

Is It Finally Fast Enough?

New Advances in Computing Hardware

R. Scott Granneman

© 2021 R. Scott Granneman

Last updated 2021-11-11

You are free to use this work, with certain restrictions.
For full licensing information, please see the last slide/page.

Apple M1

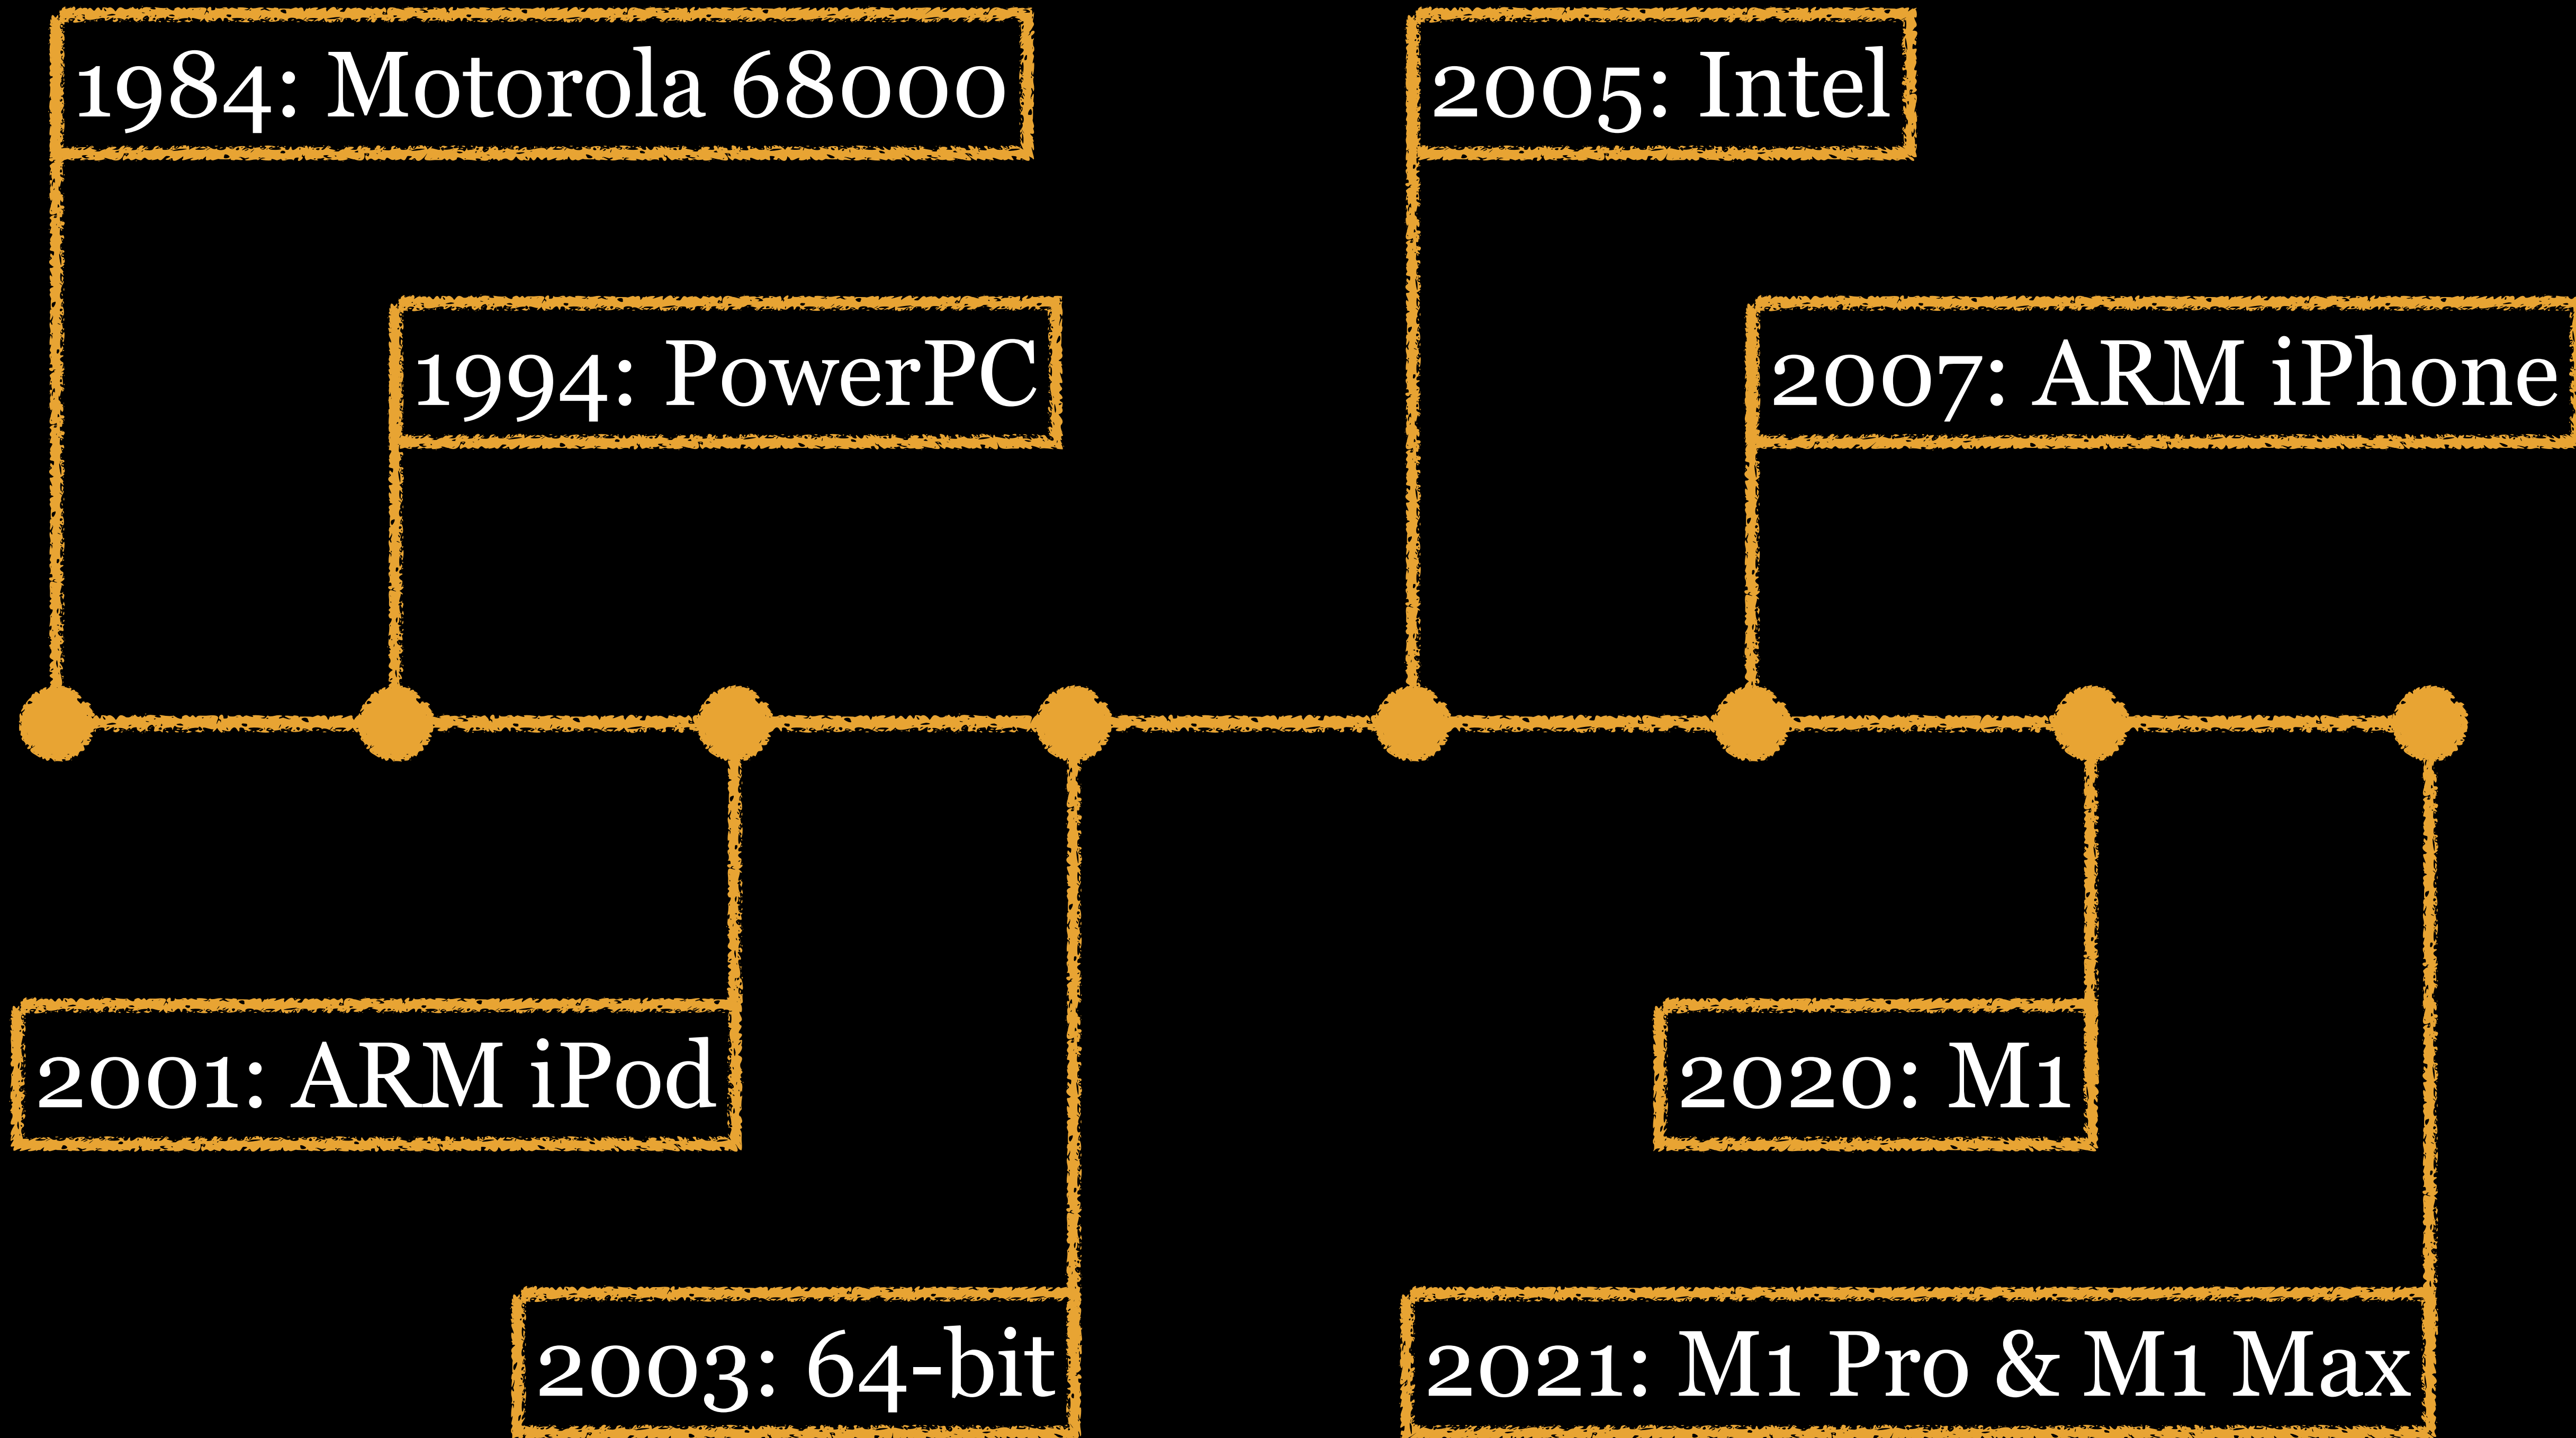
Apple M1 Pro & M1 Max

Intel 12th-gen Core (Alder Lake)

Google Tensor

Framework

Apple's CPU History



2020-06-22: Apple announces Mac switch to M1


2020-11-10: Apple announces M1 inside new MacBook Air, 13" MacBook Pro, & Mac mini

2021-04-20: Apple announces M1 inside 24" iMac, & 11" & 12.9" iPad Pro

2021-10-18: Apple announces M1 Pro & M1 Max inside MacBook Pro 14" & 16"



Apple: “The transition to Apple silicon will take about two years to complete”



CPU

**Thunderbolt
controller**

Apple T2

I/O chip

DDR4 memory

Why the movement towards SoCs?

Needed for smartphones: saves space, ↓ power, & ↑ batteries

Expertise developed from mobile can be applied to desktops & laptops as well


“A system on a chip (SoC; es-oh-SEE or sock) is an integrated circuit (also known as a ‘chip’) that integrates all or most components of a computer or other electronic system. These components almost always include a central processing unit (CPU), memory, input/output ports and secondary storage, often alongside other components such as radio modems and a graphics processing unit (GPU) — all on a single substrate or microchip.” —Wikipedia

“It features the world’s fastest CPU core in low-power silicon, the world’s best CPU performance per watt, the world’s fastest integrated graphics in a personal computer...

...up to 3.5× faster CPU performance, up to 6× faster GPU performance, and up to 15× faster machine learning, all while enabling battery life up to 2× longer than previous-generation Macs” —Apple, November 10, 2020

1st PC chip built using 5-nanometer process technology

16 billion transistors



5 nanometer process



Machine learning accelerators

16-core

Neural Engine

11 trillion operations per second



Thunderbolt / USB 4 controller



Media encode and decode engines



Apple M1

16 billion transistors

