Exponential Innovation

Understanding the Computer Chip Market and the Symbiotic Relationship Between the Processor and Software Innovation

Authors:

Braden Cox & Nora von Ingersleben

1946

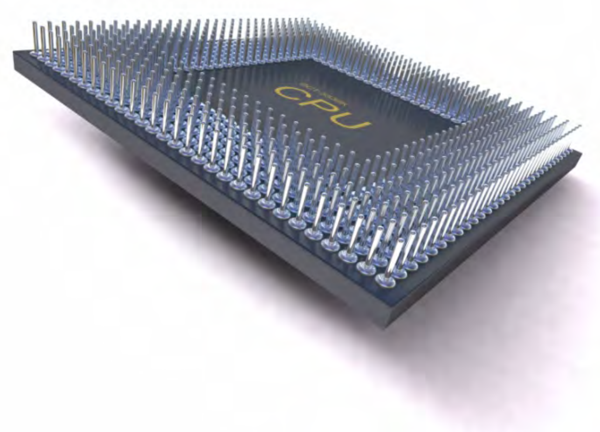

2009

Executive Summary

For more than forty years, the software development community has leveraged huge increases in computer chip performance and steady drops in chip prices to give consumers what they want: better products at lower prices.

In 1965, Intel co-founder Gordon Moore predicted that the number of transistors on a chip (essentially the speed of the chip) would double every two years - a pace of innovation unmatched by any other industry.

This American success story of rapid product improvement at steadily falling prices is no accident - companies like AMD, IBM, Intel, Sun and Texas Instruments annually invest tens of billions of dollars to ensure that software developers have the tools we need to deliver for the American consumer. The results are clear -- from 2000 to 2008, processor performance increased 28 times while prices fell by 50 percent. Just looking at Intel, the average price of an Intel microprocessor for a PC has dropped by 60 percent over the past ten years.

The successes of the software industry --built on the innovations of the computer chip manufacturers—are innumerable, but a few examples include:

- **Medical Imaging -** Breakthroughs in processing power have enabled the 3D high definition imaging and multitouch interfaces that make software like InterKnowology's VirtuView possible. Using this technology, doctors can be better prepared than ever for complicated heart procedures as they can investigate potential problems and annotate inside and outside the heart by placing stents and marking lesions.

- **Digital Animation -** When the digital animation studio Pixar was founded in 1986, the processing power necessary to produce its first hit movie "Toy Story" wasn't even available. It took another five years before the technology was cost effective enough to start working on their first full-length movie. Today, even low budget Saturday morning cartoons can take advantage of photorealistic digital animation.

- **Digital Imaging for Consumers -** Just ten years ago, the average PC could not edit a picture from your 15 megapixel camera let alone play video of your kids from your high definition video camera. Today, you can shoot and edit video on your iPhone, and watch HD movies and surf the net simultaneously on your PC.

The chip market is clearly working well for software developers and consumers, and it appears any potential competitive issues have been resolved by the recent private settlement between AMD and Intel. Therefore, ACT is extremely concerned by the potential of additional government intervention in the chip market. The health and vibrancy of the computer chip market are beyond debate. Therefore, we in the software development community ask government regulators to proceed extremely cautiously and avoid any actions that may reduce incentives for innovation or result in higher chip prices for consumers. Such regulatory action presents a clear danger to our businesses, and could have the perverse effect of stifling innovation, raising prices, and costing American jobs.

Exponential Innovation

Moore's Law and Rapid Growth of the Chip Industry

- ❑ Moore's Law states that the transistor count on a chip will double every two years, leading to exponential increases in computing speed and power

- ❑ Achieving Moore's Law is not easy, however – it requires massive investments in R&D, constant innovation and continued partnership with software developers

Computers have come a long way since the ENIAC, the world's first general-purpose electronic computer that was introduced in 1946 and consisted of 17,468 vacuum tubes, 7,200 crystal diodes, 1,500 relays, 70,000 resistors, 10,000 capacitors and around 5 million hand-soldered joints. It weighed 27 tons.

Transistors began replacing vacuum tubes in electronic equipment during the 1950s. By 1965, Intel co-founder Gordon Moore observed that the number of transistors that could be placed inexpensively on an integrated circuit had increased exponentially, leading to ever higher processor speeds.[1] Moore initially believed the density of transistors would double every 18 months, but later concluded that the time frame was closer to two years.

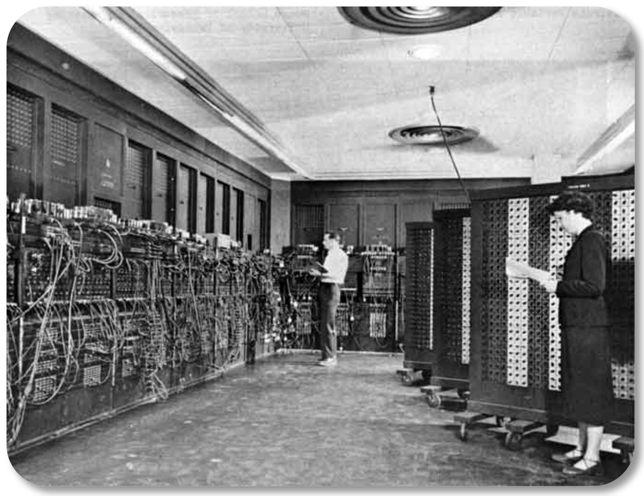

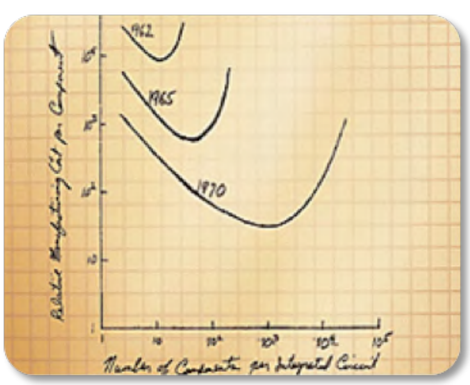

Original "Moore's Law" Graph drawn by Gordon Moore

Moore was so remarkably accurate that his prediction became known as "Moore's Law." Unlike a scientific law that expresses a universal fact of the physical world, Moore's Law expressed the aspiration of human ingenuity and served as a reliable predictor for the computer industry.

The silicon industry moves extremely fast and aggressively in order to "keep up with the Moore's." Achieving the technological breakthroughs for doubling speed requires massive investments in research & development (R&D), innovations delivered at a breakneck pace, and interactions with software developers and manufacturers who take advantage of new processor features.

