Moore's Law & More: Fast, Cheap Computing, and What it Means for the Manager

Last modified: Nov. 10, 2008
Note: this is an earlier version of the chapter. All chapters updated after July 2009 are now hosted (and still free) at http://www.flatworldknowledge.com. For details see the 'Courseware' section of http://gallaugher.com

LEARNING OBJECTIVE:
After studying this section you should be able to:

1. Define Moore's Law, as well as the rate of advancement for magnetic storage (disk drives) and telecommunications (fiber optic transmission).
2. Understand how the price elasticity associated with faster / cheaper technologies opens new markets, creates new opportunities for firms and society, and can catalyze industry disruption.
3. Provide concrete examples of firms that have leveraged the rise of faster / cheaper technologies to their advantage.
4. Recognize and define various terms for measuring data capacity.
5. Consider the managerial implication of faster / cheaper computing on areas such as strategic planning, inventory, and accounting.

INTRODUCTION

Faster and cheaper – those two words have driven the computer industry for decades, and the rest of the economy has been along for the ride. Today it's tough to imagine a single industry not impacted by more powerful, less expensive computing. Faster and cheaper puts mobile phones in the hands of peasant farmers, puts a free video game in your Happy Meal, and drives the drug discovery that may very well extend your life.

Some Definitions

This phenomenon of 'faster, cheaper' computing is often referred to as Moore's Law, after Intel co-founder, Gordon Moore. Moore didn't show up one day, stance wide, hands on hips, and declare "behold my law", but he did write a four-page paper for *Electronics Magazine* in which he described how the process of chip making enabled faster chips to be manufactured at cheaper prices.

Moore's friend, legendary chip entrepreneur and CalTech professor, Carver Mead, later coined the "Moore's Law" moniker. It was snappy, plus as one of the founders of Intel, Moore had enough geek cred for the name to stick. Moore's original paper offered language only a chip designer would love, so we'll rely on the more popular definition: *chip performance per dollar doubles every 18 months* (Moore's original paper assumed two years, but many sources today refer to the 18-month figure, so we'll stick with that).

Moore's Law applies to chips – broadly speaking, to *processors* and *RAM* chips, or the electronics stuff that's made out of silicon[1]. The *microprocessor* is the brain of a computing device. It's the part of the computer that executes the instructions of a computer program; creating, say, a web browser, word processor, video game, or virus. For processors, Moore's Law means that next generation chips should be twice as fast in 18 months, but cost the same as today's models (or, chips that are same speed should cost half as much as they do today).

RAM chips are chip-based memory (RAM stands for Random Access Memory). The RAM inside your personal computer is *volatile memory*, meaning that when the power goes out, all is lost that wasn't saved to *non-volatile device* (e.g. a hard disk or other more permanent storage media). Cameras, MP3 players, USB drives, and mobile phones use *flash memory* (sometimes called flash RAM). It's not as fast as the RAM in your PC, but holds data even when the power is off (so flash memory is also non-volatile memory). With RAM chips, Moore's Law means that in 18 months you'll pay the same price as today for twice as much storage. Computer chips are sometimes also referred to as *semiconductors*, so if someone refers to the *semiconductor industry*, they're talking about the chip business[2].

Moore's Law does not, strictly speaking, apply to other technology components. Chips are *solid-state* devices with no moving parts (that's why it's better to jog with a iPod nano, with songs stored on chips, instead of with a hard drive player that can skip if the whirring drive is jostled). But other computing components are also seeing their price/performance curves skyrocket exponentially. Data storage doubles every 12 months. Networking speed is on a tear, too. With an equipment change at the ends of the cables, the amount of data that can be squirt over a fiber-optic line can double every 9 months[3].