Up to
8-core GPU

8-core CPU



Advanced image signal processor



Secure Enclave



Unified memory architecture

Industry-leading performance per watt

Advanced power management

High-efficiency CPU cores

High-performance CPU cores

Low-power video playback

Neural Engine

High-bandwidth caches

Advanced display engine

High-performance GPU

Advanced silicon packaging

Cryptography acceleration

Always-on processor



High-performance video editing

High-performance unified memory

Performance controller

Machine learning accelerators

High-quality image signal processor

Low-power design

High-performance storage

High-efficiency audio processor

Secure Enclave


8-core CPU

The highest-performance CPU
we've ever built.

Up to

3.5x

faster CPU
performance¹




4 high-performance cores

Ultra-wide execution architecture

192KB instruction cache

128KB data cache

Shared 12MB L2 cache



4 high-efficiency cores


Wide execution architecture

128KB instruction cache

64KB data cache


Shared 4MB L2 cache

1/10th of power



For “lightweight,
everyday tasks” like
email or web browsing
at 1/10th the power

Can work together
with 4 high-
performance cores



A detailed die photograph of the Apple M1 chip, showing its intricate circuitry. The chip is divided into several functional blocks. The top-left corner features the Apple logo and the text 'M1'. The top-right corner is highlighted with a white border and contains the text '8-core GPU'. The rest of the die is filled with complex patterns of lines and rectangles representing various components.

 M1

8-core
GPU

Up to 8 cores

128 execution units

Up to 24,576 concurrent threads

2.6 teraflops


82 gigatexels/second

41 gigapixels/second

“World’s fastest
integrated graphics
in a personal
computer”

The logo for the Apple M1 chip, featuring the Apple logo and the text "M1" in a stylized font.The text "8-core GPU" is displayed in a white, sans-serif font. The background behind the text is a dark blue grid pattern.

Designed to “handle extremely demanding tasks with ease, from smooth playback of multiple 4K video streams to rendering complex 3D scenes”