Through billions of dollars in investment, chipmakers such as AMD, Intel, Texas Instruments (TI) and IBM achieve Moore's Law. Without the spending, Moore's observation would fail.

Companies continue to make the massive investments necessary to double the number of transistors every two years. This tireless march led Moore himself to admit that he was "perpetually amazed" at how technologists continue to push the limits.[2]

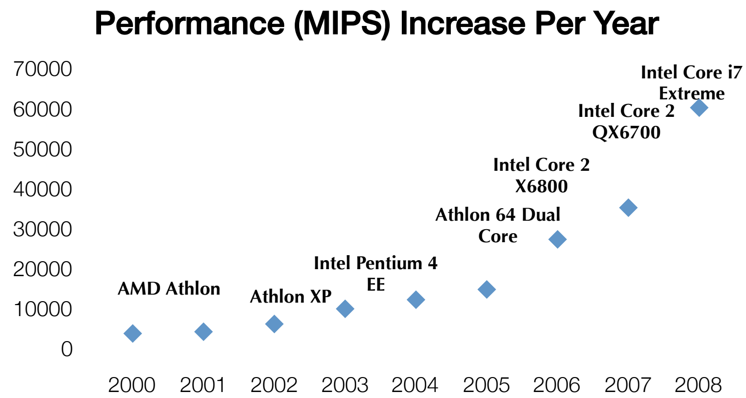

The results are amazing. Processor speeds, memory capacity, sensors, and even the number and size of pixels in digital cameras have all improved at (roughly) exponential rates. This graph demonstrates just how rapidly chip performance has increased over the past 8 years.

Compare a 1989 computer with a modern computer, and it is easy to see how computers have progressed.

In 1989, Fort Worth-based Tandy Corporation offered what it advertised as "the most powerful computer ever." The Tandy 5000 MC Micro Computer had a 20 MHz Intel 80386 Intel microprocessor[3], and offered 2MB Random Access Memory (RAM). However, the Tandy 5000's technology and performance came at a high price of $8,499 ($15,000 in today's dollars). Moreover, the monitor and the mouse were not even included in this price and needed to be bought separately.

Today, $15,000 would buy six Dell XPS 630 desktops with a 3.0 GHz Intel Core 2 Extreme QX9650 processor, a 21.5 inch Full HD Widescreen Monitor and more than 4GB of memory per machine.

And it is not just computers that have made great strides—the processing power of a common game cube today, priced at $299, dwarfs the capacity of a $2,000 PC of five years ago.[4]

Of course, Moore's law doesn't happen automatically. In the next section, we will detail how aggressive R&D spending and continuous innovation enables chipmakers to keep achieving Moore's Law.

Keeping Up With The Moores

Massive Investments in Innovation & Falling Prices

- ❑ Keeping Pace with Moore's Law Requires Massive Investments in R&D and Rapid Technological Advancement
- ❑ R&D Spending in the Semiconductor Industry Exceeds 15% of Revenues and is Among the Highest of any Industry

The semiconductor industry is in a constant technological arms race. For the past forty years the chip industry has been delivering on Moore's promise through aggressive research and development and continuous innovation.

Chipmakers spent $45.7 billion in 2007 on R&D and related engineering activities.[5] With R&D investments greater than 15 percent of sales, chipmakers rank among the top R&D investors in a cross-industry comparison.[6] In absolute numbers, the semiconductor industry is expected to spend nearly $35 billion in 2009, placing it below the pharmaceutical industry but above the software industry, which includes such R&D heavyweights as Microsoft and Google.

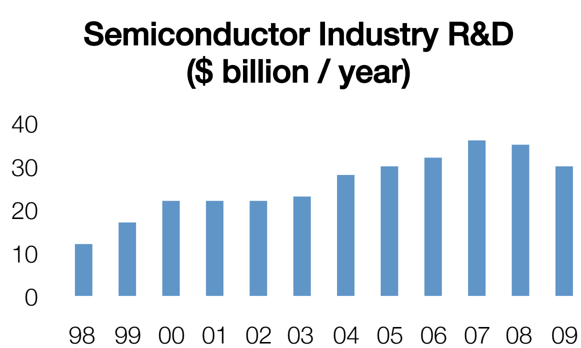

On a per-company basis, there are a number of notable spenders on R&D. In 2008, Intel spent $5.7 billion on R&D, AMD spent $1.8 billion, TI spent $1.9 billion.[7] Other players in the chip market include IBM, who dedicated almost half of its $6 billion R&D budget toward chip development.[8] Sun Microsystems—developer of the SPARC processor—invested $1.8 billion in R&D in 2008. Motorola invested $4.1 billion.[9]

All this R&D spending goes to good use in keeping pace with Moore's Law. And as we detail in the next section, chip makers have delivered on a number of innovations.

Innovation Produces New Processors

Innovation is a buzzword for today's knowledge economy, and for good reason. Innovation occurs by understanding the implications of invention and developing ideas into something of value. Processors have seen innovations in clock speeds, multiple computing cores and lower power consumption. As a result processors have become faster processors and computing capability has increased.

Clock Speeds

Clock speed is the measure of cycles per second (measured in hertz) for the maximum frequency a processor can toggle between a *zero* and *one* state. The higher the clock speed, the faster the processor.

Historically, with each successive generation of processors, clock speeds increased along with the number of transistors. The x86—the most popular instruction set for computer processors—was introduced by Intel in 1978. Intel 8086 was a 16 bit processor with 29,000 transistors, and a variant of it (the 8088) powered the original IBM PC at a clock speed of 4.77 MHz.

For years processing power, measured in millions of instructions per second (MIPS), steadily rose because of increased transistor counts and clock speeds. Much of this has been achieved through new manufacturing technologies. While transistor size is shrinking (50% every two years), capacity is doubling.

However, clock speed is not the only measure of processor performance. Other innovations now play a huge role. Today's processors see performance improvements through innovations in pipelining, instruction sets, and the development of multi-core processors.

Multiple Cores

Gone are the days when increasing clock speeds would have been a guaranteed way to make software run faster. During the 1990s, Intel, TI, IBM and AMD steadily increased processor clock rates into speeds measured in GHz, but these fast processors were often idle because other system components could not keep pace. The processors also ran very hot, creating significant power consumption and heat problems, particularly for mobile machines.[10]

Today's chips therefore take a different approach to increasing processor speeds. Chips have processors with multiple cores to separately but simultaneously handle independent tasks. A multi-core processor is essentially a processing system composed of two or more independent cores (or CPUs).

