Reed Electronics Group

MICROPROCESSOR REPORT

www.MPRonline.com

THE INSIDER'S GUIDE TO MICROPROCESSOR HARDWARE

Why Apple Feels Chipper

Hiring More Custom-Chip Designers Makes Sense for Apple

By Tom R. Halfhill {5/26/09-01}

Print may be dying and newspapers may be dinosaurs, but a story in the *Wall Street Journal* can still stir up a tizzy. For days after the *WSJ* reported that Apple is hiring more people with chip-design experience, industry pundits either scratched their heads in wonder or questioned why "a company like Apple" might want to design its own chips.

That Apple has been accumulating expertise in chip design is not news. More than a year ago, Apple acquired P.A. Semi, a fabless semiconductor company headed by former senior engineers from DEC's Alpha and StrongARM teams. Apple paid $278 million in cash for P.A. Semi—a princely price for a short-lived 150-person company. At the time, *Microprocessor Report* speculated that Apple may have been legally obligated to compensate P.A. Semi for not using its PWRficient processors in Apple Macs or other products. (See *MPR 5/27/08-02*, "A Tale of Two Companies.")

Since then, as the *WSJ* noted, Apple has been hiring additional engineers and managers with chip-design backgrounds. Among the heavy hitters are Bob Drebin, former chief technology officer of AMD; Raja Koduri, former chief technology officer of AMD's ATI group; and Mark Papermaster, a Power Architecture architect whose defection from IBM provoked legal action to enforce a noncompetition agreement. It hasn't escaped notice that Drebin and Koduri formerly worked in ATI's Imageon division (mobile graphics), which AMD sold to Qualcomm.

On top of that, Apple's recent job listings and the online profiles of Apple employees suggest a stronger tilt toward semiconductor engineering. It all seems to add up: Apple is designing chips. But judging from the way some observers are interpreting these tea leaves, you would think Apple is committing a tragic error that will undermine the company's (*ahem*) core business.

MPR Photo by Tom R. Halfhill

On June 29, 2007, early adopters besieged Apple Stores to buy the first iPhone. To create this kind of frenzy, Apple must differentiate its products from those of competitors.

In reality, there's no cause to be alarmed or mystified. Apple has good reasons for acquiring more chip designers. Nor is it out of character. Indeed, it's consistent with Apple's strategy since 1984 and is more important today than ever. It may also be instructive for other companies trying to decide whether in-house chip design is a sensible strategy.

'Think Different' With Differentiation

Apple cofounder Steve Wozniak designed the Apple I and Apple II computers with off-the-shelf parts in the 1970s because they were all he could afford. Indeed, he scrounged the first 6502 microprocessor for the Apple I prototype as a free sample from MOS Technology. (Guerrilla engineering at its best.) But since introducing the Macintosh in 1984, Apple has relied heavily on custom engineering to differentiate its hardware from that of commodity-oriented competitors.

Macintoshes have used custom chips for 25 years, mainly for system logic, I/O, and graphics. Apple has even dabbled in CPUs. In 1991, Apple formed an alliance with IBM and Motorola to help define the PowerPC architecture. (See *MPR* 10/16/91, "Apple, IBM, and Motorola Sign Contracts for Far-Reaching Collaboration.") Large chunks of the Mac OS once resided in custom ROMs, until execution-in-place lagged behind the rising speed of DRAM. Recent Macs tend to use more standard silicon, especially for graphics, but Apple's consumer-electronics products are stuffed with custom chips. Apple strives to differentiate its products in three ways: optimized hardware, easy-to-use software, and trend-setting industrial design.

Today, it's almost impossible to build a competitive smartphone or any other cutting-edge consumer-electronics product without using SoCs. This is especially true for mobile consumer electronics. The cramped dimensions of these systems and their tiny batteries require highly optimized silicon to avoid a clunky appearance and disappointing battery life. Off-the-shelf SoCs like the Texas Instruments OMAP devices make it possible for many vendors to create competitive products without designing their own chips. But anyone can buy an OMAP, so differentiation narrows to the system software and industrial design.