Unified memory architecture

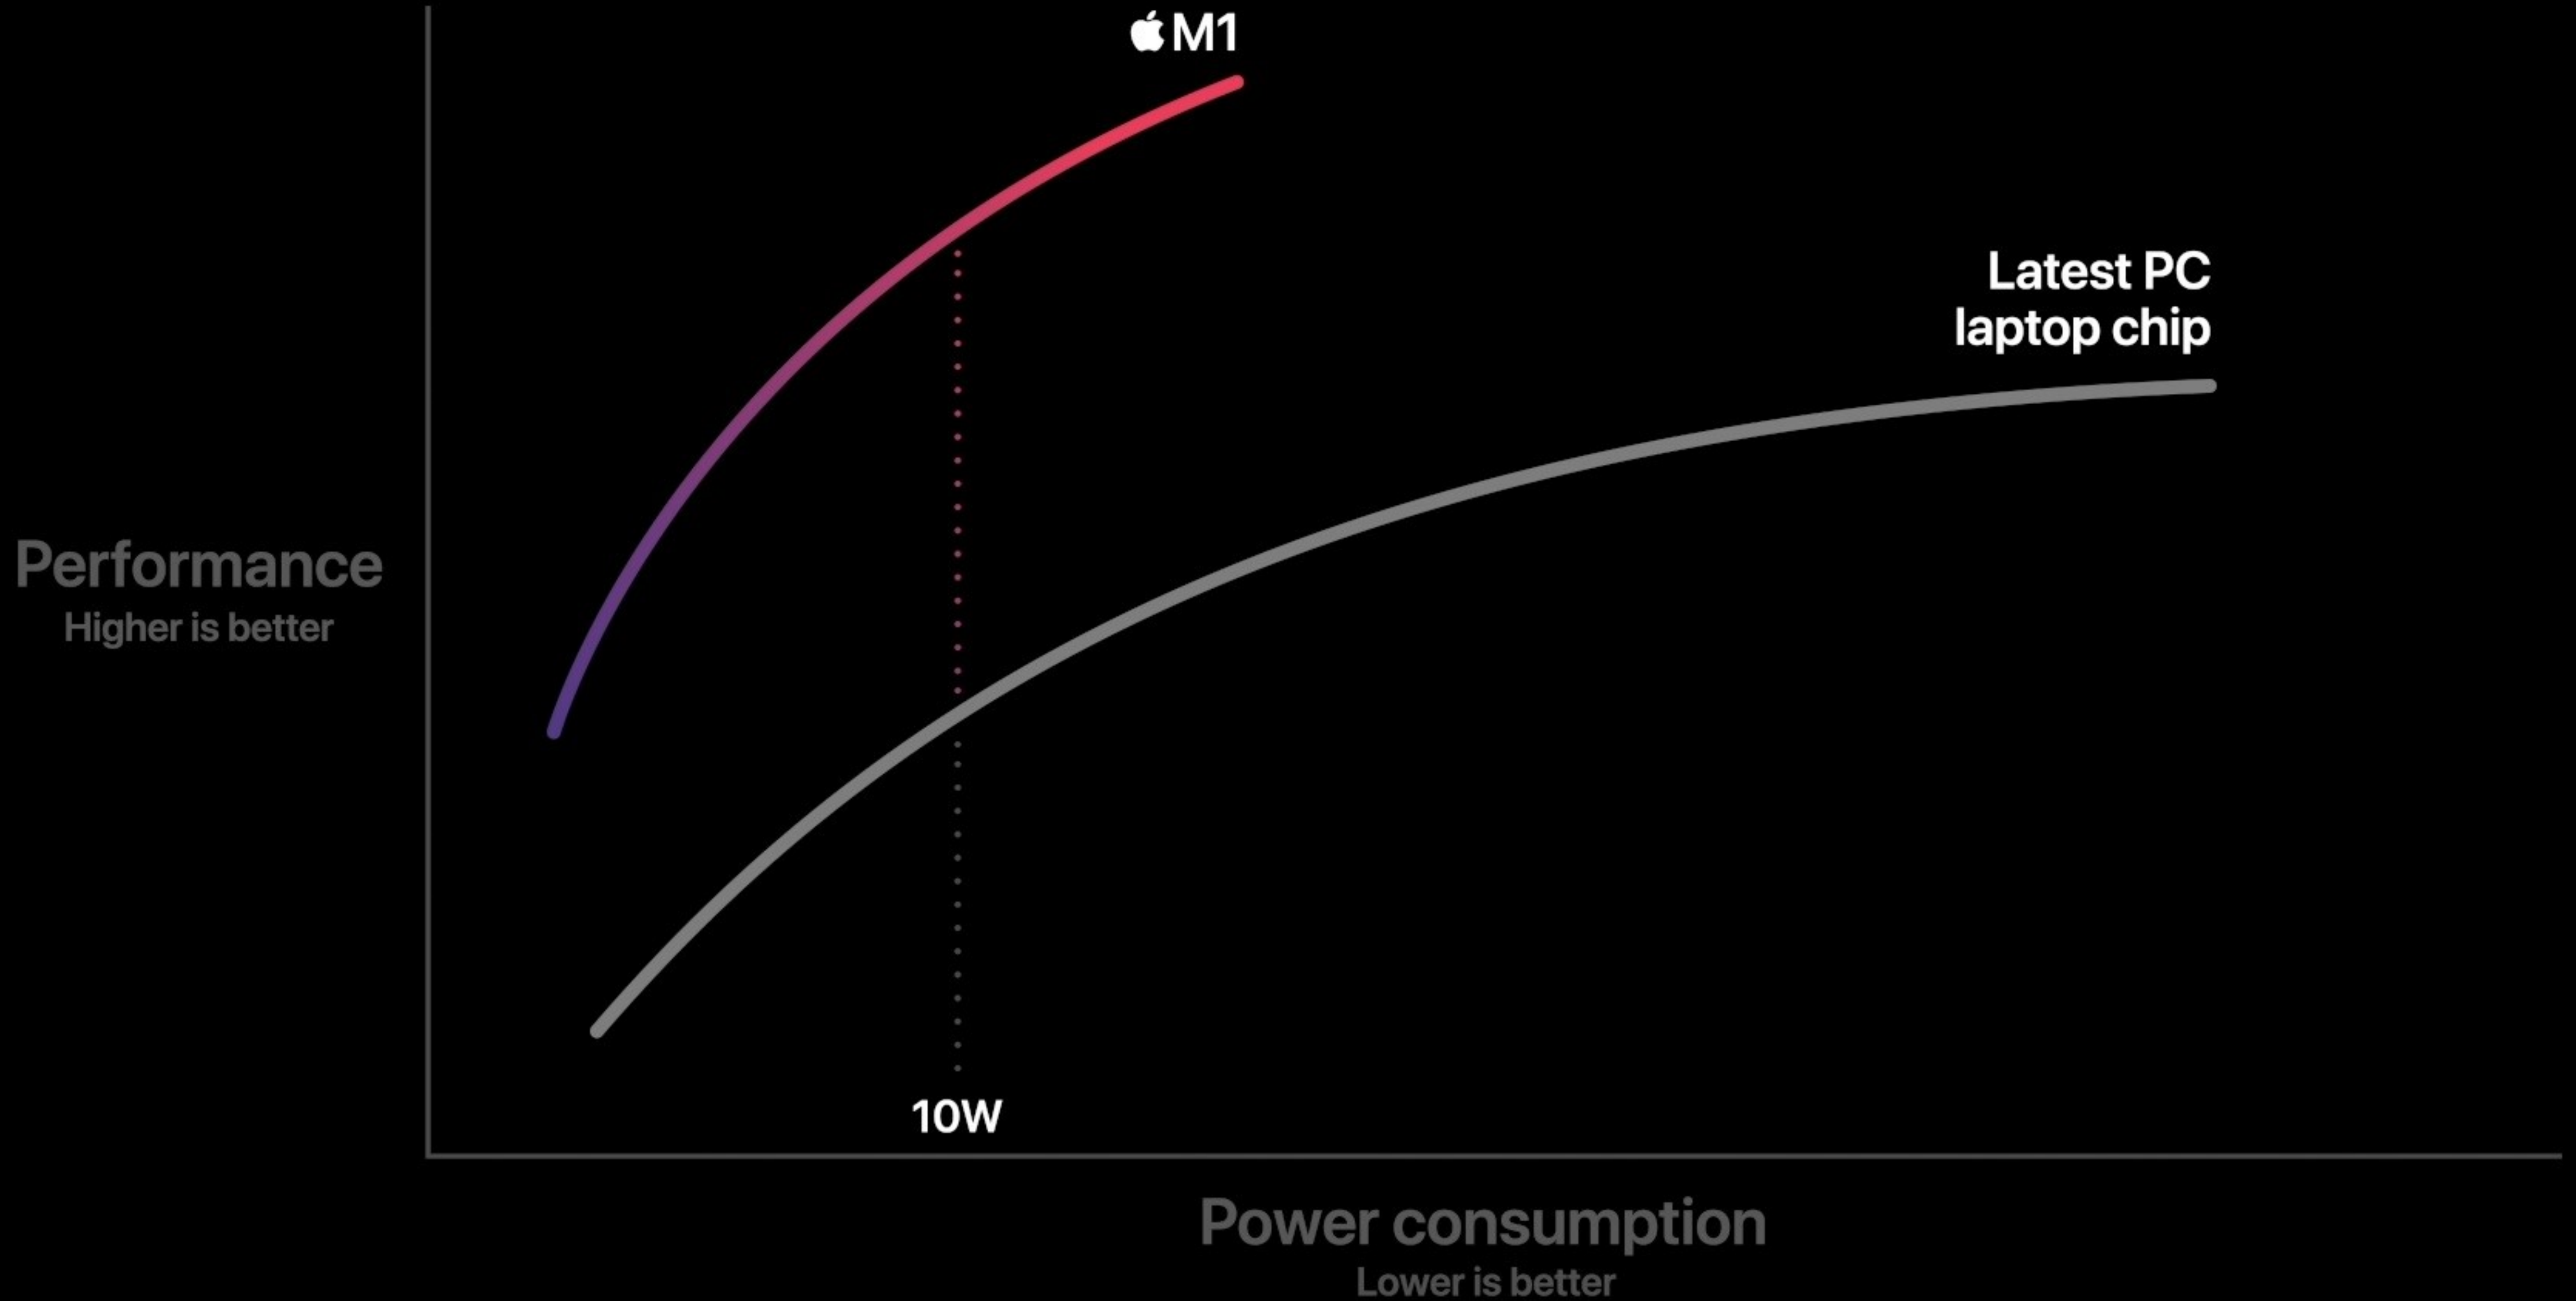
High bandwidth, low latency

Apple-designed package


Accessible to entire SoC

Benefits of UMA: All technologies in M1 SoC can access the same data without copying it between multiple pools of memory, improving performance & efficiency

CPU Performance vs. Power



GPU Performance vs. Power



The image shows a detailed die photograph of the Apple M1 chip. The chip is a square silicon die with a complex grid of circuitry. The top-left corner is highlighted with a red-to-orange gradient, and the bottom-right corner is highlighted with a blue-to-purple gradient. The text 'Apple M1' is overlaid in the top-left, and '16-core Neural Engine' is overlaid in a white box in the bottom-left.

 M1

16-core
Neural
Engine

“The future of machine learning is at the ‘edge,’ which refers to the edge of computing networks, as opposed to centralized computing.

In a centralized machine learning network, users send data to a server, which makes a prediction, and sends that back to the user. This is slower, more expensive, less reliable, and less secure than edge computing, where predictions are made directly on the user’s device.” —Frederik Bussler

“...capable of 11 trillion operations per second, the Neural Engine in M1 enables up to 15x faster machine learning performance. In fact, the entire M1 chip is designed to excel at machine learning, with ML accelerators in the CPU and a powerful GPU, so tasks like video analysis, voice recognition, and image processing will have a level of performance never seen before on the Mac.” —Apple, November 10, 2020

What does the Neural Engine do?

“Final Cut Pro can intelligently frame a clip in a fraction of the time. Pixelmator Pro can magically increase sharpness and detail at incredible speeds.” —Apple

The Neural Engine can be used for: video analysis • voice & facial recognition • artificial intelligence • *computational photography* • recognizing objects in photos • identifying purposes of words in sentences for dictation • create captions on the fly for videos • *Night mode* to capture photos in low-light environments • augmented reality • *deep fusion*: snap a picture & the best pixels from 8 bursts are combined into 1 image • sleep tracking • translation • identify background sounds • handwriting recognition • *palm rejection* for Apple Pencil • virtual assistant response voices

**Faster than 98%
of PC laptops**



Let's see what happens when you try to open 18 default apps on an M1 MacBook Air...