Beginning in 2004, the major processor manufacturers started turning to multicore architectures, allowing them to more effectively manage power and keep system costs down while providing the high processing speeds needed by new applications. This allows processor manufacturers to continue delivering the exponential performance growth suggested by Moore's Law, and has enabled improvements in virtualization and graphically-intense technologies.

Lowering Power Consumption

Power consumption has long been a bane to users of portable devices—who hasn't experienced a cell phone or laptop running out of battery life during the middle of something important? While researchers continue to stretch the limits of battery technology, processors have also advanced to become more efficient and better use battery power. The trend began with Intel's 2003 introduction of the Pentium M processor, which was the first PC microprocessor designed "from the ground up" as a notebook PC processor. This processor and its successor products opened the way to thin and light notebook PCs with long battery life and performance that rivals many desktop PCs. Notebook PCs, which in 2003 accounted for a small fraction of the PC market, now outsell desktops.

Other companies have also introduced innovations

Survey of ACT members

August 2009

"Multicore Makes Programming Harder, But When My Apps Leverage it, They Can Do More."

Jeffrey Richter, Founder of Wintellect

In a recent survey of ACT's members, 58% of the respondents identified mulitcore technology as the processor advancement that has most improved their software products. In the past, most applications were based on a single-threaded architecture, and applications depended on clock speeds to increase performance. Today's applications are increasingly designed to take advantage of the increased multiple computing engines available on new chips, resulting in performance improvements when run on multi-core systems.

for small form factor computing. Texas Instruments has developed what it calls “SmartReflex technologies”—a range of “intelligent and adaptive hardware and software techniques that dynamically control voltage, frequency and power based on device activity, modes of operation and temperature.”[11] Samsung recently introduced a new mobile processor core implementation of the ARM Cortex-A8 processor architecture, in what it describes as a “45 nanometer (nm) Low Power (LP), low leakage process technology.”[12]

Increasingly, power consumption has also become a major issue for servers and other computers that power massive cloud computing and data centers. A large data center running at full capacity uses the same energy as 24,000 homes.[13] Roughly 50 cents is spent on energy for every dollar of computer hardware.[14]

Chip manufacturers have tried to improve chip performance while cutting every watt possible in order to reduce the power draw of data centers. Advances such as chip multi-threading, slower disk drives, and automated power-down technology are options for driving down the power consumption of datacenter systems. AMD and Intel have developed processors that manage power consumption by placing systems into a lower power state when the system workload decreases.[15]

In the next section, we will discuss the forces that drive companies to keep on investing in research and innovation.

The Drive to Innovate

A Dynamic Market and Strong Competitors Push Innovation

- ❑ Intense competition forces chip makers to invest and innovate, particularly as new competitors emerge from converging markets
- ❑ A vigorous and dynamic market rewards innovation and lower prices with increased market share

Moore's Law is not being met simply because a bunch of engineers have the need for speed. Users always want their computers to do more and do it faster; otherwise, there is no compelling reason to replace the machine they already have.

As former Intel CEO Craig Barrett pointed out in a 2009 *Wall Street Journal* article:

> Every few years some company will say, 'What's with the pell mell rush to improve our technology every two years? Let's slow down to, say, four years, and only have to invest half as much capital.' It always sounds like a cool idea, and it always ends up with that company losing market share.[16]

In this section, we analyze the competitive forces at work that force processor makers to respond to consumer demand, innovate or lose market share, and keep reducing prices.

Customer Demand

At the heart of chip manufacturers' efforts to keep pace with Moore's Law are the demands from their customers.

Original Equipment Manufacturers (OEMs), such as Dell, HP, and Sony, are the primary customers for chip manufacturers. OEMs place ever-increasing demands on chip manufacturers so that the OEMs can provide compelling new machines to consumers by delivering new capabilities, speeds, longer battery life, and lower costs. In order to best compete against each other, OEMs play chip manufacturers off of one another and typically keep purchase agreements to quarterly contracts. The longest agreements tend to be for a year. Even these short-term contracts are constantly at risk because OEMs regularly try to renegotiate terms due to constant technological improvements and changing competitive conditions. This gives OEMs massive power to influence the chip market.

Software developers are also an important customer and innovation driver for chip companies. In order to deliver faster versions of existing applications and radical new innovations, developers demand faster, improved processors. Perhaps the best example of this part of the demand side equation is Pixar animation studios. At the time of Pixar's founding in 1986, animator John Lassiter already had visions for the kind of animations they wanted to create through their software, but the processing power simply didn't exist at reasonable prices. It took years before Pixar's proprietary animation technology could be properly used to create full-length feature films.[17] Even when Toy Story was introduced in 1995, it took 117 Sun Microsystems computers and 70 Silicon Graphics machines to create and edit the movie.[18]

> "If I knew in 1986 how much it was going to cost to keep Pixar going, I doubt if I would have bought the company…The problem was, for many years the cost of the computers required to make animation we could sell was tremendously high. Only in the past few years has the price come down to the point that it makes business sense."[19]
>
> *Steve Jobs, CEO of Apple:* 1995

Finally, users are the ultimate customers in the processor distribution chain. Users are always looking for machines that enable them to do new things, do existing tasks faster and less expensively, and to be more mobile.

Competition is Fierce

The market for microprocessors is hyper-competitive. Chipmakers drive the market forward through massive R&D investments and constant innovations. Software firms take advantage of the pressures on manufacturers to offer new features, improve quality, increase computing speed, lower power usage, and drive down overall costs just to remain competitive.

Chipmakers compete against each other in ways that benefit product development and innovation. As AMD notes on its website, competition results in more innovation:[20]

> When Intel came under competitive pressure from AMD, the result was an accelerated decline in the price of microprocessors that translated into productivity acceleration.[21] There was a shift in the product cycle from three years to two years, a consequence of intensifying competition in the semiconductor market.[22]

And this competition is fierce. Remember the IBM PowerPC? As recent as 2003, Apple CEO Steve Jobs said that "[t]he PowerPC G5 changes all the rules. This 64-bit race car is the heart of our new Power Mac G5, now the world's fastest desktop computer. IBM offers the most advanced processor design and manufacturing expertise on earth...."[23] Yet Intel and AMD kept up the pressure and by 2005 IBM had fallen behind, and Apple made the switch from IBM's PowerPC processors to x86-based processors.