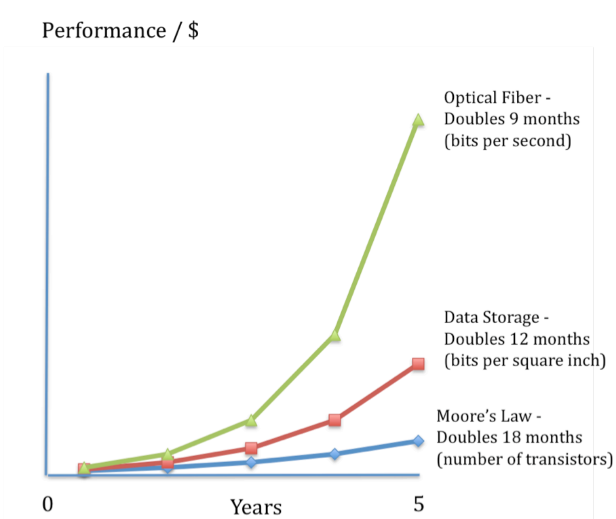

Advancing rates of technology (Silicon, Storage, Telecom)[4]

[1] Although other materials besides silicon are increasingly being used.
[2] Semiconductor materials, like the silicon dioxide used inside most computer chips, are capable of enabling as well as inhibiting the flow of electricity. These properties enable chips to perform math or store data.
[3] Fiber optic lines are glass or plastic data transmission cables that carry light. These cables offer higher transmission speeds over longer distances than copper cables that transmit electricity.
[4] Adopted from Shareholder Presentation by Jeff Bezos, Amazon.com, 2006

Get out your Crystal Ball

Faster and cheaper makes possible the once impossible. That means that as a manager, your job will be about predicting the future. First, consider how the economics of Moore's Law opens new markets. When technology gets cheap, *price elasticity* kicks in. Tech products are highly *price elastic*, meaning consumers buy more products as they become cheaper[5]. And it's not just that existing customers load up on more tech, entire *new markets* open up as firms find new uses for these new chips.

Just look at the *five waves of computing* we've seen over the previous five decades[6]. In the *first wave* in the 1960s, computing was limited to large, room-sized mainframe computers that only governments and big corporations could afford. Moore's Law kicked in during the 1970s for the *second wave*, and minicomputers were a hit. These were refrigerator-sized computers that were as speedy or speedier than the prior generation of mainframes, yet were affordable by work groups, factories, and smaller organizations. The 1980s brought *wave three* in the form of PCs, and by the end of the decade nearly every white-collar worker in America had a fast/cheap computer on their desk. In the 1990s *wave four* came in the form of Internet computing – cheap servers and networks made it possible to scatter data around the world, and with more power, personal computers displayed graphical interfaces that replaced complex commands with easy-to-understand menus accessible by a mouse click. At the close of the last century, the majority of the population in many developed countries had home PCs, as did most libraries and schools.

Now we're in *wave five*, where computers are so fast and so inexpensive that they become ubiquitous – woven into products in ways few imagined years before. Silicon is everywhere! It's in the throw-away RFID tags that track your luggage at the airport. It provides the smarts in the world's billion-plus mobile phones. It's the brains inside robot vacuum cleaners, next generation Legos, and the table lamps that change color when the stock market moves up or down. These digital shifts can rearrange entire industries. Consider that today the firm that sells more cameras than any other is Nokia, a firm that offers increasingly sophisticated chip-based digital cameras as a give-away as part of its primary product, mobile phones. This shift has occurred with such sweeping impact that former photography giants Pentax, Konica, and Minolta have all exited the camera business.

Ambient Devices and the Fifth Wave

Carl Yankowski almost never gets caught in the rain without his umbrella. That's because Yankowski's umbrella regularly and wirelessly checks weather reports on its own. If the umbrella gets word it will rain in the next few hours, the handle blinks with increasing urgency, warning its owner with a signal that seems to declare "you will soon require my services". Yankowski is CEO of 'fifth wave' firm Ambient Devices, a Massachusetts startup that's embedding computing and communications technology into everyday devices in an attempt to make them 'smarter' and more useful.

Ambient's ability to pull off this little miracle is evidence of how quickly innovative thinkers are able to take advantage of new opportunities and pioneer new markets on the back of Moore's Law. The firm's first product, the Orb, is a lamp that can be set up to change color in real time in reaction to factors such as the performance of

[5] As opposed to goods and services that are *price inelastic* that consumers will try their best to buy even if prices go up, like healthcare and housing.
[6] Copeland, 2005

your stock portfolio or the intensity of the local pollen count. In just six months, the 10 refugees from MITs Media Lab that founded Ambient Devices took the idea for the Orb, designed the device and its software, licensed wireless spectrum from a pager firm that had both excess capacity and a footprint to cover over 90 percent of the US, arranged for manufacturing, and began selling the gizmo through Brookstone and Nieman Marcus[7]. Ambient has since expanded the product line to several low-cost appliances designed to provide information at a glance. These include the Ambient Umbrella, as well as useful little devices that grab and display data ranging from sports scores to fluctuating energy prices (so you'll put off running the dishwasher 'til evening during a daytime price spike). The firm even partnered with LG on a refrigerator that can remind you of an upcoming anniversary as you reach for the milk.