Favourites

Apple iCloud Yahoo Bing Google Wikipedia Facebook Twitter LinkedIn

The Weather Channel Yelp TripAdvisor

Privacy Report

Safari has not encountered any trackers in the last seven days.

Mail

Favourites

Apple iCloud Yahoo Bing Google Wikipedia Facebook Twitter LinkedIn

The Weather Channel Yelp TripAdvisor

Privacy Report

Safari has not encountered any trackers in the last seven days.

Mail

Battery life is outstanding

M1 MacBook Pro 13"

17 hours wireless web browsing

20 hours movie playback

M1 MacBook Air

15 hours wireless web browsing

18 hours movie playback




Apple M1
PRO


The logo for the Apple M1 Pro processor. It features the Apple logo on the left, followed by the text "M1" in a large, white, sans-serif font. Below "M1" is the word "PRO" in a smaller, light blue, sans-serif font. The entire logo is centered on a solid black background.

Apple M1
MAX


The logo for the Apple M1 Max processor. It features the Apple logo on the left, followed by the text "M1" in a large, white, sans-serif font. Below "M1" is the word "MAX" in a smaller, purple, sans-serif font. The entire logo is centered on a solid black background.



🍏 M1




🍏 M1 Pro



🍏 M1 Max

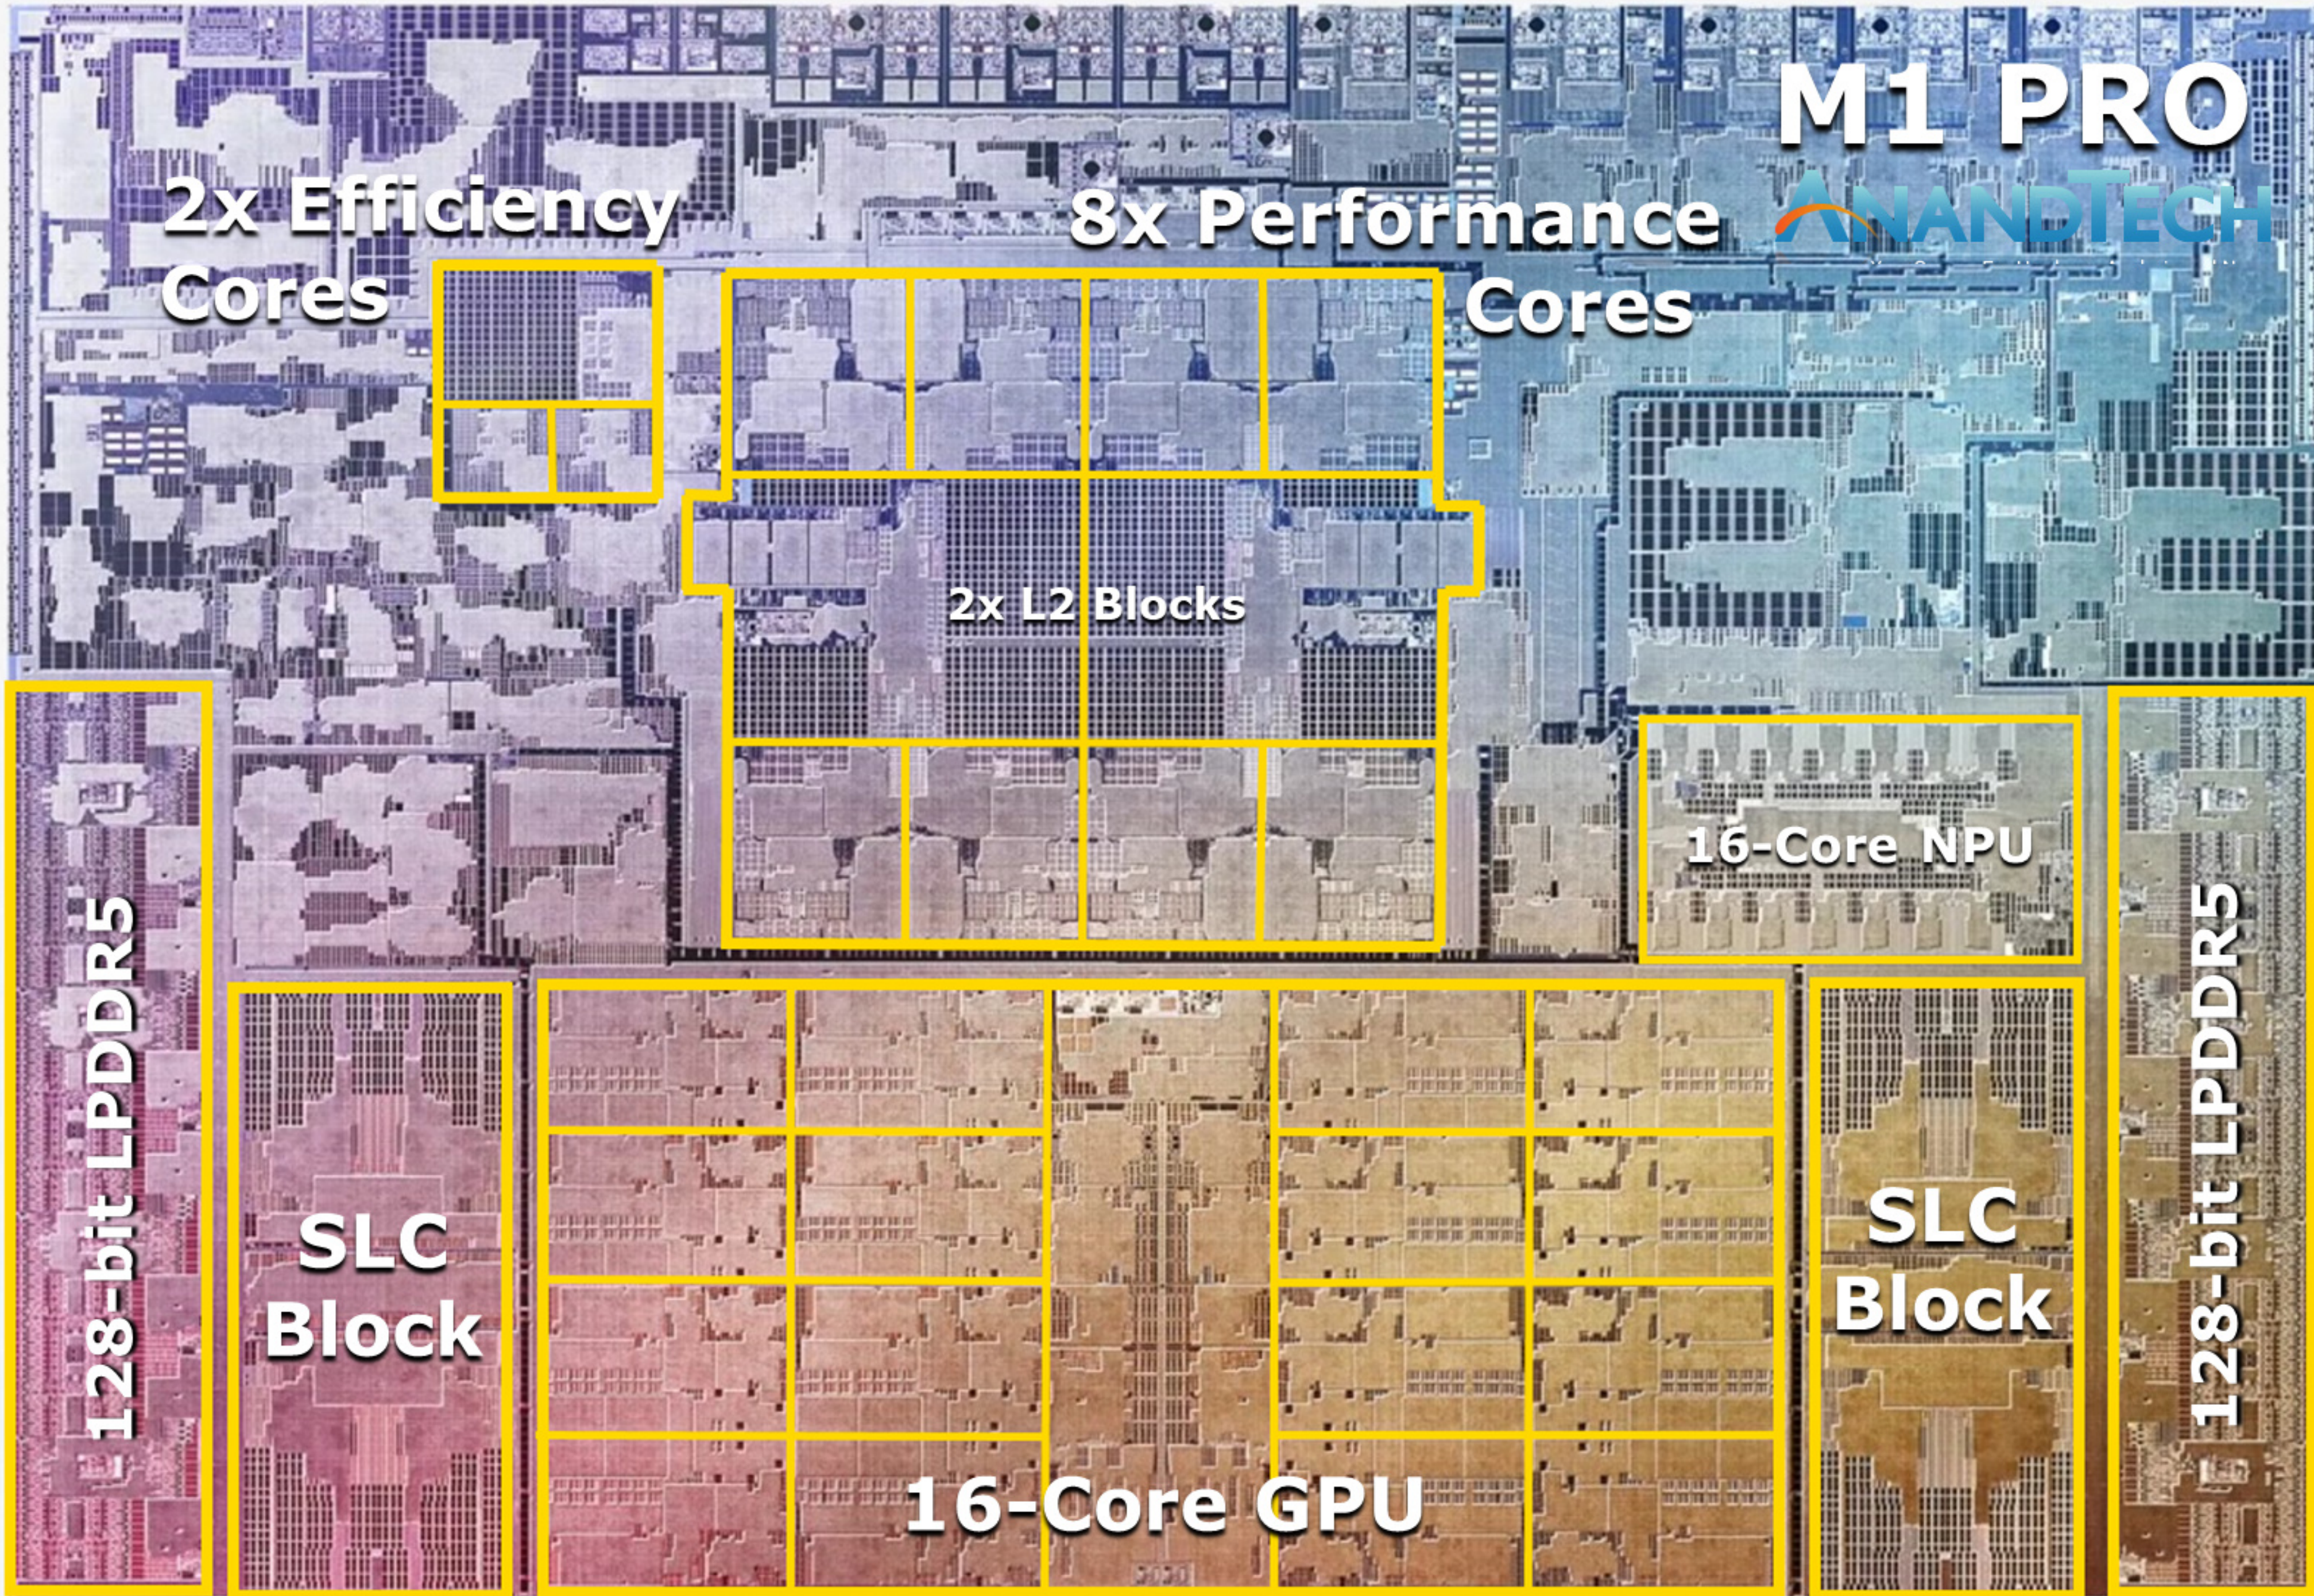
● CL4 NZNM DB71 CLZLM

Z58NNCTJGK 19 VDAKJ



● CL4 NZNM DB71 CLZLM

Z58NNCTJGK 19 VDAKJ



M1 PRO



2x Efficiency Cores

8x Performance Cores

2x L2 Blocks

16-Core NPU

128-bit LPDDR5

SLC Block

16-Core GPU

SLC Block

128-bit LPDDR5

ProRes


encode and
decode



Thunderbolt 4



Secure Enclave



Support for two external displays


Up to

32GB

Unified memory

33.7 billion
Transistors

 **M1**
PRO



Up to
10-core
CPU



Up to
16-core
GPU

16-core
**Neural
Engine**

11 trillion operations per second

Industry-leading
performance per watt

5 nm process


200GB/s
Memory bandwidth