While no longer a supplier to Apple, IBM is still a competitor—after all, only four years have passed and IBM still produces chips. Its most recent processor for the server market is the Octo-Core Power 7, which competes against other high-end processors. According to one market watcher at Insight64, "[The] Power7 will be the fastest processor around, probably faster than Intel's Nehalem in some benchmarks."[24]

Competition is also Converging

Anyone with a smartphone can observe that convergence is here today. The processor market is evolving, and as mobile, netbook and other devices converge, chipmakers are increasingly invading each other's "turf" to compete against each other. Software developers welcome this converging competition.

The processor industry is commonly segmented into desktop, server, mobile, and embedded markets. But strict market segmentation overlooks the competition that chipmakers face from other

companies that are not direct competitors today. Markets are often "contestable" because they are constrained by potential entrants (ie. possible competitors) in addition to current participants.

There are many industry players involved with the design and manufacture of semiconductors. As smartphones and other mobile devices become more PC-like, vendors such as ARM, Marvell, Qualcomm, Samsung and Texas Instruments produce more powerful processors. At the same time, PC processor vendors such as Intel and VIA are creating improved chips for ultra-mobile PCs, netbooks, smartphones, and other portable and mobile devices.

The market is reacting quickly to this convergence. ARM has traditionally been focused on low-power processors for mobile devices, but recently introduced dual-core, quad-core, and eight-core processor designs "aimed at everything from netbooks to servers."[25]

In addition, Qualcomm CEO Paul Jacobs recently described how his company is embracing convergence. In an interview with the Financial Times, he stated:[26]

> We [Qualcomm] almost look at it as a perfect storm…The benefit for us of being the leading chipset manufacturer is that we can afford to do a lot of investment in R&D – 20 per cent of our revenues go into R&D. That allows us to stay ahead on the radio technology side and we are also able to put a very large investment into the microprocessing technology.

Even Apple is described as a future rival to Intel and AMD. Apple acquired chip design firm P.A. Semi in 2008, which is capable of designing chips for iPhones, iPods, and tablets.[27] Intel will soon be producing a more power efficient Atom chip that is targeted for smartphones, mobile Internet devices, and tablets.

As the above indicate, if Intel or AMD stop innovating in processors, there are many others to come in and take market share. Simply put, if a chipmaker stumbles in today's fast-paced market, it will likely fall.

Innovation Produces Market Share

The processor market responds as any marketplace should—those who innovate increase their share of the market. Those that fall behind and fail to keep pace with the innovations of their competitors lose market share.

Take for instance AMD and the graphic below. When AMD introduced its Athlon chip in 1999, it was an instant success. The Athlon chip was praised by technology reviewers and won performance benchmark tests versus Intel's Pentium III. As a result AMD grabbed market share. It had only 13.6% of the market in 1999, but had over 20% market share two years later - **a nearly 50% growth in market share.**[28]

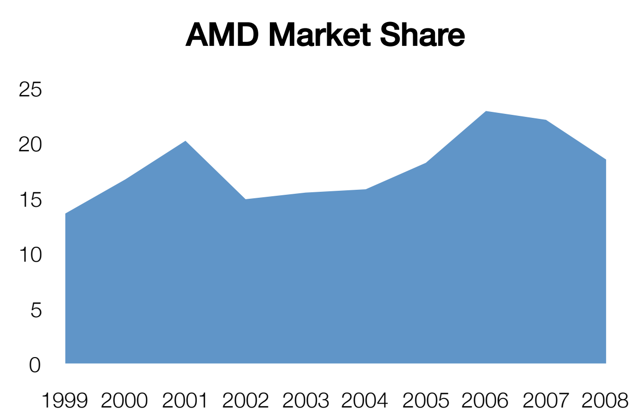

AMD's share declined, however, when Athlon fell behind the performance of Intel's Pentium 4 processor in 2001. But following the introduction of AMD's innovative K8 architecture in 2003, AMD's share of the market again grew rapidly by almost 50% to 23% of the market. Further growth was probably capped only by its production capacity. Market share fell in 2007 after Intel introduced its highly innovative Core microarchitecture and AMD experienced growing pains.

As a result of the competition from AMD, Intel's share of the market decreased precipitously. The above graph highlights how Intel lost eight percent of its overall share of the market by the end of 2006.

The key here is to highlight the dynamism of the market for computer processors. Despite its dominance by two firms, the market quickly rewards innovative products introduced by each company.

Innovation and Competition Drive Prices Down

Competitive pressures and intensive R&D have been the driving forces behind ever decreasing prices. Processor price drops have been dramatic:

- Computing power that cost $2.73 in 1996 cost a penny in 2006.[29]
- In the last ten years, the average price of Intel's microprocessors for personal computers has fallen 60 percent. In 1999, the average selling price of a computer processor was $192. In 2009 it was $73.[30]
- From 2000 to 2008, relative processor performance among high-end chips became approximately 28 times better while the price was cut in half.

As an example, compare two Intel CPU products, the Pentium III from 2000 and the Core 2 Quad that was introduced in 2008. In 2000, a Pentium III was more than twice the price of a Core 2 Quad, yet it had only 3% of the transistors—28 million versus 820 million. The time frame involved is staggering—in only eight years, processor performance for these high-end chips was about 28 times better. And all these performance gains occurred as prices were cut by more than half.

Prices Decline as Performance Rises

Processor	Price	Transistors	Relative Perform.
2000 Pentium III 850MHz	$715	28 million	1
2008 Core 2 Quad 9450	$295	820 million	28
Change	-59%	30x	28x

The following graph from U.S. Bureau of Labor Statistics is a producer price index of tech-related industries.[31] As the graph shows, the technology industry as a whole is characterized by falling prices over time. However, the price drop in microprocessors is by far the most marked. Indeed, according to the U.S. Department of Labor, the quality-adjusted price of CPUs has declined more than any of the 1,200 products it has tracked, including software, storage devices, portables and laptops, and PCs. The real cost of processing power has dropped roughly 40% annually over the past 10 years.