Products developed by 'fifth wave' firm Ambient Devices include the weather-reading Ambient Umbrella, the Energy Joule, LG's Information Center refrigerator, and the Orb lamp

One of the most agile surfers of this *fifth wave* is Apple, Inc. (see the Apple case) – a firm with a product line that is now so broad that in January 2007, it dropped the word 'Computer' from its name. Apple's breakout resurgence owes a great deal to the iPod. Launched in October 2001, the original iPod sported a 5GB hard drive that Steve Jobs declared would "put 1,000 songs in your pocket". Cost? $399. Less than six years later, Apple's highest-capacity iPod sold for $50 less than the original, yet held 40 *times* the songs. By that time the firm had sold over 150 million iPods – an adoption rate faster than the original Sony Walkman. Low-end iPods now go for less than $50 a piece, and high-end models have morphed into Internet browsing devices capable of showing maps, playing videos, and gulping down songs from Starbucks WiFi while waiting in line for a latte.

The original iPod has also become the jumping off point for new business lines including the iPhone, Apple TV, and iTunes store. As an online store, iTunes is always open. iTunes sold over 20 million songs on Christmas Day 2007, a day when virtually all of its US-based competition was closed for holiday. In a short five years after its introduction, iTunes has sold over 4 billion songs and has vaulted past retail giants Wal-Mart, BestBuy, and Target to become the number one music retailer in the US. Today's iTunes is a digital media powerhouse, selling movies, TV shows, games, and other applications. And with podcasting, Apple's iTunes University even lets students at participating schools put their professors' lectures on their gym playlist for free. Surfing the fifth wave has increased the value of Apple stock sixteen-fold six years after the iPod's launch. Ride these waves to riches, but miss the power and promise of Moore's Law and you risk getting swept away in its riptide. Apple's rise occurred while Sony, a firm once synonymous with portable music, sat on the sidelines unwilling to get on the surfboard. Sony's stock stagnated, barely moving in six years. The firm has laid off thousands of workers while ceding leadership in digital music (and video) to Apple.

[7] Copeland, 2005 & Miller, 2003.

Top US Music Retailers

1992	2005	2006	2008
1. Musicland	1. Wal-Mart	1. Wal-Mart	1. iTunes
2. The Handleman	2. Best Buy	2. Best Buy	2. Wal-Mart
3. Tower Records	3. Target	3. Target	3. Best Buy
4. Trans World Music	**7. iTunes**	**4. iTunes**, Amazon tie	4. Amazon, Target tie

Moore's Law restructures industries. The firms that dominated music sales when you were born are now bankrupt, while one that had never sold a track now sells more than anyone else.[8]

While the change in hard-drive prices isn't directly part of Moore's Law (hard drives are magnetic storage, not silicon chips), as noted earlier, the faster/cheaper phenomenon applies to storage, too. Look to Amazon as another example of jumping onto a once-impossible opportunity courtesy of the price/performance curve. When Amazon.com was founded in 1995, the largest corporate database was one terabyte or TB (see "Bits & Bytes" below) in size. In 2003, the firm offered its "Search Inside the Book" feature, digitizing the images and text from thousands of books in its catalog. "Search Inside the Book" lets customers peer into a book's contents in a way that's both faster and more accurate than browsing a physical bookstore. Most importantly for Amazon and its suppliers, titles featured in "Search Inside the Book" enjoyed a seven percent sales increase over non-searchable books. When "Search Inside the Book" launched, the database to support this effort was 20 TB in size This means that in just eight years, the firm found that it made good business sense to launch an effort that was a full *20 times* larger than anything used by *any* firm less than a decade earlier. For Amazon, the impossible had not just become possible, it became good business.

Bits & Bytes[9]

Computers express data as bits that are either one or zero. Eight bits form a byte (think of a byte as being a single character you can type from a keyboard). A kilobyte refers to roughly a thousand bytes, or thousand characters, Megabyte = 1 million, Gigabyte = 1 billion, Terabyte = 1 trillion, Petabyte = 1 quadrillion, Exabyte = one quintillion bytes.

