How Advanced RISC Machines' Disruptive Innovation May Cost Intel Its Lead in the Mainstream Processor Market

Eric Mesa

26 October 2011

In *The Innovator's Dilemma* Clayton Christensen explores why established companies seem unable to cope with disruptive technologies and tend to cede their leading positions to newcomers, assuming they even survive the transition to the new technology. His conclusion is that companies are unable to adapt because disruptive technologies do not fit within the established companies' established markets. At times the market for the disruptive technology does not yet exist. Even if the market does exist, it does not make sense for the established company to pursue the new markets the disruptive technology would serve because the margins are usually much lower than with their current technology. All companies have fiduciary responsibilities to their investors to grow larger as quickly as can be responsibly done and these disruptive technologies with their low margins cannot meet those financial goals. A summary of the disruptive technology time line Christensen constructs would go as follows. A new company develops disruptive technology. At first the disruptive technology is inferior to the current market leader in some key way. Because of competitive forces, the companies innovate the mainstream technology beyond what the average person needs. As the disruptive technology companies innovate to compete in **their** market, they eventually reach a level where they can be used by the mainstream market. They have all the old features the consumers needed in addition to their disruptive technology and this leads consumers to abandon the old technology.[6] This paper will explore how the ARM central processing unit (CPU) is a disruptive CPU architecture that will potentially unseat the mainstream x86/x86-64 CPU architecture.

At the outset of the personal computer (PC) revolution nearly every computer company not only had their own operating system, but also had their own custom CPU. One of these companies was a British computer company, Acorn Computers, which, in 1985, created the ARM CPU for their latest computer.[5] While Intel had chosen for their x86 CPUs to be Complex instruction set computing (CISC), Acorn chose for their CPUs to be Reduced instruction set computing (RISC). RISC is a simpler architecture which uses fewer transistors, reducing design costs and consuming less power. "The ARM (Acorn RISC Machine) CPU had performance comparable to the Motorola 68000 used in early Macs, but with half as many transistors, which meant it consumed less power and generated less heat."[14] However, RISC CPUs require more instructions per operation and, therefore use more RAM. At that time, RAM was very expensive while there was no advantage to consuming less power.[8] CISC designs also require less complex compilers to turn the high level code to machine code. Additionally, CISC designs have "fast floating point performance and improved cache performance."[10] Cache, the fastest memory available to CPUs is very expensive so there was a huge benefit to being able to use a smaller cache. These various properties gave the advantage to Intel. This, together with Windows targeting the Intel platform, left Acorn computers as the only PCs running on the ARM architecture. In time all PCs other than Apple's came to run on Intel (or Intel clone) CPUs. (And even Apple eventually switched from PowerPC CPUs to Intel CPUs in the 2000s)

As PC customers continued to demand faster and faster processors Intel and its competitors complied by increasing processor speed. This began to cause a differentiation between chips used for PCs and chips used for more general purpose appliances. Processor speed increases were achieved by the sustaining innovation of increasing the numbers of transistors on a chip. Transistors are small electrical switches and the process of switching them on and off requires electricity to do the work and, as in most times when work is done, causes heat to be released as a by-product. Therefore, as processors increase in speed by increasing the number of transistors on a chip, they also increase in thermal output and power consumption. For example, the Pentium 4 630 consumes 86 Watts of power while idle and 155 watts under load. The next generation CPU, the Intel Core 2 Extreme QX6850 consumes 94 watts while idle and 195 watts while under load.[17] Since the main PC platform was the desktop, this increase in power consumption did not matter to the user. The market forces they exerted on Intel were demanding faster and faster processors no matter the power consumption. The increased heat these faster processors produce is one of the reasons the processors released in the last few years (at the time paper is being written) have been multi-core systems. A 2007 article on Tom's Hardware stated, "The Pentium 4 500 series was available at up to 3.8 GHz (model 570) while the 600 series maxed out at 3.6 GHz (Pentium 4 660). Plans to cross the 4 GHz line were discussed, but heat dissipation put an end to that idea, so Intel worked towards getting the first dual core processors ready."[18] This extra heat and power consumption were a drag on laptop battery lifetimes, but as Windows was only available on the Intel platform, users did not have a choice. Just as Christensen predicted, Intel's continuous sustaining innovation meant they were over-providing for the customer's needs and consumers were now ready to consider features other than processor speed.[6] This oversupply of processing power has become especially pronounced in recent years as a lot of compute power has been offloaded to Internet-based programs. Customers are now ready for something disruptive.

