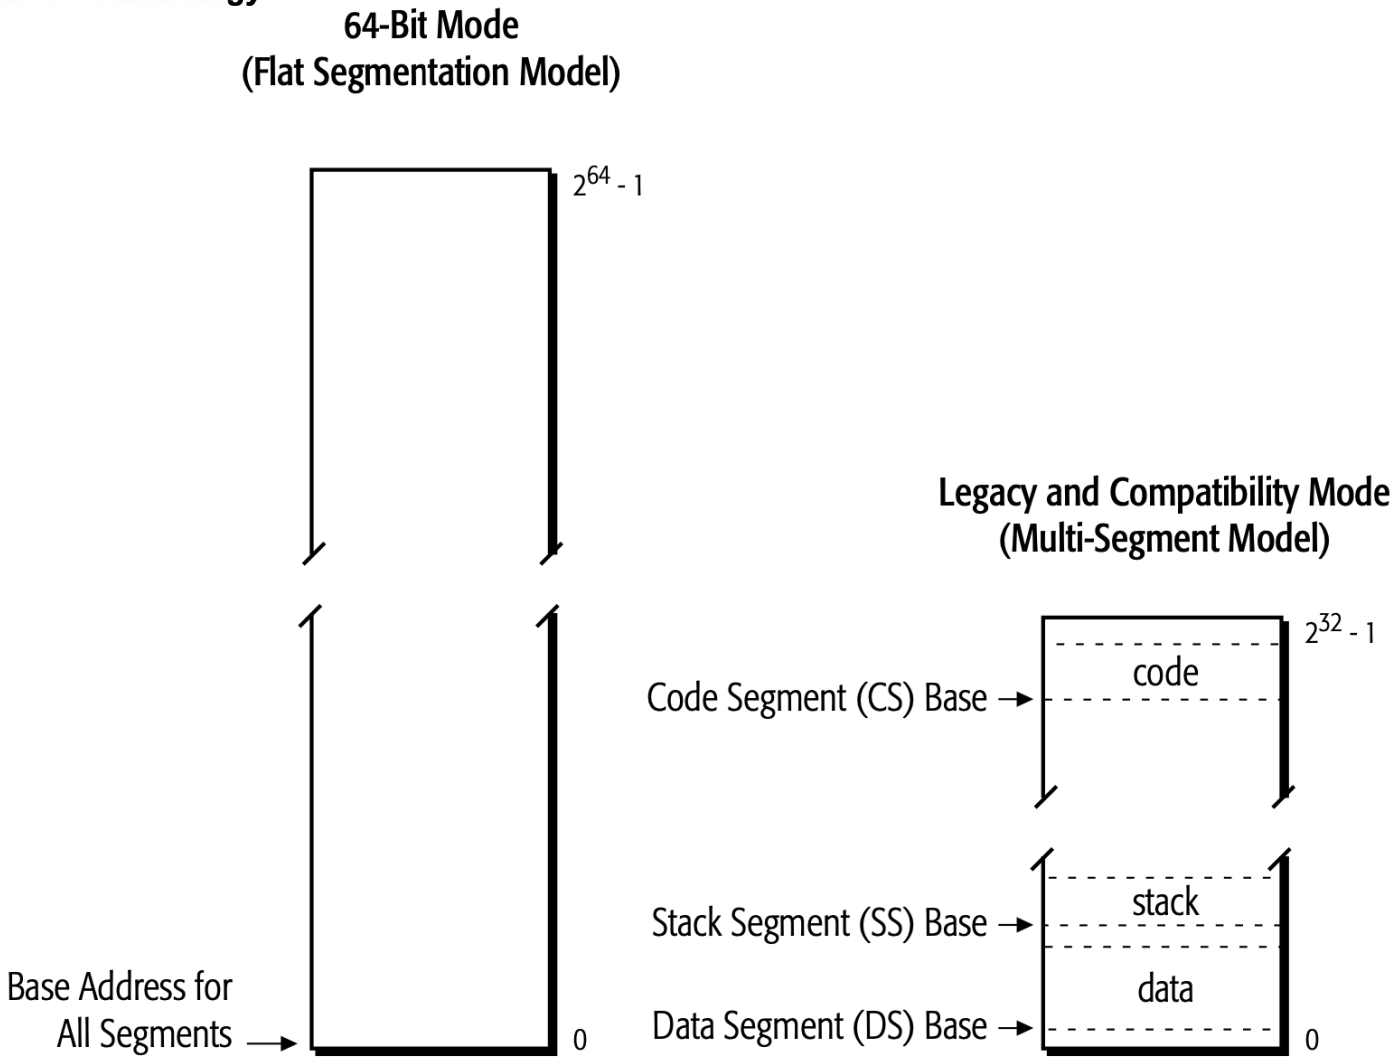

Figure 2-1. Virtual-Memory Segmentation

TLB Data

5.13 Real stuff: The ARM Cortex-A8 and Intel Core i7 — P&H 5.13

Figure 5.13.1: Address translation and TLB hardware for the ARM Cortex-A53 and Intel Core i7 920 (COD Figure 5.42).

A53

Both processors provide support for large pages, which are used for things like the operating system or mapping a frame buffer. The large-page scheme avoids using a large number of entries to map a single object that is always present.

Characteristic	ARM Cortex-A53	Intel Core i7
Virtual address	48 bits	48 bits
Physical address	40 bits	36 bits
Page size	Variable: 4, 16, 64 KiB, 1, 2 MiB, 1 GiB	Variable: 4 KiB, 2/4 MiB
TLB organizatio n	1 TLB for instructions and 1 TLB for data Both L1 TLBs are fully associative, with 10 entries, round robin replacement Unified L2 TLB with 512 entries, 4-way set associate TLB misses handled in hardware	1 TLB for instructions and 1 TLB for data per core Both L1 TLBs are four-way set associative, LRU replacement L1 I-TLB has 128 entries for small pages, 7 per thread for large pages L1 D-TLB has 64 entries for small pages, 32 for large pages The L2 TLB is four-way set associative, LRU replacement The L2 TLB has 512 entries TLB misses handled in hardware

State of the Art

Micro Controllers

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE
COMP222

AVR 8-bit MCU

AVR microcontrollers

From Wikipedia, the free encyclopedia

AVR is a family of microcontrollers developed since 1996 by Atmel, acquired by Microchip Technology in 2016. These are modified Harvard architecture 8-bit RISC single-chip microcontrollers. AVR was one of the first microcontroller families to use on-chip flash memory for program storage, as opposed to one-time programmable ROM, EPROM, or EEPROM used by other microcontrollers at the time.
AVR microcontrollers find many applications as embedded systems. They are especially common in hobbyist and educational embedded applications, popularized by their inclusion in many of the Arduino line of open hardware development boards.

XMEGA

Flash size	Frequency [MHz]	Package	SRAM	EEPROM	Release year
16–256 KB	32	44–100-pin package	1–32 KB	512–2048 bytes	—

the ATxmega series offers a wide variety of peripherals and functionality such as:

- Extended performance features, such as DMA, "Event System", and cryptography support
- Extensive peripheral set with ADCs

ADC

Application-specific AVR

- megaAVRs with special features not found on the other members of the AVR family, such a

FPSLIC (AVR with FPGA)

FPGA

- FPGA 5k to 40k gates
- SRAM for the AVR program code, unlike all other AVRs
- AVR core can run at up to 50 MHz[9]