While storage is most often listed in bytes, telecommunication capacity (bandwidth) is often listed in bits per second (bps). The same pre-fixes apply (Kbps = kilobits, or one thousand bits, per second, Mbps = megabits per second, Gbps = gigabits per second, Tbps = terabits per second).

These are managerial definitions, but technically, a kilobyte is 2^{10} bytes, or 1,024. Mega = 2^{20}, Tera = 2^{30}, Giga = 2^{40}, Peta = 2^{50}, and Exa = 2^{60}. To get a sense for how much data we're talking about, see the table below.

8 Note: 12-tracks equivalent to one CD. Source: Los Angeles Times.
9 Schuman 2004, Huggins 2008

	Managerial Definition	Exact Amount	To put it in perspective
1 Byte	One keyboard character	8 bits	1 byte = a letter or number
1 Kilobyte (KB)	One thousand bytes	2^{10} bytes	1 typewritten page = 2 KB
1 Megabyte (MB)	One million bytes	2^{20} bytes	1 digital book (Kindle) = approx. 500-800 KB
1 Gigabyte (GB)	One billion bytes	2^{30} bytes	1 digital photo (7 megapixels) = 1.3 MB
1 Terabyte (TB)	One trillion bytes	2^{40} bytes	1 MP3 song = approx. 3 MB
1 Petabyte (PB)	One quadrillion bytes	2^{50} bytes	1 CD = approx. 700 MB
1 Exabyte (EB)	One quintillion bytes	2^{60} bytes	1 DVD movie = approx. 4.7 GB
			Wal-Mart data warehouse (2004) = 1/2 PB
			Printed collection of the Library of Congress = 20 TB
			Amount of data created in 2006 = 160 EB

Here's another key implication – if you are producing products with a significant chip-based component, the chips inside that product rapidly fall in value. That's great when it makes your product cheaper and opens up new markets for your firm, but it can be deadly if you overproduce and have excess inventory sitting on shelves for long periods of time. Dell claims its inventory depreciates as much as a single percentage point in value each week[10]. That's a big incentive to carry as little inventory as possible, and to unload it, fast!

While the strategic side of tech may be the most glamorous, Moore's Law impacts mundane management tasks, as well. From an accounting & budgeting perspective, as a manager you'll need to consider such questions as, how long will your computing equipment remain useful? If you keep upgrading computing and software, what does this mean for your capital expense budget? Your training budget? Your ability to make well-reasoned predictions regarding tech's direction will be key to answering these questions.

Tech for the Poor

Nicholas Negroponte, the former head of MIT's Media Lab, is on a mission. His OLPC project (One Laptop Per Child) aims to create a $100 laptop for distribution to the world's poor. A rubberized keyboard and entirely solid-state design (flash RAM rather than hard drive) help make the machine durable. The laptop's ultra-bright screen is readable in daylight and can be flipped to convert into an e-book reader. And a host of open source software and wiki tools for courseware development, all aim to keep the costs low (The first generation of the product, called the XO, was priced around $175). Mesh networking allows laptops within a hundred feet or so to communicate with each other, relaying a single internet connection for use by all. And since the XO is targeted at the world's poorest kids in communities where power generation is unreliable or non-existent, several battery-charging power generation schemes have been tested including a hand crank, pull chord, and foot pedal. While the XO remains controversial and has failed to meet its initial deployment targets, it is a product made possible by the rapidly falling price of computing.

While the success of the OLPC effort is yet to be determined, another tech product containing a microprocessor is already transforming the lives of some of the world's most desperate poor – the cell phone. There are 3 billion people worldwide that don't yet have a phone, but they will, soon. In the ultimate play of Moore's Law opening up new markets, mobiles from Vodaphone now sell for $25, while Indian telecom provider Spice has a $20 "people's phone" on deck. While it took roughly 20 years to sell a billion mobile phones worldwide, the second billion sold in four years, and the third billion took just two years. Today some 80% of the world's population lives within cellular

[10] Breen, 2004

network range, double the 2000 level. By year-end 2006, the International Telecommunications Union (ITU) claimed that 68 percent of mobile subscriptions were in developing countries[11].