Throughout the 1980s and 1990s consumers were moving away from other operating systems such as Amiga, Acorn, etc and moving towards Windows. That meant using an Intel (or clone) machine. Meanwhile, as Christensen predicts would happen, Acorn looked for new markets that would appreciate its low heat and low power consumption CPU [6] Around that time Apple was looking for a low power, low heat chip for its innovative personal digital assistant (PDA), the Newton. They first tried to use the AT&T processor named Hobbit for their PDA. However, it was not powerful enough for only one processor to be used in the PDA. When Apple found out about the ARM line of processors, it bought a stake in ARM and Acorn spun off Advanced RISC Machines as a new company to manage the design.[14] ARM found a partner appreciative of their chip's advantages. Although the Netwon turned out to be a market failure, Advanced RISC Machines had found a market that appreciated their strengths: the portable and embedded market. The portable market needed low power consumption so that the batteries could last a reasonable amount of time before recharging. (Even more important with the battery technology of the early 1990s than it is today) And embedded devices are usually very small, and, thus, require as little heat dissipation as possible in order not to ruin the other chips in the device. Advanced RISC Machines had also found their new business model. They would create the basic intellectual property (IP) behind their chips and then license it to other companies to integrate into their own products.[1]

Of course, as Christensen tells us in *The Innovator's Dilemma*, companies often know about disruptive technologies and develop their own version.[6] Intel was not blind to the advantages of the RISC CPU, especially as RAM prices began to fall. So Intel developed the i860. However, Christensen tells us that under most circumstances a company cannot support both sustaining and disruptive technologies.[6] One market or the other must prevail within a company. Intel's then-CEO Andy Groves recalled that challenge in his book *Only the Paranoid Survive*:

> We now had two very powerful chips that we were introducing at just about the same time: the 486, largely based on CISC technology and compatible with all the PC software, and the i860, based on RISC technology, which was very fast but compatible with nothing. We didn't know what to do. So we introduced both, figuring we'd let the marketplace decide. However, things were not that simple. Supporting a microprocessor architecture with all the necessary computer-related products – software, sales, and technical support – takes enormous resources. Even a company like Intel had to strain to do an adequate job with just one architecture. And now we had two different and competing efforts, each demanding more and more internal resources. Development projects have a tendency to want to grow like the proverbial mustard seed. The fight for resources and for marketing attention (for example, when meeting with the customer, which processor should we highlight) led to internal debates that were fierce enough to tear apart our microprocessor organization. Meanwhile, our equivocation caused our customers to wonder what Intel really stood for, the 486 or i860?[11]

So, when confronted with the fact that everything up to that point was compatible with their x86 chips, which had retained backwards compatibility across the entire line, or starting from scratch, they decided to stick with the 486. After all, AMD was creating clone chips and they might have lost the entire market to AMD if the vendors hadn't also jumped over to the new architecture. So the market kept Intel from moving over to a RISC CPU at that time. Christensen mentions that these market forces often keep companies from adopting disruptive technology. [6] But it would not be the last time that Intel worked with this architecture.

In the mid-1990s the PDA market that Apple had attempted to spearhead was finally booming. Digital Equipment Corporation (DEC) licensed ARM technology to create the StrongARM processor to power these new devices. ARM's continued innovation had led to a chip that was perfect for this market and which could not be matched by Intel's x86 line. The StrongARM, based on the ARM4, ran at 160 MHz while consuming only 500mW of power. Montanaro, et al wrote, in a paper about the StrongARM, "This represents a significantly higher performance than is currently available at this power level."[16] In 1997 as the result of the settlement of a lawsuit between DEC and Intel over patent infringement, Intel acquired the right to produce and sell DEC's chips, including StrongARM. [15] Intel then worked with Marvell to develop a follow-on chip called XScale, based on the ARM5 technology. However, in 2006, Intel once again divested themselves of this technology by selling the entire stake in the venture to Marvell for $600 million. AMD also rid itself of similar technology it had acquired. Intel stated it wished to do this in order to focus on its x86 chips. [7]