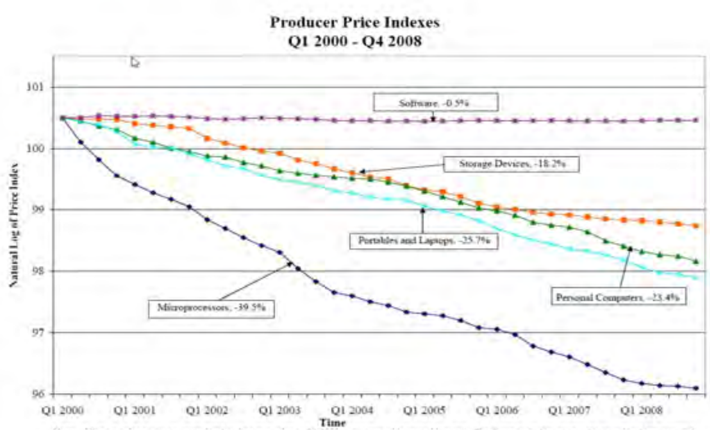

We Demand Chip Innovation

Developers Rely on Continued Innovation By Chipmakers

- ❑ Developers and vendors rely on hardware as an integral part of the software ecosystem
- ❑ New processor capabilities enable software innovation
- ❑ Chipmakers interact with programmers so that software maximizes the benefits from new processor features such as multicore and power management

Innovation in semiconductor construction has led to vast technological improvements and cost reductions over the past twenty years. But the benefits are not merely limited to hardware. The opportunities for software developers to deliver new and compelling products increase with every successive processor that is faster and cheaper. As a result, we have seen rapid improvements in everything from consumer software applications to sophisticated IT systems.

In this section we identify some of these software innovations and connect them to specific advances in hardware technology. While software as an industry is distinct from hardware manufacturers and chip designers, software and hardware increasingly depend on each other to build and sell better products.

Hardware Innovations Enable Software Features—and Vice Versa

The route to faster applications, reduced power consumption, and high definition video is paved by a symbiotic relationship between processor makers and software developers. Software developers, including ACT's membership of small and mid-size information technology firms, rely on an environment that inspires and rewards innovation.

Software developers take advantage of a robust hardware ecosystem that produces better and faster hardware platforms. Better processors enable software innovations that are integrated into new software versions. Consumers buy these new products and reward developers for their efforts.

Consider the case of a new mom and dad that want to edit photos of their newborn to send to grandma. They use Adobe Photoshop, a photo-storing and editing program that needs fast processors and lots of memory to run properly. Let's see how hardware technology makes a difference what they can do with photos of their new baby.

Twenty years ago, the Tandy 5000 we described earlier would have balked at today's conventionally-sized digital photo. The Tandy had VGA graphics with a resolution of 640x480, which couldn't display—much less edit—a 2MB sized 1600x1200 photo. A decade ago, it would take nearly two hours to generate a CD slide show of two dozen photos with a soundtrack. Today, Photoshop can generate higher quality in just minutes.

Today's computers are much more advanced, of course. And so is today's Photoshop. When compared to version 5, which was released just over a decade ago, the features of Adobe Photoshop version 11 CS4 are improved and much more sophisticated. Whereas v.5 could for the first time enable editable type and multiple undo, v. 11 uses perspective-based editing and content-aware scaling. But these software features require advances in computer hardware. V.5 required only a Pentium processor, 32MB of RAM and 72MB of hard disk space. V.11 needs a 1.8 GHz processor, 512MB RAM, and 1GB hard disk space.

There are numerous other examples of software benefitting from hardware. Just think how long a Google or Bing search would take if not for the thousands of advanced database servers working behind the scenes. Or whether a computer could even run a Java application to upload photos and tag photos, such as what users do every day on Facebook. Even the Human Genome Project benefitted from hardware advancements, speeding up the process of cataloguing DNA as hardware speeds increased.

Stanford computer science professor and venture capitalist David Liddle described the relationship between Moore's Law and the creation of the commercial software industry as a meaningful force in the economy:[32]

> The ever lower cost and wide availability of powerful processors dramatically increased the use of computers, enabling software developers to sell their products in huge numbers at affordable prices. Secondly, far fewer architectural variations meant that successful software products needed to run on only one or two different CPU types as this would cover almost the entire market of vendors and customers. Moreover, software developers knew "that there would be inexorable steady improvements in cost/performance which would seldom require any significant changes to the programs, thus allowing larger software investments to be made, in products which would surely perform better and better over time, courtesy of Moore's Law."

But the reverse is also true. Chipmakers depend on software programmers to maximize new hardware features. Without software, all that power and speed goes to waste because it is software that allows users to interact with technology.

Without Software, Silicon is Just Sand

While software depends on hardware, it is clear that hardware also relies on new software features. The release of Windows 7 will increase PC sales by 6.9 percent worldwide in the fourth quarter, according to research firm IDC—even despite the slow economy.[33] Likewise, industry experts predicted new computer sales would increase 10% when Vista was introduced in 2007, the first Windows operating system to be optimized for multicore.[34]

While new software drives hardware sales, it can also be a matter of a device manufacturer's survival. Motorola is hoping to reverse its multiyear sales decline by introducing Motoblur, a layer of software that sits above the Android operating system. This application will coordinate incoming messages and news feeds on future Motorola handsets, enhancing the hardware feature set. According to Motorola CEO Sanjay Jha, "Motoblur is going to become very important to Motorola"—a statement that highlights how software and Internet connectivity are important mobile phone features.[35]

Indeed, today there is increased interaction between hardware and software. In the past, software applications were instantly faster when processors increased clock speed. Today, software that was not designed to take advantage of multiple computing engines available on chips will not see performance improvements. Applications designed for single-core systems may run slower despite the increased computing power.

As an analogy, imagine that a four-lane highway is added to a region with only single-lane highways. The drivers in that region must learn to use all four lanes to take advantage of the additional capacity. If everybody stayed in a single lane, there is no benefit from the added space.

Therefore, software developers need to change the way they code to take advantage of hardware improvements. And they need the help of chip designers.

The major chip manufacturers have all established support programs for the developer community:

- Sun sells computer servers and workstations based on its own SPARC processors as well as AMD's Opteron and Intel's Xeon processors. It has developer support programs and a Sun Developer Network to assist programmers.[36]

- Intel offers developers free bulletins with relevant technical and strategic information, online and in-house training programs, consulting, software support, software development products, community forums, and updates and drivers. Developers can also get access to pre-release processor and communication software and hardware through the Intel Software Partner Program.[37] In addition, Intel offers a tool called Threading Building Blocks which helps software developers take advantage of multi-core processor performance without having to be a threading expert.[38]

- AMD has tools – such as the AMD CodeAnalyst Performance Analyzer – which analyze software performance on AMD microprocessors and help developers optimize application performance. In addition, AMD offers developers access to drivers and downloads, docs and articles, forums and newsletters, and developer guides and manuals.[39] With each new processor advancement, increased attention must be devoted to application design, testing, and optimization. Chip makers and computer manufacturers work closely with operating system designers and software developers to provide these tools.[40]

New Applications and Better Features—the Software Benefits of Hardware Innovation

New processor advancements lead to identifiable software features and have had a significant effect on enterprise and consumer applications for:

- **Photo and Video Editing/Creation** — Faster CPUs and graphics processors enable

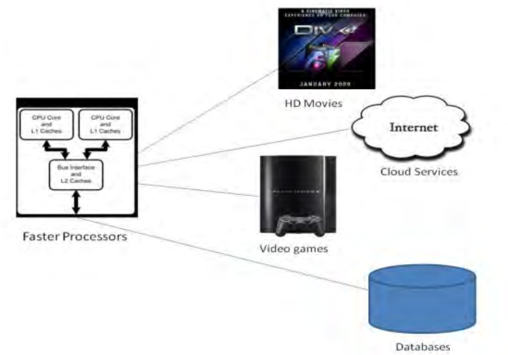

seamless photo and video editing.