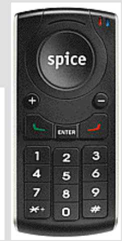

The XO PC, logo for grameen phone of Bangladesh, and the People's Phone by Spice

Why such demand? Mobiles change lives for the better. According to Columbia economist Jeffry Sachs, "the cell phone is the single most transformative technology for world economic development"[12]. Think about the farmer who can verify prices and locate buyers before harvesting and transporting perishable crops to market, the laborer who was mostly unemployed but with a mobile is now reachable by those who have day-to-day work, the mother who can find out if a doctor is in and has medicine before walking in three hours later (?) with her sick child, or the immigrant laborer serving as a housekeeper who was "more or less an indentured servant until she got a cell phone" enabling new customers to call and book her services[13].

As an example of impact, look to poor fishermen in the Indian state of Kerala. By using mobile phones to find the best local marketplace prices for sardines, these fishermen were able to increase their profits by an average of eight percent even though consumer prices for fish *dropped* four percent. The trends benefiting both buyer and seller occurred because the fishermen no longer had to throw away unsold catch previously lost by sailing into a port after all the buyers had left. A 2005 London Business School study found that for every 10 mobile phones per 100 people, a country's GDP bumps up 0.5%[14].

Bangladeshi economist Mohammed Yunus won the Nobel Peace Prize based on his work in the microfinance movement, an effort that provides very small loans to the world's poorest entrepreneurs. Microfinance loans grew the market for Grameen Phone Ltd., a firm that has empowered over 250,000 Bangladeshi "phone ladies" to start businesses that helped their communities become more productive. Phone ladies buy a phone on microcredit for about $150 each. These special long-life battery phones allow them to become a sort of village operator, charging a small commission for sending & receiving calls. Through phone ladies, the power of the mobile reaches even those too poor to afford buying one outright. Grameen Phone now has annual revenues of over $1 billion and is Bangladesh's largest telecom provider.

In another ingenious scheme, phone minutes become a proxy for currency. The NY Times reports that a person "working in Kampala, for instance, who wishes to send the equivalent of $5 back to his mother in a village will buy a $5 prepaid airtime card, but rather than entering the code into his own phone, he will call the village phone operator and read the code to her. She then uses the airtime for her phone and completes the transaction by giving the man's mother the money, minus a small commission"[15].

South Africa's Wizzit and GCash in the Philippines allow customers to use mobile phones to store cash credits sent from another phone or purchased through a post office or kiosk operator. When phones can be used as currency for purchases or payments, who needs Visa? Vodafone's Kenyan-based M-Pesa mobile banking program landed 200,000 new customers in a month – they'd expected it'd take a year to hit that mark. With 1.6 million customers by that time, the service is spreading throughout Africa. The 'mobile phone as bank' may bring banking to a billion unserved customers in a few years.

[11] Corbett, 2008
[12] Ewing, 2007
[13] Corbett 2008
[14] Ewing 2007
[15] Corbett 2008

LEARNING OBJECTIVE:
After studying this section you should be able to:

1. Describe why Moore's Law continues to advance, and discuss the physical limitations of this advancement.
2. Name and describe various technologies that may extend the life of Moore's Law.
3. Discuss the limitations of each of these approaches.

THE DEATH OF MOORE'S LAW?

Moore's Law is possible because the distance between the pathways inside silicon chips is getting smaller. Moore simply observed that we're getting better over time at squeezing more stuff into tinier spaces. While chip plants (semiconductor fabrication facilities, or *fabs*) are incredibly expensive to build, each new generation of fabs can crank out more chips per silicon wafer. And since the pathways are closer together, electrons travel shorter distances. If electronics now travel half the distance to make a calculation, that means the chip is twice as fast.

But the shrinking can't go on forever, and we're already starting to see three interrelated forces – *size*, *heat*, and *power* – threatening to slow down the Moore's Law gravy train. When you make processors smaller, the more tightly packed electrons will heat up a chip – so much so that unless today's most powerful chips are cooled down, they will melt inside their packaging. To keep the fastest computers cool, most PCs, laptops, and video game consoles need fans, and most corporate data centers have elaborate and expensive air conditioning and venting systems to prevent a melt-down. A trip through the Facebook data center during its recent rise would show that the firm was a "hot" startup in more ways than one - the Plexiglass sides of the firm's server racks were warped, having begun to melt from the heat![16] The need to cool modern data centers draws a lot of power and that costs a lot of money.