Apple has built its entire business on being different, sometimes gratuitously so. Apple is also willing to sacrifice market share in pursuit of higher profit margins. (With the iPod, unlike the Mac, Apple has won a rare victory on both fronts: high margins and industry-dominant market share.) Apple is confident it can create better system software, superior user interfaces, and innovative industrial design. If, in addition, Apple can deliver better hardware by using better SoCs, market success is almost guaranteed. And market success allows Apple to charge margin-rich prices.

For many years, Apple has outsourced the design of some custom chips to other companies. Apple provided the specifications and some oversight. For as long, Apple has employed chip engineers to be liaisons with outside design houses and semiconductor suppliers. Now, Apple is hiring more chip engineers.

One possibility is that Apple is merely adding talent to keep pace with its expanding product catalog. Another possibility is that Apple's chip suppliers may have missed some recent project deadlines, prompting Apple to pull more chip design in house to regain control over its product schedule. The strongest possibility is that Apple is plunging deeper into custom chip design. Such a plunge wouldn't be out of character for a growing consumer-electronics company with big ambitions.

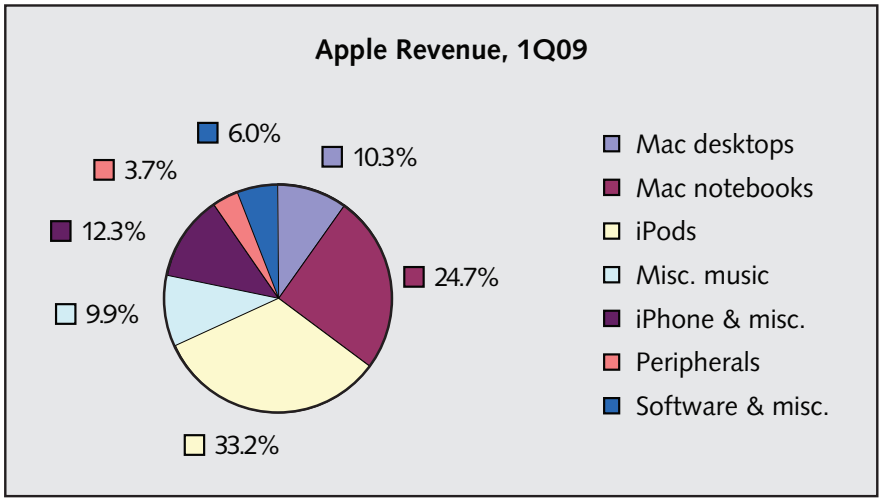

Figure 1. Apple revenues, 1Q09. All those little iPods add up. At $3.37 billion for the quarter, they account for the biggest share (33.2%) of Apple's revenues. Mac notebook computers account for the next-largest slice, at $2.51 billion. Revenues from Mac desktops ($1.04 billion) barely exceed revenues from Apple's iTunes store and other music-related sources ($1.01 billion). Even the iPhone generates more revenue ($1.24 billion) than Mac desktops. (Data source: Apple.)

Silicon Is the Key Ingredient

Big consumer-electronics companies commonly depend on custom SoCs for differentiation. No one questions why Canon, Panasonic, Sony, and Toshiba design custom chips. With the iPod and iPhone, Apple is entering the same league. Apple is becoming a consumer-electronics company that also makes personal computers—not a computer company that also makes consumer electronics. (Actually, the distinction between personal computers and consumer electronics is largely historical and obsolete.)

Figure 1 illustrates the revenues generated by Apple's current product mix. The iPod and iPhone, not the Mac, are Apple's prime revenue sources.

Canon, a major camera manufacturer since the 1950s, is a good example of technology evolution in consumer products.

As film cameras incorporated more automation and gave way to digital cameras, silicon became a vital component. Canon differentiates its digital cameras by designing its own image sensors and image-processing chips.

In a rare example of SoC marketing to the general public, Canon has branded a line of image processors under the DIGIC name. DIGIC reportedly stands for "Digital IC." Like other brand names for chips (Opteron and Core i7 come to mind), "DIGIC" is technically meaningless. It's merely a springboard for consumer marketing.