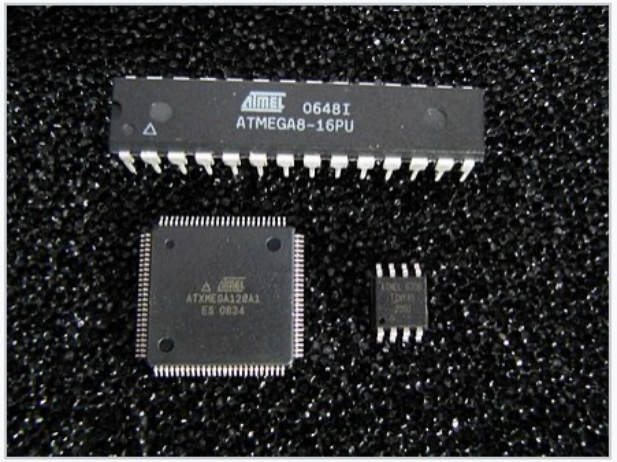

Various older AVR microcontrollers: ATmega8 in 28-pin narrow dual in-line package (DIP-28N), ATxmega128A1 in 100-pin thin quad flat pack (TQFP-100) package, ATtiny45 in 8-pin small outline (SO-8) package.

AVR 32-bit MCU

AVR microcontrollers

From Wikipedia, the free encyclopedia

•In 2006, Atmel released microcontrollers based on the 32-bit AVR32 architecture.

This was a completely different architecture unrelated to the 8-bit AVR, intended to compete with the ARM-based processors. It had a 32-bit data path, SIMD and DSP instructions, along with other audio- and video-processing features. The instruction set was similar to other RISC cores, but it was not compatible with the original AVR (nor any of the various ARM cores). Since then support for AVR32 has been dropped from Linux as of kernel 4.12; compiler support for the architecture in GCC was never mainlined into the compiler's central source-code repository and was available primarily in a vendor-supported fork. At the time that AVR32 was introduced, Atmel had already been a licensee of the ARM architecture, with both ARM7 and ARM9 microcontrollers having been released prior to and concurrently with the AVR32; later Atmel focused most development effort on 32-bit chips with ARM Cortex-M and Cortex-A cores.

SIMD/DSP

ARM

- Analog comparator
- 10 or 12-bit A/D converters, with multiplex of up to 16 channels
- 12-bit D/A converters

Maxim/ADI MCU

maxim integrated | NOW PART OF ANALOG DEVICES

Microcontrollers

100 MHz

Part Number ↑↓ Filter by part number	Microcontroller Type	MCU Core ↑↓	Internal Flash (KBytes) ↑↓	Speed (max) (MHz) ↑↓	Data Processing ↑↓	Internal SRAM (KBytes) ↑↓
MAX32672 High-Reliability, Tiny, Ultra-Low Power ARM Cortex-M4F Microcontroller with 12-Bit 1MSPS ADC	General Purpose, Low Power	ARM Cortex-M4F	1000	100	32-bit	200
MAX32675 Ultra-Low-Power Arm Cortex-M4F with Precision Analog Front End for Industrial and Medical Sensors	General Purpose, Industrial, Low Power	ARM Cortex-M4F	384	100	32-bit	160
MAX32670 High Reliability, Ultra Low Power Microcontroller Powered by ARM Cortex M4 w/ FPU for Industrial and IoT	General Purpose, Low Power	ARM Cortex-M4F	384	100	32-bit	160

ARM

Maxim/ADI Clock Gen

maxim integrated | NOW PART OF ANALOG DEVICES

Part Number (Filter by part number)	End Equipment	f_{IN} (min) (MHz)	f_{IN} (max) (MHz)	f_{OUT} (min) (MHz)	f_{OUT} (max) (MHz)
MAX31180 Spread-Spectrum Crystal Multiplier	General Purpose	16	33.4	16	134
DS1080L Spread-Spectrum Crystal Multiplier	General Purpose	16	33.4	134	134

Jitter

Output Jitter (RMS) (ps)	V_{SUPPLY} (V)
75	3.3
75	3.3

CSUN CALIFORNIA STATE UNIVERSITY NORTHRIDGE

Quora

MCU

DR JEFF SOFTWARE INDIE APP DEVELOPER

COMP2

Tom Crosley · Follow

Embedded systems programmer for 45 years · 4y

What is the smallest CPU currently?