● SM1SRCB ZM87 KMEIGL

T13NNPCRV10 MJMCL

● SM1SRCB ZM87 KMEIGL

T13NNPCRV10 MJMCL

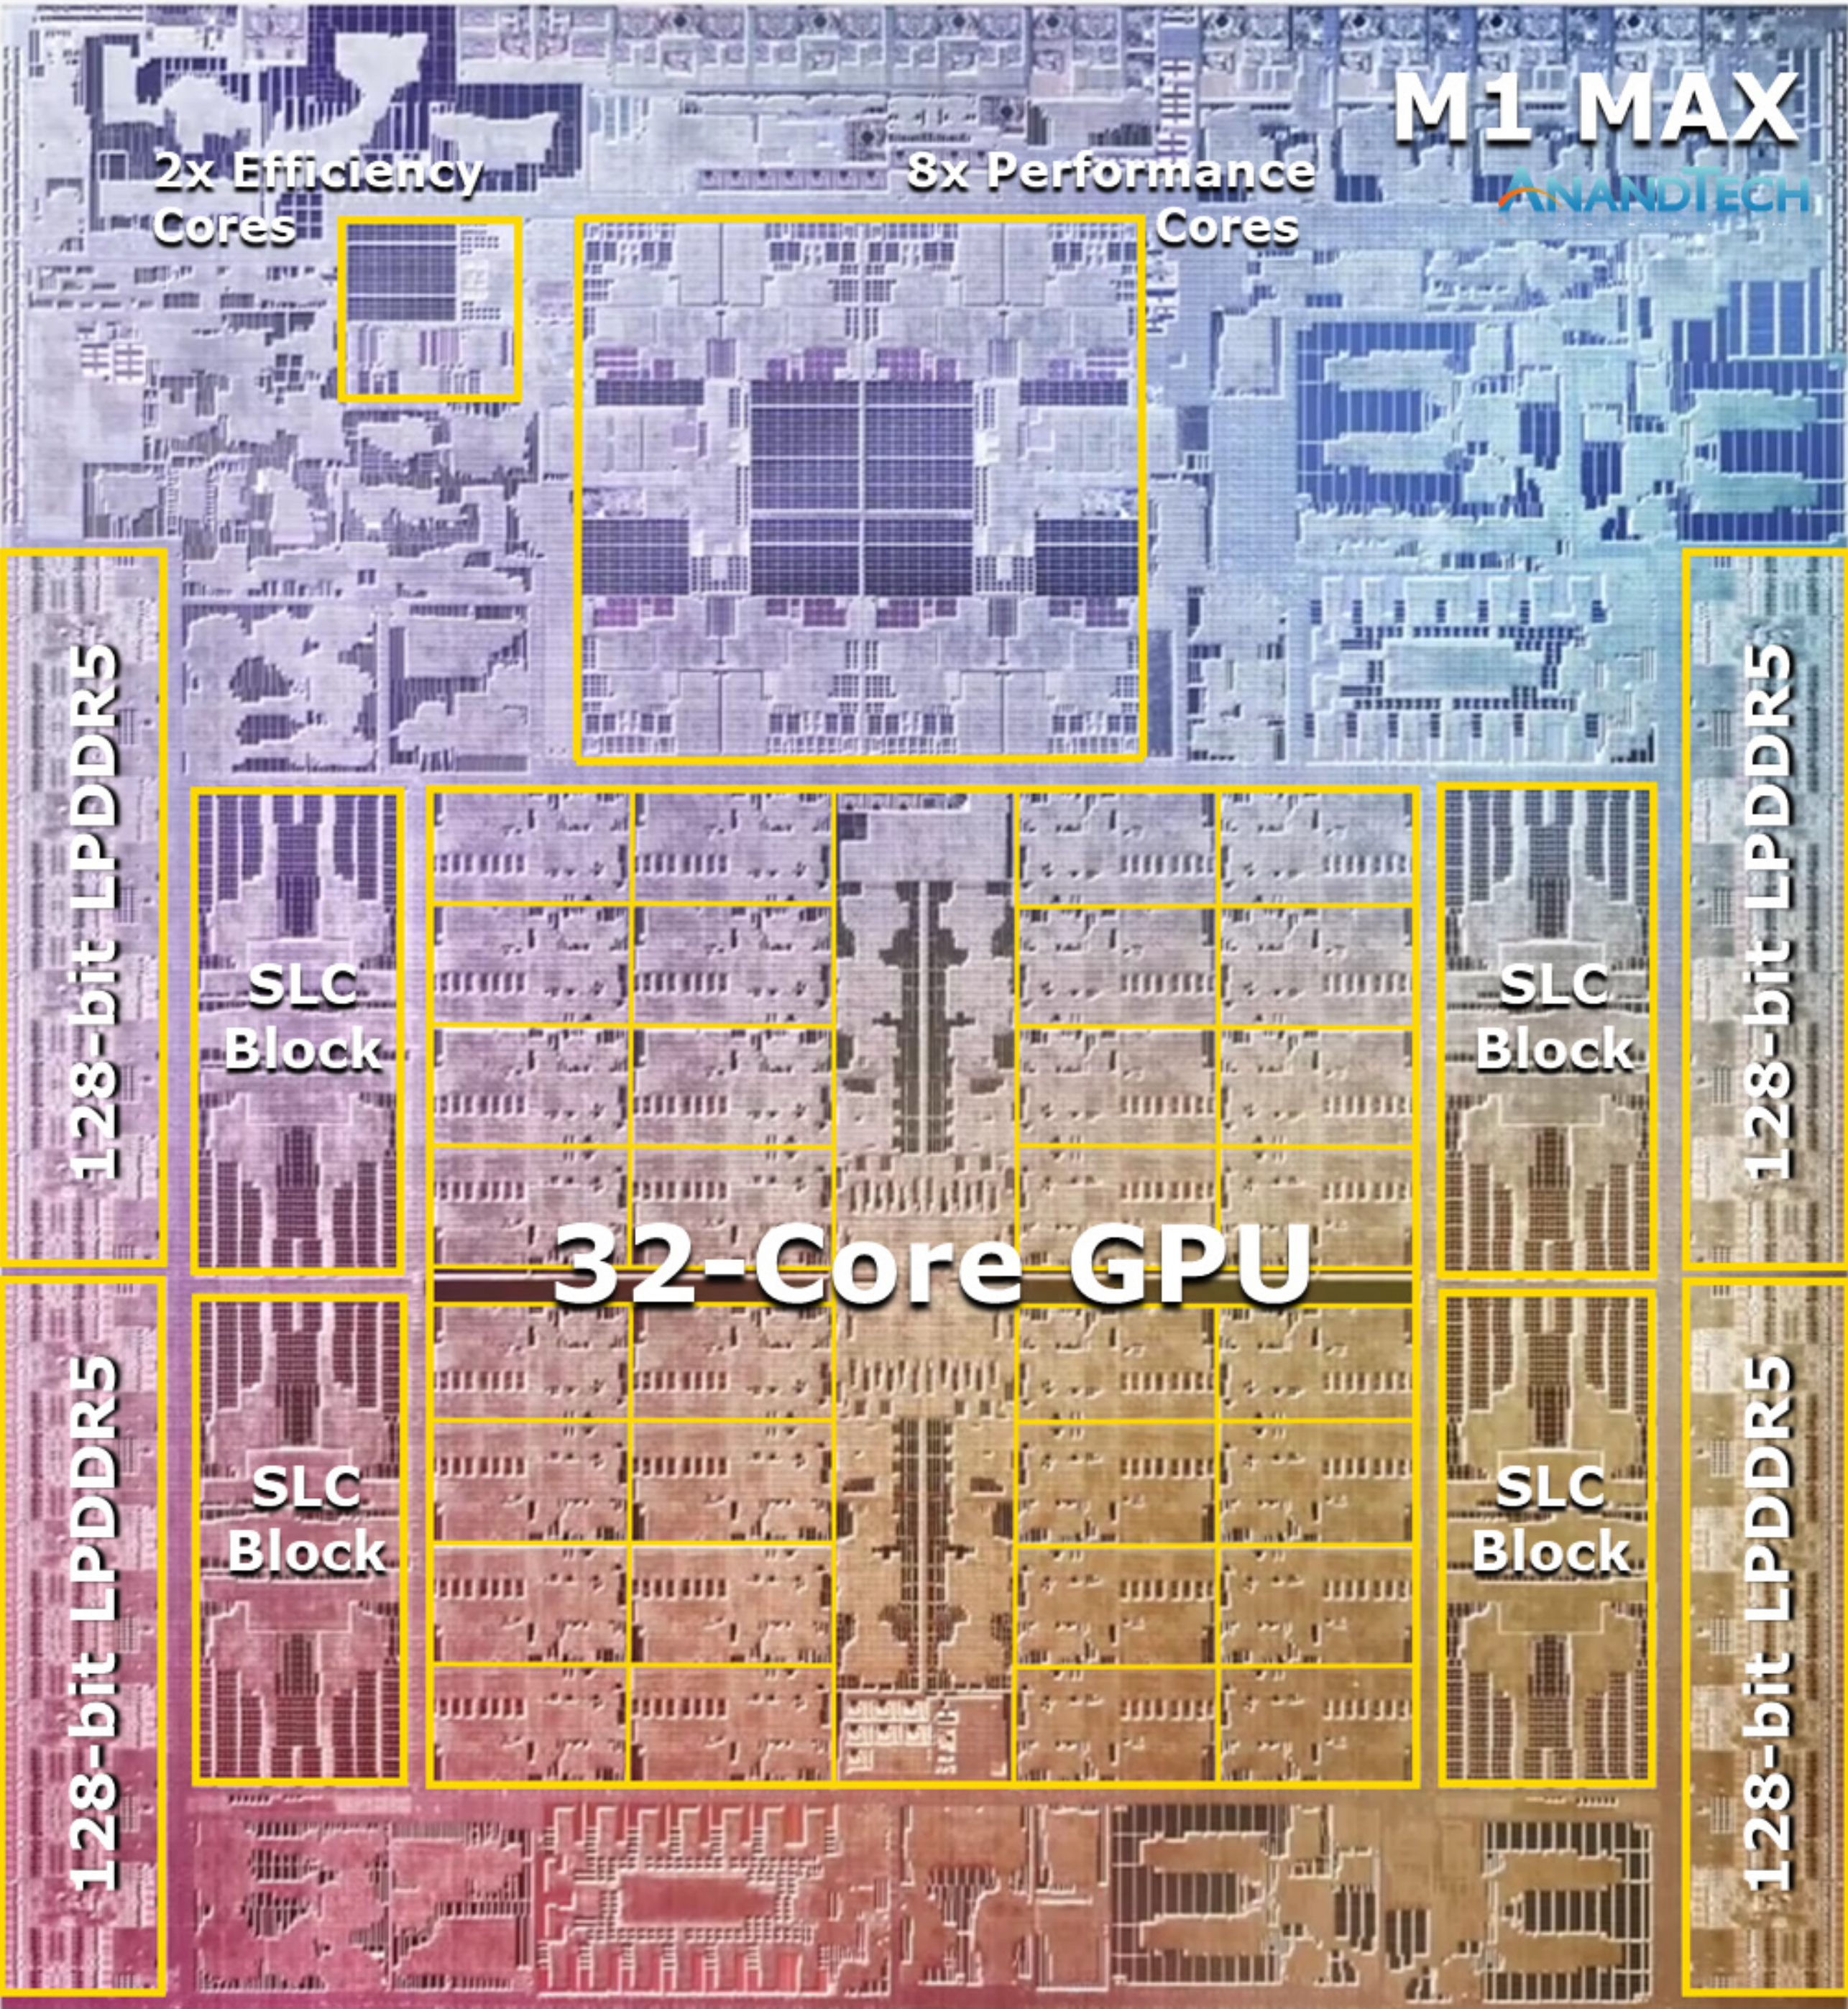


● SM1SRCB ZM87 KMEIGL

T13NNPCRV10 MJMCL

● SM1SRCB ZM87 KMEIGL

T13NNPCRV10 MJMCL



M1 MAX

ANANDTECH

2x Efficiency Cores

8x Performance Cores

128-bit LPDDR5

SLC Block

SLC Block

128-bit LPDDR5

32-Core GPU

128-bit LPDDR5

SLC Block

SLC Block

128-bit LPDDR5

ProRes


encode and
decode



Thunderbolt 4



Secure Enclave



Support for four external displays

Up to

64GB


Unified memory

57 billion
Transistors

 **M1**
MAX



10-core
CPU



Up to
32-core
GPU

16-core
**Neural
Engine**

11 trillion operations per second

Industry-leading
performance per watt

5 nm process

400GB/s
Memory bandwidth



Johnny Srouji, Senior Vice President, Hardware Technologies



Reviewers

“The M1 chips make laptops as powerful as some of the fastest desktops on the market yet so efficient that their battery life beats that of just about any other laptop. ...

The result is something like the difference between a muscle car and a Tesla. The muscle car achieves high speeds with a huge engine that burns a lot of gasoline. The Tesla can hit even higher speeds while consuming less power because its electric motor is inherently more efficient than a gas engine. For years, Intel was making muscle cars; Apple’s big innovation was to build the Tesla of computer chips.” —Farhad Manjoo, *The New York Times*

“I’ll just cut to the chase — for content creation and creative work*, these are the most powerful laptops we’ve ever seen.” —Monica Chin, The Verge

* But not really for gaming, as she makes clear

“But ... when I tell you these laptops are next level, I mean that because the real world performance and the real world capabilities of these laptops is dramatically better than they have any business doing. ... I haven't been able to say something's truly next level in a really long time, but these are. It's one of my favorite products of the year, to be honest.” —Marques Brownlee