The first DIGIC processor appeared in Canon cameras about ten years ago. The latest generation is the DIGIC 4. Canon amortizes the development cost of these chips by marching them down its extensive product line. A new DIGIC chip debuts in Canon's professional DSLRs, then appears in Canon's consumer-level DSLRs, and finally reaches end of life in Canon's point-and-shoot digicams.

Canon reveals little about these SoCs. The DIGIC 4's prime feature is a 14-bit analog-to-digital converter (ADC) that extends dynamic range by capturing 16,384 shades per RGB color channel. The older DIGIC III chip has a 12-bit ADC that captures only 4,096 shades per channel. Canon says the enhanced processing power of the DIGIC 4 enables additional features, such as automatic face detection, high-ISO noise reduction, shadow lightening, and high-definition video recording. In other words, it's a faster image processor.

By designing its own image processors and sensors, Canon can produce optimized designs that are difficult for smaller camera companies to match with off-the-shelf parts. And Canon has the large volumes necessary to sustain a custom silicon strategy. Nikon competes in the same manner, though Nikon cameras tend to use sensors from outside suppliers.

It's noteworthy that the historic transition from film to digital photography hasn't overthrown Canon and Nikon from their leading positions in the camera industry. Often, a major technology transition radically reshuffles an industry. In this case, the market leaders possess the technological firepower to defend their positions.

Below Canon and Nikon, however, the camera industry is roiling. Storied companies like Konica and Minolta have dropped out. Leica, which pioneered small cameras, has partnered with Panasonic to survive. Olympus and Pentax fight over table scraps of market share. Sony—perhaps the world's most famous name in consumer electronics—acquired Minolta's photographic technology and is making new cameras with Minolta lens mounts under the Sony name. Because a modern camera is essentially a mobile computer with a lens, a camera company without semiconductor expertise is severely handicapped. It's not surprising that Apple views custom silicon as a key ingredient in its growing line of consumer-electronics products.

Projects Must Be Pipelined

When a company commits to using custom silicon—whether the chip design is internal or outsourced under supervision—a long-term development pipeline is mandatory. It's not enough to design one good custom chip and rest on the laurels. Creating a complex SoC takes years, whereas product turnover in the consumer-electronics market is about one year. To keep the product line refreshed, two or three generations of an SoC must move through a staged development pipeline. In that way, an improved version of the chip can hit the market every year or two, and the pipeline allows some wiggle room for project delays.

Of course, multiple projects require multiple development teams. Some personnel can overlap as the projects move through different phases, but one team focused on one project isn't sufficient. This requirement may explain the pace of Apple's recent hiring. The expense of sustaining multiple teams shouldn't be an obstacle for Apple—a debt-free company that in 2008 reported a $4.8 billion profit on $32.4 billion of revenue. Currently, Apple has $29 billion cash in the bank.

To amortize development costs and keep the end products current, the product design teams must be pipelined, too. Although Apple likes to keep its product lines relatively sparse, those lines will probably expand somewhat to prolong the usefulness of the custom chips. Following the DIGIC example, an audio SoC that debuts in a high-end iPod Touch could eventually make its way into a low-end iPod Shuffle. A communications chip or applications processor that debuts in a higher-end iPhone may eventually appear in a future lower-end iPhone.

At the heart of today's iPhone is a custom ARM-based SoC reportedly developed by Apple and Samsung. Without this SoC, the iPhone probably wouldn't exist, or would suffer from such poor performance that it would fail. Although nearly two years have passed since Apple introduced the first iPhone—a wildly popular product—traditional cellphone makers have been slow to respond. It's a telling example of the limitations of relying on standard parts when a custom SoC allows one player to redefine a product category.

Which CPU Architecture?

It will be interesting to see which processor cores Apple uses in future SoCs. Right now, Apple favors ARM, and *MPR* thinks Apple will stick with ARM. There's no known technical reason to change architectures. Although the P.A. Semi engineers who joined Apple last year brought recent experience with the Power Architecture, that's not a compelling reason to switch. Indeed, some P.A. Semi engineers were on the StrongARM team at DEC, so they are no strangers to ARM, the world's most popular 32-bit microprocessor architecture.