Probably this little guy, ATTINY20 from Microchip:

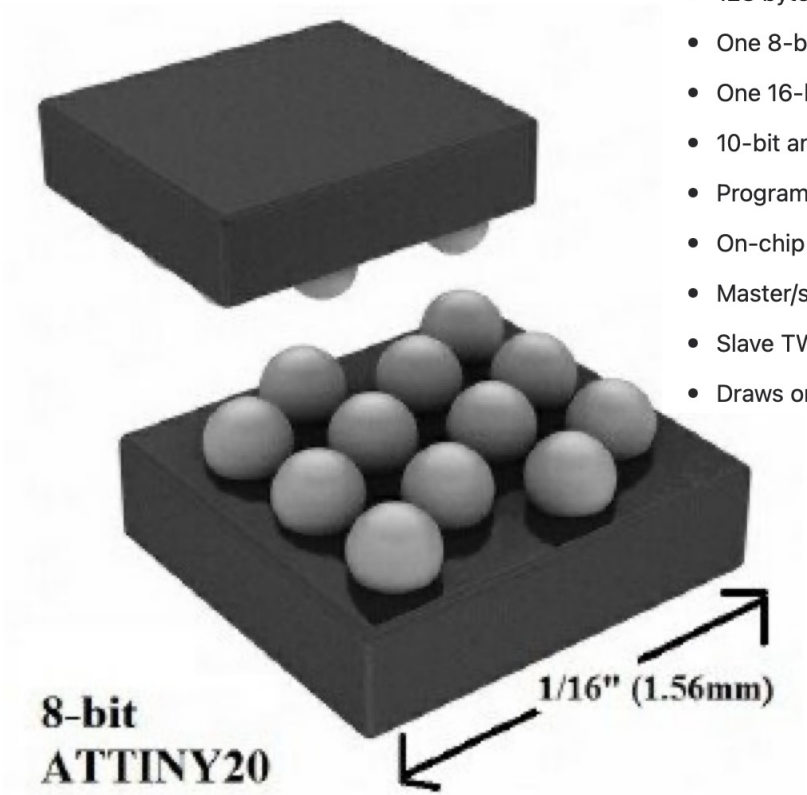

Capable microcontroller with the following features:

- 16 x 8 general purpose working registers
- Up to 12 MIPS throughput at 12 MHz
- 2K bytes of programmable flash program memory
- 128 bytes internal SRAM
- One 8-bit timer/counter with two PWM channels
- One 16-bit timer/counter with two PWM channels
- 10-bit analog to digital converter
- Programmable watchdog timer with separate on-chip oscillator
- On-chip analog comparator
- Master/slave SPI serial interface
- Slave TWI serial interface
- Draws only 200 µA at 1MHz and 1.8V

Less than 1/16″ square (1.40 mm x 1.56mm). Look at a ruler and see how small that really is.

State of the Art

Supercomputers

New Top Super

Quora 273 606 Search Quora

Currently, as of the time of writing this (August 2022), the world's fastest computer is the HP Frontier supercomputer, owned by Oak Ridge National Laboratory.

Frontier

- AMD EPYC
- AMD MI250

Summit

- AMD EPYC
- Nvidia A100

Supers: #1 Frontier

Quora

Dan L. Oom · Sep 13

The "Summit" is now in fourth place. The fastest computer is the Frontier, also at ORNL at 1.1 EFLOPS.

It uses AMD EPYC 64C 2GHz processors, with a total of over 8 million cores, so its a bit like a million PCs with a highspeed Slingshot-11 interconnect.

9 Reply

Drazen Zoric · Sep 14

Frontier uses only 9472 Epycs with total of 606,208 cores but also uses AMD MI250X Instinct Accelerators (GPU) with total of 8,335,360 "GPU" cores.