“We expected large performance jumps, but we didn’t expect the [*sic*] some of the monstrous increases that the new chips are able to achieve. ... [T]he M1 Pro & Max ... achieve performance figures that simply weren’t even considered possible in a laptop chip. The chips here aren’t only able to outclass any competitor laptop design, but also competes against the best desktop systems out there, you’d have to bring out server-class hardware to get ahead of the M1 Max — it’s just generally absurd. (con’t.)

“What’s clearer, is that the new GPU does allow immense leaps in performance for content creation and productivity workloads which rely on GPU acceleration. ... The combination of raw performance, unique acceleration, as well as sheer power efficiency, is something that you just cannot find in any other platform right now, likely making the new MacBook Pro’s not just the best laptops, but outright the very best devices for the task.” —AnandTech

“The 16-inch MacBook Pro with the M1 Pro is the longest lasting laptop I’ve ever tested in my career as a hardware reviewer, like period. I got 16 hours of continuous use with the 16-inch M1 Pro model. I was jumping between a dozen-ish Chrome tabs, using a couple apps like Slack and Spotify, and often running Zoom calls and YouTube videos over that and I have never seen a laptop last this long.” —Monica Chin, The Verge

Developers







John Szumski

@jszumski



I'm excited to be rolling out fully loaded M1 Max MBPs to all of Twitter's iOS & Android engineers! We're seeing improvements in both top line performance and thermal throttling that currently plague our Intel builds.