So why would Intel continuously discard this RISC technology over it's older CISC-based technology? Is it just because everything up until now has run on the x86 architecture? Christensen tells us to take a look at the profit margins because disruptive technology usually has low profit margins that will not allow the company to meet its growth targets for its investors.[6] Indeed, in the 2006 deal Marvell cited a reason for Intel's sale of XScale technology: "$100 million per quarter, albeit with low margins caused by production constraints"[7] In 2008 an article on Tom's Hardware stated that "Intel Builds Atom CPU For $8, Sells It For Up To $135" giving it margins of $127.[12] ARM, on the other hand, makes less than $10 in licensing costs according to Gassée who states that "Intel might harbor the usual cannibalization fear: cheaper ARM processors might displace x86 CPUs"[9] And as commodity hardware has made its way into supercomputers, Intel is chasing higher profit margins by invading the space once held by supercomputer manufacturers like Cray. It does not have time to concern itself with these low margin technologies that have nothing to do with its market. But should it?

As mentioned above, Advanced RISC Machines does not manufacture its own chips. Consistent with the properties of a disruptive technology company as described in *The Innovator's Dilemma*, it essentially out-sources all production by simply being a license holder.[6] Therefore, unlike Intel, ARM does not need to maintain chip fabrication facilities, allowing it to better deal with the lower margins. This also allows ARM to change their chip designs to match market needs without having to worry about discarding expensive fabrication equipment. Additionally, since it does not create any chips, nearly all of its design is done in software, decreasing costs to the company. Up until now Intel has not had to worry about ARM eating into its x86 market. As with all disruptive technologies, it was not powerful enough in the established market (power desktop PCs) so it found its own market. It started off with PDAs and other embedded technology. According to ARM's Company Overview for investors, their "designs are used in more than 95% of the world's mobile phones"[1]. Browsing their corporate website shows their chips are being used in TVs, routers, car entertainment devices, digital cameras, and as control chips on hard drives and solid state drives. Graphics company Nvidia has based their new Tegra 2 system on a chip (SOC) on ARM technology.[3] A SOC combines the CPU and graphics processing unit (GPU) onto one chip and presents direct competition to Intel's x86-based SOC offerings. As I mentioned above, and consistent with Christensen's observations, as Intel has made their processors more powerful than customers demand, power consumption has moved to become the new desire of customers who increasingly use mobile devices as their main computing platform.[6] Also consistent with Christensen, ARM has greedily eyed the higher margins in the PC market. As reported earlier this year, there are rumors that Apple will be replacing Intel CPUs with ARM CPUs for its laptops and may do so for its desktops as well.[4] One of the trends that worked against ARM in the late 80s and early 90s was that all the operating systems were designed for Intel (and clone) chips. That trend is changing as Microsoft announced that Windows 8 would run on ARM on tablets. And, although, they later backtracked a bit and said that only certain Microsoft programs would be compiled for the ARM chip it was still blow to Intel.[13] ARM also has support of the Linux Kernel and many of the major Linux distributions fully support ARM. [2] In fact, the author of this paper owns a computer running the Arch Linux distribution on an ARM chip created by Marvell and its performance is nearly indistinguishable for most basic computing tasks from its x86 brethren while consuming two orders of magnitude less power. With the current "green" technology trends power consumption may become important to desktop users as well.

I believe the biggest takeaway from studying the company is the market flexibility they gained by being a licensor of technology. This allows them to excel at design and allowing others to excel at integration, marketing, and all the other aspects of bringing a technology to market. Each company is free to integrate the chip as best fits the product they are designing. I believe it is such a powerful way to come up with new technology – leveraging the strengths of many companies – that it may well become the new paradigm, an extension of where business have already been going with the way they source parts globally. Innovators will design core bits of technology that will be combined by other companies to create the final product the user sees.

So, at the time of this paper's completion, it appears that ARM is on the cusp of making the jump from the market it needed to create as a disruptive technology to the market Intel has dominated. It has already gained near 100% market share on portable devices. It is too early to tell if ARM will be successful or if Intel will be able to innovate to match ARM's benefits and stem the losses. Christensen's examples in *The Innovator's Dilemma* do not place the odds in Intel's favor, but it is not completely unprecedented for the company to be able to survive. Intel is said to finally be pursuing chips that may be able to compete with the power draw and heat dissipation of ARM's chips.[4] Additionally, it is the niche Linux OSes that have the best support for ARM such that the experience would be exactly the same for the user regardless of CPU. Microsoft and Apple have been making motions in the direction of ARM, which would be a huge game changer. Only time will tell. However, ARM certainly has the hallmarks of a disruptive technology company that may win. It is smaller than Intel and AMD, its technology is cheaper than the competition, it provides a need the market is demanding now that it is sated on processor speeds (ARM's traditional disadvantage is no longer a disadvantage), and it has caught up on processor speeds to a level that is acceptable to the customers. By organizing itself as a research and licensing company it has virtually eliminated the cost of being nimble while Intel maybe have to change its fabrication processes and equipment in order to produce chips that compete with ARM. Nearly the entire trajectory of Advanced RISC Machines has tracked with that of a disruptive technology company as described in *The Innovator's Dilemma*.