- **Gaming** — New processor technologies help games render graphics more quickly, allow for more complex images where objects interact in ever so realistic ways, and enable massive multiplayer games.

- **Database Applications** — Database applications are faster, can hold more data, and use less server capacity thanks to multi-core and hyperthreading technologies.

- **Office and Enterprise Productivity** — Productivity applications such as Microsoft Excel have more features because they effectively utilize multiple threads and take advantage of multicore architecture.

- **Cloud Services** — Platform-level multicore driver innovation in network interface controller (NIC) related technologies reduce overhead by processing more information packets. NIC hardware drivers can use multiple queues to balance load and keep networking overhead as low as possible.

- **Operating Systems** — Benefits in performance and responsiveness, energy efficiency and power management, and graphics and multi-media.

From Pac-Man to Realistic Men: Online Gaming, Photorealism and HD Video

Computer games are often on the leading edge of software development and therefore are among the most hardware-intensive of applications. New processor technology helps games quickly render realistic images.

Compare the drastic differences in levels of detail in the following images: Pac-man, which was introduced in 1980, with Uncharted 2, which won the E3 2009 Video Games award for best graphics.[41]

But perhaps the biggest gaming breakthrough has been cloud computing. Powerful server processors help gaming companies optimize their products for streaming over the Internet, and even enable richer graphics than equivalent games installed locally on a user's computer. For instance, AMD's "Fusion Render Cloud" provides high-end CPUs, graphical processor units (GPUs), and other platform technologies for fully-interactive high-definition gaming experiences over the Internet. Now, multiple players can play games in a Web browser.

New processor technology has also helped online gaming companies save money by dramatically reducing infrastructure, energy bills, and server rack space.

Case Study: VitruView

InterKnowlogy is a software development company that created a groundbreaking application for the Microsoft Surface. The Surface is a software and hardware technology that allows for advanced manipulation of digital content. On top of this advanced Microsoft platform, InterKnowlogy built VitruView, an application for doctors and other health care professionals. VitruView was designed to help with angiography procedures in a catheter lab. Using the Windows Presentation Format application.

VitruView is a three dimensional, multi-touch application that gives doctors the ability to zoom and rotate a 3D virtual image of the human heart. Doctors can annotate inside and outside the heard by placing stents and marking lesions, and can even add and remove arteries using their fingers. All of this would not be possible without advanced hardware capabilities.

Finally, new hardware also helps software developers deliver lifelike visual effects. Multicore and threading, new microarchitecture design, and advanced media instruction sets benefit applications such as Pegasys TMPEGEnc, Pinnacle Studio 12 and Cyberlink PowerDirector 8.

One example is OTOY, which is a software company that uses AMD's supercomputers to power its server-side rendering technology.

According to OTOY's CEO, Jules Urbach:

"Video editors don't need to spend 40 hours rendering one frame, you can do it instantly. Now you can render 100 cinema-quality viewpoints on one graphics card. That delivers you high-end games that are higher quality and more realistic than what you see today."[42]

Movies benefit in similar ways. A critical phase in the digital cinema production pipeline is to create realistic waterfalls, clouds, fire, and clothing. Special effects applications help to simulate fluid flow, particle simulation and cloth dynamics. Parallelizations on multi-core as well as other optimizations deliver up to four times the speed for cloth simulations and up to eight times the speed for fluid flow simulations.

Procedural animations in the balloon scene of the recent Pixar movie "Up" depended upon advanced processor features.[43] The movie "The Curious Case of Benjamin Button" featured optical motion scanning to depict the changing of age for Brad Pitt's character. The skin was scanned and captured like a hologram, and the data was rendered using supercomputers.

Fast and Powerful: Operating Systems on Multicore

The new Mac OS X 10.6 operating system is one example of how software has been explicitly designed to take advantage of new hardware technologies.[44] Apple has a project called Grand Central Dispatch to take advantage of multicore processor technology and minimize the difficulties of parallel programming. The operating system handles complicated administrative chores so programmers do not have to do as much work to optimize their programs.

For the development of Windows 7, Microsoft and Intel worked together to optimize the next Windows version for multicore processor technology.[45]

Microsoft designed Windows 7 to divide tasks like video encoding for simultaneous execution over multiple cores and threads. This makes applications run more quickly. Microsoft also made changes to the kernel to improve the power management of a chip's cores. DVD playback on a battery-powered Windows 7 laptop is 16 watts, compared to a drain of 20 watts on a similar Windows Vista laptop.[46]

Microsoft also makes better use of a timer system that puts cores in Intel processors back into sleep mode when idle. Depending on usage, the OS can intelligently put different cores into different power states, and processors remain in idle mode for a longer period in Windows 7 compared to Windows Vista. Another performance benefit in Windows 7 was collaboration with Intel to optimize Solid State Drive technology.

The operating system is not the only software that takes advantage of multicore. Application developers also change their code to detect and use multiple processors.

Virtualization

Virtualization technology baked into chip design has eliminated the need for complex software workarounds.[47] In the past, virtualization implementations required time-consuming programming and translations of existing code. It also required software changes to match virtual machine monitor code and avoid hardware architectural limitations. New processors make software simpler, efficient, and more robust by eliminating the need for extensive software changes, and support a long list of unmodified operating systems including Windows and many Linux distributions. Today all vendors embrace hardware virtualization technology, including those vendors—such as VMware—that originally had solutions based on software workarounds.

Highlighting processor technology improvements and price reductions helps reveal how applications, users, and the software industry as a whole have benefited from hardware advancements. Moore's law is seemingly an impossible challenge, yet it is being met—increasingly, through close interaction between hardware and software developers.