3:08 PM · Nov 5, 2021 · Twitter Web App

251 Retweets **86** Quote Tweets **2,003** Likes

Staff Software Engineer at Twitter



Jameson

@softwarejameson



We recently found that the new 2021 M1 MacBooks cut our Android build times in half.

So for a team of 9, \$32k of laptops will actually save \$100k in productivity over 2022. The break-even point happens at 3 months.

TL;DR Engineering hours are much more expensive than laptops!

1:52 PM · Nov 3, 2021 · Twitter for iPhone

3,274 Retweets 592 Quote Tweets 18.2K Likes

Staff Engineer at Reddit; compared to 2019 i9 32GB MBP



Mahyar McDonald

@mahyarm8



All active iOS Engineers at Uber are getting upgraded to 16" M1 Max MacBook Pros with 64GB of RAM, which includes new hires! Looking forward to a faster machine.

2:39 PM · Nov 1, 2021 · Twitter Web App

78 Retweets **39** Quote Tweets **805** Likes

Type to enter text



Swati

@swatiswoboda



Step 1. Someone mentions how it'd be lovely to have M1 Mac as our new work laptops in our developers slack channel.

Step 2. SEVEN minutes later, there is an email in our inbox. Not only are we getting M1s, but we can keep our existing Macs because why not do the simple thing?

5:05 PM · Oct 22, 2021 · Twitter Web App

68 Retweets **39** Quote Tweets **939** Likes

Development Manager at Shopify



Tobi Lütke 

@tobi



That was a really fun thing to announce. We put in a huge order for Apple Silicon.

We are toolmakers here at Shopify. We celebrate and appreciate others who make brilliant tools. Those inspire us to do better ourselves.

10:29 AM · Nov 8, 2021 · Twitter Web App

44 Retweets **14** Quote Tweets **685** Likes

CEO of Shopify



Apple still has to release an M1 Mac Pro

(Supposedly) based on M1 Max, but with 2 dies instead of 1

\$\$\$

Options

M1 Pro (10-core CPU & 16-core GPU) or
M1 Max (10-core CPU & 24- or 32-core GPU)

16, 32, or 64GB unified memory

1, 2, 4, or 8TB SSD storage

8-Core M1 Pro CPU

14-Core GPU

16GB Unified Memory

512GB SSD Storage

14" Liquid Retina XDR display

3 Thunderbolt 4 ports,
HDMI port, SDXC card
slot, MagSafe 3 port

\$1,999



10-Core M1 Pro CPU

16-Core GPU

16GB Unified Memory

1TB SSD Storage

14" Liquid Retina XDR display

3 Thunderbolt 4 ports,
HDMI port, SDXC card
slot, MagSafe 3 port

\$2,499

10-Core M1 Pro CPU

16-Core GPU

16GB Unified Memory

512GB SSD Storage

16" Liquid Retina XDR display

3 Thunderbolt 4 ports, HDMI port, SDXC card slot, MagSafe 3 port

\$2,499

10-Core M1 Pro CPU

16-Core GPU

16GB Unified Memory

1TB SSD Storage

16" Liquid Retina XDR display

3 Thunderbolt 4 ports, HDMI port, SDXC card slot, MagSafe 3 port

\$2,699

10-Core M1 Max CPU

32-Core GPU


32GB Unified Memory

1TB SSD Storage

16" Liquid Retina XDR display

3 Thunderbolt 4 ports, HDMI port, SDXC card slot, MagSafe 3 port

\$3,499



Intel's 12th-generation Core processors, codenamed Alder Lake

Intel's 10-nm i9-12900K offers 16 cores

- » 8 “performance” cores featuring SMT (*symmetric multithreading*, aka *hyperthreading*)
- » 8 lower-performance “efficiency” cores without SMT

AMD's best consumer CPU, the 7-nm 5950X, has 16 cores, all high-performance with SMT

Intel beats AMD's best consumer CPU offerings for single-& multi-threaded performance

Intel expects to release 12th-generation Core processors for laptops in early 2022

According to Geekbench 5 benchmarks, Alder Lake is 1.5× faster than the M1 Pro & M1 Max in multi-core performance

Average multi-core score:

- » Core i9: ~18,500
- » M1 Pro & M1 Max: ~12,500

According to Intel, Alder Lake uses 125W at Base Power, & 241W at Maximum Turbo


Ars Technica found that Intel's i9-12900K drew more than 300W of system power — over 100 watts higher than the Ryzen 9 5950X at full tilt

Alder Lake requires a ton of power, & it runs 🔥🔥🔥

“Intel is seemingly only capable of operating at the extremes: very fast ‘performance at all costs’ chips that consume inordinate power, and power-efficient chips that run very slow. The sweet spot is clearly a proper balance in the middle.” —John Gruber, Daring Fireball

Apple never said Apple Silicon would be the fastest chips on the market

Apple promised industry-leading performance per watt



About Asahi Linux

Began in January 2020

Asahi Linux is a project and community with the goal of porting Linux to Apple Silicon Macs, starting with the 2020 M1 Mac Mini, MacBook Air, and MacBook Pro.

Our goal is not just to make Linux run on these machines but to polish it to the point where it can be used as a daily OS. Doing this requires a tremendous amount of work, as Apple Silicon is an entirely undocumented platform. In particular, we will be reverse engineering the Apple GPU architecture and developing an open-source driver for it.

Asahi Linux is developed by a thriving community of free and open source software developers.

The name

Asahi means “rising sun” in Japanese, and it is also the name of an apple cultivar. 旭りんご (*asahi ringo*) is what we know as the McIntosh Apple, the apple variety that gave the Mac its name.



Search or jump to...



[Pull requests](#) [Issues](#) [Trending](#) [Explore](#)



Asahi Linux

Porting Linux to Apple Silicon macs

[Overview](#) [Repositories 11](#) [Packages](#) [People 3](#)

Popular repositories

m1n1

Public

A bootloader and experimentation playground for Apple Silicon

● C ☆ 924 🍴 56

gpu

Public

Dissecting the M1's GPU for 3D acceleration

● C ☆ 776 🍴 27

linux

Public

Forked from torvalds/linux

Linux kernel source tree

● C ☆ 597 🍴 12

docs

Public

Hardware and software docs / wiki

☆ 434 🍴 14

AsahiLinux.github.io

Public

macvdmtool

Public

People



Top languages

● C ● C++ ● Python ● CSS

● Makefile

[Report abuse](#)

“Is this a Linux distribution?”

Asahi Linux is an overall project to develop support for these Macs. We will eventually release a remix of Arch Linux ARM, packaged for installation by end-users, as a distribution of the same name.” —Asahi FAQ



Alyssa Rosenzweig

@alyssarzg



Everything just happens... instantly? What?

Computers haven't felt this fast since before I was born.

5:58 PM · Sep 29, 2021 · Twitter Web App

7 Retweets **2** Quote Tweets **174** Likes

Asahi Linux GPU lead

October 6, 2021: Progress report for September

“With these drivers, M1 Macs are actually usable as desktop Linux machines! While there is no GPU acceleration yet, the M1’s CPUs are so powerful that a software-rendered desktop is actually faster on them than on e.g. Rockchip ARM64 machines with hardware acceleration. ... Remember, there are still many missing bits (USB3, TB, camera, GPU, audio, etc.) as well as patchsets a bit too problematic to bundle as-is at this time (WiFi, which needs significant rewrites), so don’t expect this to be anywhere near the polished experience that is the goal of our project.”



Chris Wade ✓

@cmwdotme



Linux is now completely usable on the Mac mini M1. Booting from USB a full Ubuntu desktop (rpi). Network works via a USB c dongle. Update includes support for USB, I2C, DART. We will push changes to our GitHub and a tutorial later today. Thanks to the [@CorelliumHQ](#) team ❤️🙏

4:28 AM · Jan 20, 2021 · Twitter for iPhone

1,245 Retweets 216 Quote Tweets 4,742 Likes

CTO of Corellium

June 27, 2021: Linux kernel 5.13 released with initial support for Apple M1 SoC

Still a work-in-progress, with no mainline support yet for the Apple M1 GPU

Apple computer owners can now run Linux natively

M2?

1st generation: M1

2nd generation: M1 Pro & M1 Max

3rd generation: M?

Expected to be made with new 3nm process

Some will have 2 dies instead of 1, so more processor
cores

3rd gen M? CPUs code names & purposes:

- » Lobos & Palma for MacBook Pro & Mac desktops
- » Ibiza will feature lower performance for iPads & MacBook Air

Google Tensor

Virtually all flagship Android phones use Qualcomm's Snapdragon 888; e.g., Samsung Galaxy S21


Google's 1st SoC in a smartphone

Google Tensor

Google's machine learning engine

Tensor security core

Powerful CPU
2 high-performance cores
2 mid cores
4 high-efficiency cores



Advanced image signal processor

Ultra-low power context engine

20-core GPU

2×2.8 Ghz Cortex-X1 cores for peak single-threaded performance



2×2.25 Ghz Cortex-A76 cores for “medium” foreground work

4×1.8 GHz Cortex-A55 cores for low-power background work

5 nm



GEEKBENCH 5: SINGLE-CORE

(Higher is better)



GEEKBENCH 5: MULTI-CORE

(Higher is better)



“...peak CPU and GPU speeds look great in benchmarks but don't always reflect real-world user experience.” —
Google

Tensor GPU is an off-the-shelf ARM Mail G78 MP20 — comparable to Qualcomm's latest in the Samsung Galaxy S21

Modem is a separate chip — not ideal

2 reasons to have a separate modem

1. Introducing new, immature modem technology for a single generation; e.g., Qualcomm introducing 5G in 2019 with Snapdragon 855
2. So-called IP; e.g., Apple isn't allowed to integrate Qualcomm modems in A15 for iPhones, which is why Apple is developing its own modems

“...the Tensor is not dramatically different from Qualcomm’s best SoC in day-to-day usage, which is a great compliment. It’s fine. Google didn’t screw it up.”
—Ars Technica

Pixel 6 Pro with camera bar




Framework

The Framework Laptop is now shipping!