Although ARM won't confirm it, *MPR* believes Apple is an ARM architectural licensee, not merely a core licensee. An architectural license gives Apple the freedom to design its own ARM-compatible microarchitecture, as Marvell and Qualcomm have done. An ARM architectural license is rumored to cost about $20 million, and the cost of developing a new microarchitecture is even greater. It's unlikely that

Apple would commit such resources if it weren't equally committed to the ARM architecture. A high-performance ARM-compatible processor tailored for Apple's needs would give Apple products an advantage over competitors using off-the-shelf ARM cores and standard parts.

Intel's Atom is frequently mentioned as an alternative to ARM. Certainly, Intel is itching to get the x86 into smartphones and other consumer products. Apple's iPhone would be a landmark design win for Intel. But a discrete microprocessor like Atom still burns too much power for a small smartphone. Although Atom is an impressive low-power processor by PC standards, it's not in the ARM class, and Intel's Atom chipsets are energy hogs. Intel lacks a licensable low-power x86 core—a formidable obstacle to designing the kinds of custom SoCs that Apple needs. (See *MPR 4/7/08-01*, "Intel's Tiny Atom.")

Intel recently took a step in the right direction by announcing a collaboration with TSMC to design Atom-based SoCs for favored customers. It falls short of ARM's open licensing model, but it opens the door for a wider variety of x86-compatible SoCs. (See *MPR 3/30/09-01*, "Intel Will Customize Atom.")

However, Apple may want more control over chip design for crucial products like the iPhone. Like almost everyone, Apple sees smartphones and mobile Internet devices as the future of personal computing. We suspect Apple isn't comfortable outsourcing its destiny to another company, especially one as territorial as Intel. If Apple worries about exposing proprietary intellectual property (IP) to outsiders, the Intel/TSMC collaboration is a less-than-ideal solution. The reluctance of Apple or its partners to reveal their IP to Intel may be as great as Intel's reluctance to share its IP (in the form of a licensable x86 core) with them.

Is Apple Reinventing Sun?

As Apple hires more chip designers, it's moving closer to becoming a vertically integrated company in the mould of Sun Microsystems. Like Sun, Apple already creates its own system software, application software, and system hardware. Critics say this resemblance doesn't bode well for Apple. It's a business model that seems dated, and Sun recently staggered into the acquisitive arms of Oracle.

But there are two important differences, both in Apple's favor. First, Apple isn't burdened with an in-house CPU architecture (e.g., SPARC) that demands relentless, expensive development to stay competitive. Apple is free to use any openly available CPU architecture and to switch architectures when necessary, as it has done twice with the Mac. (See *MPR 6/27/05-01*, "Apple Drops PowerPC for Pentium.")

Second, Apple focuses on consumer products, not servers. Although the consumer electronics market is even more competitive than the enterprise server market, it's also more fluid. In the past ten years, we've seen whole new product categories like digital cameras, MP3 players, portable video players, cellphones, smartphones, and netbooks soar to popularity. Whipping up frenzy over a hot product like the iPhone is easier than getting corporate IT buyers excited about a new blade server.

Figure 2 illustrates the ramp rate and seasonal cycles of Apple's iPod sales. In contrast, the server market is much