Conclusion

The Market is Working For Software Developers

❑ Chipmakers are Meeting the Needs of Software Developers and Consumers Through Constant Innovation and Price Decreases

❑ Government Intervention in This Market Could Derail These Benefits

Technology has improved rapidly, particularly over the last decade. At the same time, prices have decreased. Processor design and performance have enabled ACT's members to create better and faster software. Simply put, software developers expect—and need—processors to be the enabler for software and IT innovation.

Chipmakers are meeting this demand. The processor market is among the most dynamic of any industry and the current state of innovation is robust. Chip makers invest heavily in R&D to deliver new hardware features, lower prices and keep innovating:

1. Technology improves—processors pack more features on smaller chips to drive Moore's Law

2. Prices decrease—processor prices have fallen more than any other tech-related industry

3. Innovation continues—chip makers invest heavily in R&D

From a software developer's perspective, the IT ecosystem is working. Two major developments are currently working in the favor of software firms:

Convergence—The processor market is more than just AMD and Intel. ARM, IBM, Qualcomm, Samsung, Texas Instruments and others that specialize in the embedded or mobile market that are well positioned for future growth. The result of all this competition is more choice for software developers.

Cooperation—Chipmakers work closely with software developers to ensure that applications and IT services can take advantage of new hardware features. Likewise, chipmakers hear from programmers who require faster, improved processors for new applications.

At a time when the European Union and U.S. are analyzing the competitiveness of the processor market, we caution regulators to not interfere. The microprocessor market serves as a prime example of the type of competition that produces innovation at a breakneck pace year after year, with spillover benefits for other industries. The widely recognized and unprecedented increase in quantity, quality, speed, functionality, and choice of microprocessors with a corresponding steady and significant decrease in prices, has greatly benefitted the software industry. Public policies should focus on markets where such benefits are not being produced. Where markets are working, governments should allow innovation to continue and the competitive process to flourish.

About the Authors

Braden Cox is Research & Policy Counsel at the Association for Competitive Technology. He earned a Bachelor's degree in Finance and a J.D. from the University of Georgia.

Nora von Ingersleben is Policy Analyst at the Association for Competitive Technology. She holds a Bachelor's degree in Politics and Philosophy from the University of York and a Master's degree in International Relations from the London School of Economics.

End Notes

[1] Gordon E. Moore (1965), "Cramming more components onto integrated circuits," Electronics, Volume 38, Number 8, April 19, 1965, available at ftp://download.intel.com/museum/Moores_Law/Articles-Press_Releases/Gordon_Moore_1965_Article.pdf.

2 Gordon E. Moore (2007), as quoted in Bryan Gardiner, "IDF: Gordon Moore Predict End of Moore's Law (Again)," Wired, September 18, 2007, available at http://www.wired.com/epicenter/2007/09/idf-gordon-mo-1/.

[3] A processor is the logic circuitry that responds to and processes the basic instructions that drive a computer. Processors are the "brains" of every computing device on earth. The processor in a personal computer or embedded in small devices is often called a microprocessor.

[4] David Yen (2004), "Perspective: End of Moore's Law? Wrong question," *CNet News*, February 18, 2004, available at http://news.cnet.com/2010-1006-5160336.html.

[5] Ann Steffora Mutschler, "Semiconductor R&D spending growth to show 8% rise in 2008, IC Insights forecasts," *Electronics Design, Strategy, News*, July 10, 2008, available at http://www.edn.com/article/CA6577202.html.

[6] Semiconductor Industry Association (n.d.), "The Research Partnership," *Choose to Compete: Semiconductors and America's Innovation Leadership*, available at http://www.choosetocompete.org/research_partnership.html.

[7] Josh Cable (2009), "Holding the Line on R&D," *Industry Week*, July 22, 2009, available at http://www.industryweek.com/articles/holding_the_line_on_rd_19612.aspx?ShowAll=1.

[8] David Manners, "Top Ten Chip Companies For R&D Spending," *EW Blogs*, November 28, 2007, available at http://www.electronicsweekly.com/blogs/david-manners-semiconductor-blog/2007/11/top-ten-chip-companies-for-rd.html.

[9] Unknown (2008), "Annual Earnings Motorola Inc. (MOT)," *The Wall Street Journal*, available at http://online.wsj.com/quotes/a_earnings.html?mod=2_0483&symbol=mot&news-symbol=MOT.

[10] Freescale Semiconductor (n.d.), *The Multifaceted World of Multi-Core Processing*, available at http://www.freescale.com/webapp/sps/site/overview.jsp?nodeId=02VS0l09206140.

[11] Texas Instruments (2009), *OMAP35x Applications Processors*, Product Bulletin, available at http://focus.ti.com/lit/ml/sprt457b/sprt457b.pdf.

[12] Samsung (2009), *SAMSUNG and Intrinsity Jointly Develop the World's Fastest ARM Cortex™-A8 Processor Based Mobile Core in 45 Nanometer Low Power Process*, Press Release, July 27, 2009, available at http://www.samsung.com/global/business/semiconductor/newsView.do?news_id=1030.

[13] Sharon Pian Chan (2009), "As data centers suck up energy, money, tech industry looks to turn down the heat," *The Seattle Times*, May 11, 2009, available at http://seattletimes.nwsource.com/html/microsoft/2009204767_greendata11.html.

[14] Sun Microsystems (n.d.), *Eco-managing Your Data Center*, available at http://www.sun.com/emrkt/innercircle/newsletter/0109/feature-itm.html.

[15] Intel (2009), *Balancing Performance and Power Consumption: Power management capabilities give developers control over power consumption*, Technology Brief, available at http://www.intel.com/technology/advanced_comm/322286.pdf and Advanced Micro Devices (2007), ACP – The Truth About Power Consumption Starts Here, AMD White Paper, available at http://www.amd.com/us-en/assets/content_type/white_papers_and_tech_docs/43761C_ACP_WP.pdf.

[16] Craig Barrett (2009), quoted in Michael S. Malone, "From Moore's Law to Barrett's Rules," *The Wall Street Journal*, May 16, 2009, available at http://online.wsj.com/article/SB124242845507325429.html.

[17] Owen W. Linzmayer (2004), *Apple Confidential: The Real Story of Apple Computer, Inc.*

[18] Ibid.