Finally, a high-performance, thin and light notebook designed to last.

Choose your Framework Laptop

Framework Laptop starts at \$999. DIY Edition configurable from \$749.



Marketplace

- All
- Framework Laptop
- Expansion Cards
- Mainboards
- Memory & Storage
- Keyboards
- Parts
- Customization
- Tools
- Software

Framework Laptop ^{/ 2 items}


Framework Laptop

Starting at \$999



Framework Laptop DIY Edition

Starting at \$749





Storage (250GB or 1TB)



DisplayPort



HDMI



SD card



USB-A



USB-C

Customization / 3 items

Bezel - Black

\$39


Coming Soon



Bezel - Gray

\$49

Coming Soon



Bezel - Orange


\$49

Coming Soon



Audio Board Kit

\$14



Fingerprint Reader Kit

\$29



Heatsink and Fan Kit

\$39



Hinge Kit

\$24

Intel® Wi-Fi 6E AX210 vPro®

\$25

Power Adapter - US/Canada

\$49




Mainboards / 3 items

Mainboard - i5-1135G7

\$449


Coming Soon



Mainboard - i7-1165G7

\$699


Coming Soon



Mainboard - i7-1185G7

\$1,049

Coming Soon



Framework Laptop / 2 items

Framework Laptop

Starting at \$999



Framework Laptop DIY Edition

Starting at \$749






Pre-made laptops you can configure

\$999: i5-1135G7 | 8GB Memory | 256GB Storage | WiFi 6 | Windows 10 Home

\$1,399: i7-1165G7 | 16GB Memory | 512GB Storage | WiFi 6 | Windows 10 Home

\$1,999: i7-1185G7 | 32GB Memory | 1TB Storage | WiFi 6 with vPro | Windows 10 Pro



Made with the Framework

Meet Paul
Evan Lee
James Rao
Loren Mitchell
Cassie Hsieh
Alye Lin

Join us!
<https://jobs.framework.com/>

Scan Framework QR Codes with your phone's camera to:

- Get help and repairs
- Order upgrade parts
- Create Marketplace listings

Try it out!

framework

FRANBBAT8
11032002V

framework

FRANBBAT8
CP8111
45005T
<https://fr.mw/franbbat8>

SPEAKER

2 x 2W

To see more details:
<https://fr.mw/franbbat8>

framework

framework

BATTERY

Rechargeable Li-Ion Polymer Battery

Model: Framework Laptop

Typical Capacity: 3678mAh/56.64Wh

額定電容量 Rated Capacity: 3572mAh/55.0Wh

額定電壓 Nominal Voltage: 15.4VDC

充電限制電壓 Limited Charging Voltage: 17.6VDC

FRANBBAT8
11032002V

二次鋰離子電池組
Rechargeable Li-Ion Polymer Battery

警告 Caution
危險 詳情請參閱說明書
禁止折騰、刺破、擠壓、加熱或燃燒
Do not disassemble, puncture, crush, tear or burn the battery.

請使用原廠指定充電器
Please use the charger specified by the original factory.
請勿拆解、刺破、擠壓、加熱或燃燒電池
Ne pas désassembler, piquer, écraser, déchirer ou brûler la batterie.

製造商 Manufacturer: 廣東新嘉源科技有限公司 Dongguan Amperex Technology Limited
中國製造 Made in China

回收符號
Li-ion 00
原廠請回收

FRANBBAT8
11032002V

MEMORY


DDR4 Channel 1
up to 32GB each

MAINBOARD

FRANMCP

MEMORY

DDR4 Channel 0
up to 32GB each



Made with the Framework

Hiroti Paul
Ewan Lee
James Rao
Loren Mitchell
Cassian Heath
Alyk Lin

Kenji
Eddy
Handy Ni
Addy Lu

Come join us!
<https://jobs.framework.com/>

Scan Framework QR Codes with your phone's camera to:

- Get help and repairs
- Order upgrade parts
- Create Marketplace listings

Try it out!

framework

SPEAKER
2 x 2W
To see more details:
<https://fr.mw/FRANEX>

framework BATTERY
KC Certification number: 2000000000X
인증번호: 82-20-8321-0337

This battery is replaceable
To see more details:
<https://fr.mw/FRANBB>

FRANBBAT01

二次锂离子電池組
Rechargeable Li-Ion
Polymer Battery

警告 Caution
危險 詳情請參閱說明書
禁止折騰、刺破、擠壓、加熱或燃燒
Do not disassemble, puncture, crush,
tear or burn the battery.

請使用原廠指定充電器
Please use the charger specified
by the original factory.
Ne pas désassembler, piquer, écraser,
déchirer ou brûler la batterie.

製造商 Manufacturer
東莞新嘉源科技有限公司
Dongguan Amperex
Technology Limited
中國製造 Made in China

型號 Model
Framework Laptop

Typical Capacity
3678mAh/56.64Wh

額定電容量 Rated Capacity
3572mAh/55.0Wh



額定電壓 Nominal Voltage
15.4VDC

充電限制電壓
Limited Charging Voltage
17.6VDC

中國製造 Made in China
AICP6/S9/67

FRANBBAT8
11032002V

Li-ion 00





Linux on the Framework Laptop

All | Sep 01 2021 | by **Nirav Patel**

We love Linux at Framework. We decided from the start of Framework Laptop development to offer the [DIY Edition](#) without an operating system pre-loaded to give you the option to bring your favorite Linux distribution. There has been immense interest in this configuration, with it outselling pre-configured systems with Windows 10 by a wide margin. We provided pre-release hardware to developers and maintainers at Fedora, elementary OS, NixOS, and Arch to make the Linux experience as smooth as possible, and we've been impressed by the incredible variety of Linux distros (and OpenBSD too!) being used by all of you.

Since we planned for Linux support from the outset, we made sure to use hardware that is well-supported and has drivers available. There are just a few areas where support is brand new and making its way into different distributions. Intel 11th Gen Core Processors, Intel AX210 WiFi (which is optional on the DIY Edition), and our Goodix-based fingerprint reader are the three items that require a newer kernel or packages than many distros currently ship. We recommend using 5.12 or newer for a kernel to get solid platform, WiFi, and bluetooth functionality, along with libfprint 1.92.0 or newer for the fingerprint reader. All of the other hardware like speakers, microphones, headphones, webcam, hardware privacy switches, keyboard media keys, ambient light sensor, and all of the Expansion Cards should work completely.

The extremely active and vibrant Linux subforum in the Framework Community is the best place to go for the latest

Problems

...but hey, it's new, so cut it some slack

“The power button has a built-in fingerprint scanner, which is convenient though I had trouble getting it to recognize my index finger on the first try every time.

”

...

“But regardless of the considerable memory size, the Framework laptop trailed behind the Samsung and Razer notebooks in my benchmarks. ...

The Framework’s battery life is also too short. It lasted five hours and a few minutes before it petered out in Gizmodo’s battery rundown test. ...”

“I do wonder if its numbers would have ranked better if it didn’t run so hot. Throughout the benchmarking process, I clocked the Framework laptop running at around 108 degrees Fahrenheit. At some points, the computer felt too hot to touch. ...

The Framework laptop certainly does work on cooling itself, which I can tell because it’s loud enough to become a noisy distraction. I’ve only ever experienced this kind of fan activity with high-performance gaming laptops, which I expect from those machines since there’s also a GPU inside that needs temperature regulation. ... But even at the lowest setting, the thermometer was clocking in 96-degree temperatures around the function keys.”

“The Framework laptop has a standard Intel integrated graphics chip, so don’t expect to do much gaming beyond the occasional virtual card game.”

Summing Up

Apple M1

Apple M1 Pro & M1 Max

Intel 12th-gen Core (Alder Lake)

Google Tensor

Framework

Thank you!

scott@granneman.com

www.granneman.com

[@scottgranneman](#)

jans@websanity.com

websanity.com

Is It Finally Fast Enough?

New Advances in Computing Hardware

R. Scott Granneman

© 2021 R. Scott Granneman

Last updated 2021-11-11

You are free to use this work, with certain restrictions.
For full licensing information, please see the last slide/page.

Licensing of this work

This work is licensed under the Creative Commons Attribution-ShareAlike 4.0 International License.

To view a copy of this license, visit <http://creativecommons.org/licenses/by-sa/4.0/>.

You are free to:

- » *Share* — copy and redistribute the material in any medium or format
- » *Adapt* — remix, transform, and build upon the material for any purpose, even commercially

Under the following terms:

Attribution. You must give appropriate credit, provide a link to the license, and indicate if changes were made. You may do so in any reasonable manner, but not in any way that suggests the licensor endorses you or your use. Give credit to:

Scott Granneman • www.granneman.com • scott@granneman.com

Share Alike. If you remix, transform, or build upon the material, you must distribute your contributions under the same license as the original.

No additional restrictions. You may not apply legal terms or technological measures that legally restrict others from doing anything the license permits.

Questions? Email scott@granneman.com