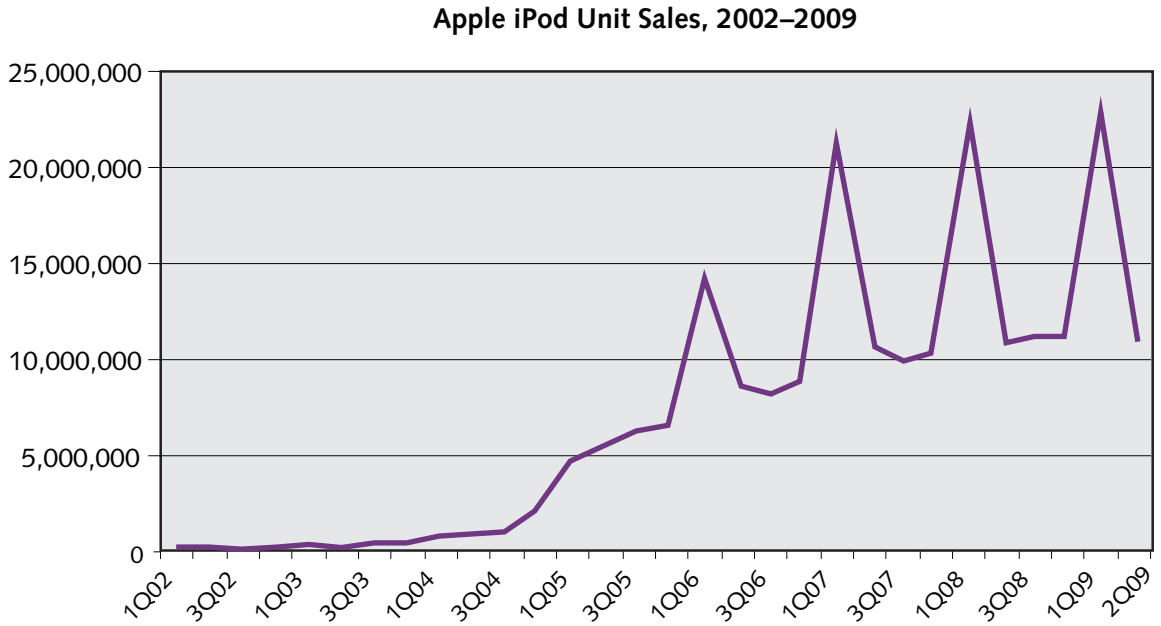

Figure 2. Apple's iPod unit sales, 2002–2009. Today, the iPod so dominates the audio market that it's easy to forget Apple wasn't the first to ship an MP3 player. Despite a late and relatively slow start, the iPod soared to popularity. This chart also reveals the seasonal pattern of iPod sales—a Christmas holiday cycle that's typical of consumer electronics. (Data source: Apple.)

For More Information

Wall Street Journal, 4/30/2009, "In Major Shift, Apple Builds Its Own Team to Design Chips"
By Yukari Iwatani Kane and Don Clark—*http://online.wsj.com/article/SB124104666426570729.html*

EDN, 5/1/2009, "Apple Looks 'Chipper': Recent Hires, Company Buys Suggest Apple May Be Readying Chip Design Unit"
By Suzanne Deffree—*www.edn.com/blog/1750000175/post/450044045.html*

EDN, 1/28/09, "Ex-IBM Engineer Papermaster Cleared to Work at Apple"
By Suzanne Deffree—*www.edn.com/blog/1750000175/post/810039881.html*

more static and is pinned to price/performance. Whereas the consumer market is a lively free-for-all, the server market is a grueling siege.

Some observers wonder if Apple is surfing a new tide of Makimoto's wave. Described in 1991 by Sony engineer Tsugio Makimoto, the "wave" is a cyclical trend that alternates between standard silicon and custom silicon at approximately ten-year intervals. In the 1990s, Makimoto's wave surged in favor of custom ASICs. According to some estimates, there were more than 10,000 new ASIC designs per year. In the 2000s, Makimoto's wave receded in favor of standard silicon. The sharply rising costs of nonrecurring engineering and wafer fabrication have considerably dampened the fervor for custom ASICs and SoCs. If Apple is moving toward in-house chip design, does it signal the next reversal of the wave?

Probably not. Despite the deflationary gravity of the global economy—a discouraging force that everyone prays is temporary—the costs of engineering and fabrication are still rising. Nothing fundamental has changed (yet). More likely, Apple, not Makimoto's wave, is entering a new phase. As Apple becomes a consumer-electronics company like Sony or Canon, it needs custom silicon to further differentiate its products from those of competitors. Without differentiation, Apple can't charge higher prices and reap larger margins than the bottom feeders.

So Apple is probably swimming against Makimoto's wave. And Apple has the strength and motivation to do it. Apple is accumulating the wealth and experience needed to design custom chips that will give its consumer products an edge over competitors using standard silicon. Other companies may not be able to afford that strategy or don't consider those product lines important enough to justify the effort. Therefore, Apple is an example only for companies whose business models, fortunes, and markets resemble Apple's.