[19] Bruce Schlender (1995), "Steve Jobs' Amazing Movie Adventure Disney is Betting Computerdom's Ex-Boy Wonder to Deliver this Year's Animated Christmas Blockbuster. Can he Do for Hollywood What He Did for Silicon Valley?", *Fortune*, September 18, 1995, available at http://money.cnn.com/magazines/fortune/fortune_archive/1995/09/18/206099/index.htm

[20] Jane Furth (1995), "Steve Jobs' Amazing Movie Adventure Disney is Betting on Computerdom's Ex-Boy Wonder to Deliver This Year's Animate Christmas Blockbuster. Can He Do for Hollywood What He Did for Silicon Valley?" *Fortune Magazine*, September 18, 1995, available at http://www.amd.com/us-en/Weblets/0,,7832_12670_13047,00.html. AMD cites Martin Neil Baily (2002), "The New Economy: Post Mortem or Second Wind," *Journal of Economic Perspectives,* 16(2), January 2, 2002, available at https://www.researchgate.net/publication/4902196_Distinguished_Lecture_on_Economics_in_Government_The_New_Economy_Post_Mortem_or_Second_Wind. It also cites Dale W. Jorgenson, Mun S. Ho, and Kevin J. Stiroh (2002), *Projecting Productivity Growth: Lessons from the U.S. Growth Resurgence*, Resources for the Future Discussion Paper 02-42, July 2002, available at http://www.rff.org/Documents/RFF-DP-02-42.pdf.

[21] Baily (2002).

[22] Jorgenson, Ho, and Stiroh (2002).

[23] Apple (2003), *Apple and IBM Introduce the PowerPC G5 Power Processor: World's First 64-Bit Desktop the Heart of the World's Fastest Personal Computer*, Press Release, June 23, 2002, available at http://www.apple.com/pr/library/2003/jun/23joint.html.

[24] Paul Lilly (2008), "IBM Gives Details on Octo-Core Power 7 Processor, *MaximumPC*, August 26, 2008, available at http://www.maximumpc.com/article/news/ibm_gives_details_octocore_power7_processor.

[25] Barry Collins (2009), "ARM launches attack on Intel's netbook stranglehold," *PC Pro*, September 16, 2009, available at http://www.pcpro.co.uk/news/351619/arm-launches-attack-on-intels-netbook-stranglehold.

[26] Paul Taylor (2009), "View from the Top: Paul Jacobs, chief executive of Qualcomm," *Financial Times*, August 28, 2009, available at http://www.ft.com/cms/s/0/3bbfc30a-9369-11de-b146-00144feabdc0,_i_email=y.html.

[27] Brooke Crothers (2009), "Intel and Apple – Future Rivals?", *CNet*, September 20, 2009, available at http://news.cnet.com/8301-1001_3-10356827-92.html

[28] Market share data compiled from publicly available sources.

[29] Calculated by dividing the MIPS of an Intel Pentium Pro at 200 MHz released in 1996 (541MIPS) by the average selling price of a processor in 1996 ($198) and dividing Intel's Core Duo running at 2.13 GHz (20,000 MIPS) with its selling price ($200). See Chris Anderson, "The Rise of 'Freeconomics'", Nov. 26, 2006, available at http://www.longtail.com/the_long_tail/2006/11/the_rise_of_fre.html

[30] According to yearly average selling prices.

[31] U.S. Bureau of Labor Statistics, industry data available at http://data.bls.gov/cgi-bin/dsrv?pc

[32] David E. Liddle (2006), "The Wider Impact of Moore's Law," *Solid State Circus*, September 2006 Issue, available at http://www.ieee.org/portal/site/sscs/menuitem.f07ee9e3b2a01d06bb9305765bac26c8/index.jsp?&pName=sscs_level1_article&TheCat=2165&path=sscs/06Sept&file=Liddle.xml.

[33] Peter Burrows, "Will Windows 7 Reboot PC Sales?"*BusinessWeek*, September 08, 2009, available at http://www.zdnetasia.com/news/software/0,39044164,62057547,00.htm.

[34] Matthew Wilkins (2007), "Vista, Dual-Core to Boost PC Growth," *ChannelWeb*, January 4, 2007, available at http://www.crn.com/showArticle.jhtml?articleId=196801311.

[35] Tom Krazit (2009), "Motorola's comeback Attempt Rests on Software," *CNet News*, September 11, 2009, available at http://news.cnet.com/8301-1035_3-10350204-94.html?tag=nl.e703.

[36] Sun Microsystems (n.d.), *SDN: A Community for Sun Developers*, available at http://developers.sun.com/.

[37] Intel (n.d.), *About the Intel Software Network*, available at http://software.intel.com/sites/support/.

[38] Intel (n.d.), Intel Threading Building Blocks 2.2 for Open Source, available at http://www.threadingbuildingblocks.org/.

[39] Advanced Micro Devices (2009), "AMD Code Analyst Performance Analyzer," *AMD Developer Central*, available at http://developer.amd.com/cpu/CodeAnalyst/Pages/default.aspx.

[40] Hewlett Packard (2007), *Multi-Core Optimization Program*, available at http://www.hp.com/techservers/hpccn/hpccollaboration/ADCatalyst/Multi-core_Optimization_Program.html.

[41] GameTrailers.com web site, available at http://www.gametrailers.com/video/best-graphics-best-of-e3/51488.

[42] Andrew Nusca (2009), "OTOY's Jules Urbach: Cloud computing "threatens" game consoles, Blu-Ray, PCs [Interview]," *ZDNet*, September 11, 2009, available at http://blogs.zdnet.com/BTL/?p=24154.

[43] Daniel Terdiman (2009), "How technology lifts Pixar's 'Up,'" *Geek Gestalt*, May 27, 2009, available at http://news.cnet.com/8301-13772_3-10249396-52.html.

[44] Apple (2009), *Mac OS X*, available at http://www.apple.com/macosx/technology/.

[45] Joakim Lialias (2009), "Your Windows 7 Experience: Rooted in Close Technical Collaboration," *Windows Partner Blog*, July 22, 2009, available at http://windowsteamblog.com/blogs/partner/archive/2009/07/22/your-windows-7-experience-rooted-in-close-technical-collaboration.aspx.

[46] Agam Shah (2009), "Microsoft, Intel Build Fast, Efficient Windows 7 Machines," *PCW Business Center*, September 1, 2009, available at http://www.pcworld.com/businesscenter/article/171275/microsoft_intel_build_fast_efficient_windows_7_machines.html.

[47] The May 2005 issue of IEEE Computer documented in detail the architectural hurdles that a software designer would have to overcome to implement a VMM without VT.