Top 5 Supers

Rank	System	Cores	Rmax (PFlop/s)	Rpeak (PFlop/s)	Power (kW)
1	**Frontier** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE DOE/SC/Oak Ridge National Laboratory United States	8,730,112	1,102.00	1,685.65	21,100
2	**Supercomputer Fugaku** - Supercomputer Fugaku, A64FX 48C 2.2GHz, Tofu interconnect D, Fujitsu RIKEN Center for Computational Science Japan	7,630,848	442.01	537.21	29,899
3	**LUMI** - HPE Cray EX235a, AMD Optimized 3rd Generation EPYC 64C 2GHz, AMD Instinct MI250X, Slingshot-11, HPE EuroHPC/CSC Finland	2,220,288	309.10	428.70	6,016
4	**Leonardo** - BullSequana XH2000, Xeon Platinum 8358 32C 2.6GHz, NVIDIA A100 SXM4 64 GB, Quad-rail NVIDIA HDR100 Infiniband, Atos EuroHPC/CINECA Italy	1,463,616	174.70	255.75	5,610
5	**Summit** - IBM Power System AC922, IBM POWER9 22C 3.07GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM DOE/SC/Oak Ridge National Laboratory United States	2,414,592	148.60	200.79	10,096
6	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz,	1,572,480	94.64	125.71	7,438

Top 10 Supers

6	**Sierra** - IBM Power System AC922, IBM POWER9 22C 3.1GHz, NVIDIA Volta GV100, Dual-rail Mellanox EDR Infiniband, IBM / NVIDIA / Mellanox DOE/NNSA/LLNL United States	1,572,480	94.64	125.71	7,438
7	**Sunway TaihuLight** - Sunway MPP, Sunway SW26010 260C 1.45GHz, Sunway, NRCPC National Supercomputing Center in Wuxi China	10,649,600	93.01	125.44	15,371
8	**Perlmutter** - HPE Cray EX235n, AMD EPYC 7763 64C 2.45GHz, NVIDIA A100 SXM4 40 GB, Slingshot-10, HPE DOE/SC/LBNL/NERSC United States	761,856	70.87	93.75	2,589
9	**Selene** - NVIDIA DGX A100, AMD EPYC 7742 64C 2.25GHz, NVIDIA A100, Mellanox HDR Infiniband, Nvidia NVIDIA Corporation United States	555,520	63.46	79.22	2,646
10	**Tianhe-2A** - TH-IVB-FEP Cluster, Intel Xeon E5-2692v2 12C 2.2GHz, TH Express-2, Matrix-2000, NUDT National Super Computer Center in Guangzhou China	4,981,760	61.44	100.68	18,482

CSUN
CALIFORNIA STATE UNIVERSITY NORTHRIDGE

COMP222

Supercomputer Pix

DR JEFF SOFTWARE
INDIE APP DEVELOPER

Quora

Andrew McGregor · Follow
Performance Measurement Lead at Fastly (company) (2019–present) · Upvot
Gaurav Saxena, MSc High Performance Computing, University of Edinburgh ·

What does a super computer look like?

From the outside:

Looks like a big industrial building with a lot of power and cooling… because that's what it is. Depending on how we look at it, that's one machine, 17 machines, 33 machines, or a classified number that might be a few hundred thousand (I didn't look it up).

Supercomputer Pix

Quora

Andrew McGregor · Follow
Performance Measurement Lead at Fastly (company) (2019–present) · Upvot

What does a super computer look like?

Each tray in the racks on the right is a very big PC-like server and some hard drives.

Supercomputer Pix

Quora

Andrew McGregor · Follow
Performance Measurement Lead at Fastly (company) (2019–present) · Upvot
Gaurav Saxena, MSc High Performance Computing, University of Edinburgh ·

What does a super computer look like?

Cooling. For scale, this space is triple-height, so those big green things on the left are about shoulder-high to a tall man.

Supercomputer Pix

Quora

Andrew McGregor · Follow
Performance Measurement Lead at Fastly (company) (2019–present) · Upvot
Gaurav Saxena, MSc High Performance Computing, University of Edinburgh ·

What does a super computer look like?

That's the network interface. All the yellow stuff? Fiber optic cable, and trays to carry it.