




Tech Toolbox: Processor Proceedings

Technology, Privacy, and eCommerce



The brains of our computers, smartphones, and tablets are microprocessors called CPUs (Central Processing Units). For the past 42 years, the most dominant [CPUs](#) inside most laptops and other personal computers have been made by Intel, using their proprietary x86 instruction set architecture. For most of that time, Intel has done a fantastic job building chips optimized for that architecture.

Even though CPUs are the “brains” of your devices, they only work when given very specific instructions — called the instruction set — which tell the processor how to move data between registers and memory or perform a calculation using a specific execution unit (such as multiplication or subtraction). Applications written in various higher-level programming languages (like Java or C++) are compiled for specific instruction sets so that they run correctly. Instruction sets must be tailored to the CPUs that use them so they can enable optimized interactions. This link between instructions and processor hardware design is what makes a CPU architecture, which can be designed for different purposes, like extreme number crunching, low energy consumption, or minimal silicon area.

One of Intel’s founders, Gordon Moore, famously predicted that the number of transistors on a semiconductor chip like a CPU would about [double every two years](#), and for more than three decades that has been fairly accurate. The first x86 microprocessors built by Intel in 1978 contained 29,000 transistors. The 286 (1982) had 134,000. The 386 (1985) had 275,000. Their Pentium (1993) had over three million. Their Core i7, released in 2018, has nearly two billion, and their 28 core Xeon Platinum (2017) has eight billion transistors!

But it turned out that jamming all those transistors into a chip operating on what is by now the rather aging x86 instruction set architecture came at a cost — they burned an awful lot of electricity, and they burned it very hot. I think most of us have had the experience of making a Zoom call, editing a video, or just having too many browser tabs open and suddenly hearing your laptop fans start to sound like a jet turbine. Or, if the laptop was actually on your lap, feeling it get pretty toasty. And (in the days when people could travel), we can remember how quickly our laptop batteries drained away.

But it turned out that jamming all those transistors into a chip operating on what is by now the rather aging x86 instruction set architecture came at a cost — they burned an awful lot of electricity, and they burned it very hot.


Recently, we have also seen striking examples of CPUs that have taken a different path. Tablets, smartphones, and smart watches needed to be able to hold a longer charge on a much smaller battery and work without burning a hole in your pocket (or on your wrist). And they had to be able to function without fans.

The prime mover in this space has been Apple. From the beginning, Apple realized that it could not build an iPhone with an Intel CPU inside. Instead, it began working with chips built around the ARM instruction set and architecture.

ARM, which stands for Advanced RISC Machine, is built on their Reduced Instruction Set Computing (RISC) architecture, which is much simpler than the Complex Instruction Set Computing (CISC) approach used by Intel's x86 architecture. Arm's RISC instructions are characterized by a close correlation between the number of instructions and the number of resulting operations they trigger. CISC, by comparison, offers much more complex instructions, many of which execute multiple operations simultaneously. This can lead to better performance in certain contexts, such as supporting multiple users in workstations or servers, but results in much greater power consumption and heat generation.

This Intel CISC approach made more sense when first created, as CPU chips had a smaller number of transistors because they could maximize the number of operations conducted by each one. But as CPU fabrication technology advanced, it became possible to run increasingly complex programs just by throwing greater numbers of transistors *into* each CPU and *at* the operations the CPU performed. So, the RISC approach began catching up with and even surpassing the CISC approach while using substantially less power and generating far less heat.

CPU performance vs. power



Up to

2x

faster CPU performance¹

Matches peak PC performance using

25%

of the power¹

The benefits of this approach have become obvious in the past several years, when benchmarks for Apple-designed ARM chips in its iPhones and iPads began to outstrip the performance of even Apple's own laptops using higher end Intel chips. Furthermore, Intel has had problems in recent years meeting their own published deadlines for introducing more advanced chipsets, which meant that Apple and other personal computer manufacturers had to wait for Intel to catch up in order to manage their own release timelines. Since Apple has always been notoriously unwilling to depend upon other companies to meet its product release goals, it became obvious that one day Apple would make the transition from using Intel CPUs in their products to using their own Apple-designed silicon based on ARM.

Well, that "one day" has arrived. Apple announced this transition at their World-Wide Developer Conference (WWDC) in June 2020, and the first M1 Macs became available in November 2020. In fact, I am writing this on one of them right now. This will not be a review of my new M1 MacBook Air — although I will tell you that it is a fantastic machine, superfast, quiet as can be, with a great keyboard and trackpad — but that is not the point of this article. The point is that this likely heralds a new era in portable computing. **Here's why.**

First, these new M1 chips are not much different or faster than their iPhone and iPad siblings, because those are already incredibly fast. They do throw more cores at the work than an iPhone can or should contain, but that just makes them able to do more in the way of multi-threaded heavy lifting.

So, going forward, the old paradigm about using smartphones for light duty and laptops for heavy duty computing won't make as much sense. The decision about when to use what will be more about screen size or storage capacity or what happens to be ready at hand than about CPU performance.

So, going forward, the old paradigm about using smartphones for light duty and laptops for heavy duty computing won't make as much sense. The decision about when to use what will

be more about screen size or storage capacity or what happens to be ready at hand than about CPU performance.

Second, battery life will for the first time actually be better on a laptop than on a tablet, because they basically use the same processors with the same energy profiles and because the laptop has a bigger battery enclosure. So, that reason for bringing an iPad or other tablet along on a business trip instead of a Mac or other laptop will also go away. Laptops will generally be heavier, but not necessarily by much – if I carry both my iPad Pro and its Magic Keyboard (which is also fantastic, by the way), the difference is just a few ounces.

Third, the apps we need for a particular purpose also may not tip the scales. Because iOS apps are designed to run on ARM chips, the new M1 Macs with relatively few developer modifications will be able to run all of the iOS apps I have come to prefer. Little by little, I expect the converse will also be true.

And, although Apple currently and rather vehemently denies that it has any plans for making touch sensitive screens for their M1 Macs, I don't believe them. There are too many hints of touch-centricity in their new versions of MacOS to find that credible. I think within one or two years, they will release some really interesting hybrid machines along the lines of the Microsoft's Surface Pros.

Even if Apple doesn't, it is crystal clear that Microsoft will once again soon begin using ARM chips in lieu of Intel ones sometime in the near future, and so will Lenovo, Dell, etc. There will then be this huge convergence and interoperability that will occur between mobile devices and personal computers.

Bonus – Because these chips are so small, cool, and powerful, and don't need huge batteries, they open the door for building really small things, like smart glasses that may enable us to finally have access to the promises of augmented reality we geeks have been drooling over. Apple actually patented lenses for such glasses for which you could adjust the prescription on the fly. Imagine that! And imagine glasses that would give you map guidance, reference information, important messages, and the like all while being lightweight and comfortable.

It won't end there either. Just as Raspberry Pi chips (tiny, low power chips used by hobbyists) are finding their ways into a huge variety of devices and use cases, these tiny but much more powerful chips could soon find homes we might never previously have considered. These new chips herald an age of potential miniaturization like we have never seen! So, keep your eyes peeled. Another step toward the future of computing has arrived.

[Greg Stern](#)



Former Global Integration Counsel

Chubb, Independent Consultant