References

[1] “ARM Company Overview.”
http://ir.arm.com/phoenix.zhtml?c=197211&p=irol-homeprofile, 2010.

[2] “Linux Support for the ARM Architecture.”
http://arm.com/community/software-enablement/linux.php, 2011.

[3] Brenesal, Barry. “An Interview with Nvidia’s Nick Stam & Sridhar Ramaswamy.” *Computer Power User* 11.9 (2011): 102–103.

[4] Burt, Jeffrey. “Apple to Dump Intel for ARM in Laptops: Report.”
http://www.eweek.com/c/a/Desktops-and-Notebooks/Apple-to-Dump-Intel-for-ARM-in-Laptops-Report-815172/,
2011.

[5] Business, Telegraph. “History of ARM: from Acorn to Apple.”
http://www.telegraph.co.uk/finance/newsbysector/epic/arm/8243162/History-of-ARM-from-Acorn-to-Apple.html,
2011.

[6] Christensen, Clayton M. *The Innovator’s Dilemma.* HarperBusiness, 2002.

[7] Devices, Linux. “Intel ditches mobile phone processors.”
http://www.linuxfordevices.com/c/a/News/Intel-ditches-mobile-phone-processors/,
2006.

[8] Faculty, Stanford CS. “RICS vs CISC.” http://www-cs-faculty.stanford.edu/ eroberts/courses/soco/projects/2000-01/risc/risccisc/,
2000.

[9] Gassée, Jean-Louis. “Intel’s bold bet against ARM: visionary or myopic?”
http://www.mondaynote.com/2010/06/27/intel%E2%80%99s-bold-bet-against-arm-visionary-or-myopic/,
2010.

[10] George, Alan D. "An overview of RISC vs. CISC." *System Theory, 1990., Twenty-Second Southeastern Symposium on* : 436–438.

[11] Grove, Andrew S. *Only the Paranoid Survive.* Doubleday Business, 1996, hardcover ed.

[12] Gruener, Wolfgang. "Intel Builds Atom CPU For 8, *SellsItForUpTo*135." http://www.tomsguide.com/us/intel-atom-silverthorne,news-937.html#, 2008.

[13] Gruman, Galen. "Windows 8 on ARM chips: It was too good to be true." http://www.infoworld.com/d/microsoft-windows/windows-8-arm-chips-it-was-too-good-be-true-173265, 2011.

[14] Hormby, Tom. "The Story Behind Apple's Newton." http://lowendmac.com/orchard/06/john-sculley-newton-origin.html, 2006.

[15] Luening, Erich. "Intel, Digital settle suit." http://news.cnet.com/2100-1023-204668.html, 1997.

[16] Montanaro, James, Witek, Richard T., Anne, Krishna, Black, Andrew J., Cooper, Elizabeth M., Dobberpuhl, Daniel W., Donahue, Paul M., Eno, Jim, Hoeppner, Gregory W., Kruckemyer, David, Lee, Thomas H., Lin, Peter C. M., Madden, Liam, Murray, Daniel, Pearce, Mark H., Santhanam, Sribalan, Snyder, Kathryn J., Stephany, Ray, and Thierauf, Stephen C. "A 160-MHz, 32-b, 0.5-W CMOS RISC Microprocessor." *IEEE Journal of Solid-State Circuits* 31.11 (1996): 1703–1714.

[17] Schmid, Patrick. "Power Consumption Tests." http://www.tomshardware.com/reviews/intel-cpu-power-consumption,1750-9.html, 2007.

[18] ———. "Processors at 3.0 GHz." http://www.tomshardware.com/reviews/intel-cpu-power-consumption,1750-2.html, 2007.