The Chief Eco Officer of Sun Microsystems has claimed that computers draw four to five percent of the world's power. Google's Chief Technology Offer has said that the firm spends more to power its servers than the cost of the servers themselves[17]. Microsoft, Yahoo, and Google have all built massive data centers in the Pacific Northwest, away from their corporate headquarters, specifically choosing these locations for access to cheap hydroelectric power. Google's location in The Dalles, OR, is charged a cost per kilowatt hour of two cents by the local power provider, less than a fifth of the 11-cent rate the firm pays in Silicon Valley[18]. This means big savings for a firm that runs more than a million servers.

And while these powerful shrinking chips are getting hotter and more costly to cool, it's also important to realize that chips can't get smaller forever. At some point Moore's Law will run into the unyielding laws of nature. While we're not certain where these limits are, chip pathways certainly can't be shorter than a single molecule, and the actual physical limit is likely larger than that. Get too small and a phenomenon known as quantum tunneling kicks in, and electrons start to slide off their paths. Yikes!

---

[16] McGirt, 2007
[17] Kirkpatric, 2007
[18] Mehta, 2006

Buying Time

One way to overcome this problem is with *multi-core* technology. This involves putting two or more lower power processor cores (think of a core as the calculating part of a microprocessor) on a single chip. Philip Emma, IBM's Manager of Systems Technology and Microarchitecture, offers an analogy. Think of the traditional fast, hot single-core processors as a 300-pound lineman, and a dual-core processor as two 160-pound guys. Says Emma, "A 300-pound lineman can generate a lot of power, but two 160- pound guys can do the same work with less overall effort." [19] For many applications, the multicore chips will out perform a single speedy chip, while running cooler and drawing less power. Multicore processors are now mainstream.

By 2007, most PCs and laptops sold had at least a two-core (dual-core) processor. The Microsoft XBox 360 has three cores. The Playstation 3 includes the so-called *cell processor* developed by Sony, IBM, and Toshiba, that runs nine cores. By 2008, quad-core processors were common in high-end desktops and low-end servers. By 2010, AMD plans a 12-core PC chip.

Multicore processors can run older software written for single-brain chips. But they usually do this by using only one core at a time. To reuse the metaphor above, this is like having one of our 160-pound workers lift away, while the other one stands around watching. Multi-core operating systems can help achieve some performance gains. Versions of Windows or the Mac OS that are aware of multi-core processors can assign one program to run on one core, while a second app is assigned to the next core. But, in order to take full advantage of multi-core chips, applications needs to be rewritten to split up tasks so that smaller portions of a problem are executed simultaneously inside each core.

Writing code for this 'divide and conquer' approach is not trivial In fact, developing software for multi-core systems is described by Shahrokh Daijavad, software lead for next-generation computing systems at IBM, as "one of the hardest things you learn in computer science."[20] Microsoft's Chief Research and Strategy Officer has called coding for these chips "the most conceptually different [change] in the history of modern computing"[21]. Despite this challenge, some of the most aggressive adaptors of multi-core chips have been video game console manufacturers. Video game applications are particularly well suited for multiple cores since, for example, one core might be used to render the background, another to draw objects, another for the 'physics engine' that moves the objects around, and yet another to handle Internet communications for multi-player games.

Another approach to breathing life into Moore's Law is referred to as *stacked* or *three-dimensional semiconductors*. In this approach, engineers slice a flat chip into pieces, then reconnect the pieces vertically, making a sort of 'silicon sandwich'. The chips are both faster and cooler since electrons travel shorter distances. What was once an end-to-end trip on a conventional chip might just be a tiny movement up or down on a stacked chip. But stacking chips present their own challenges. In the same way that a skyscraper is more difficult and

[19] Ashton, 2005.
[20] Ashton, 2005.
[21] Copeland, 2008

costly to design and build than a ranch house, 3D semiconductors are tougher to design and manufacture. In Spring 2007, IBM announced that it had perfected 3D chipsets and would begin offering commercial versions of the technology in 2008. The first market for the stacked chips is in mobile phones, where IBM claims the technique improves power efficiency by up to 40 percent.

Quantum Leaps, Chicken Feathers, and the Indium Gallium Arsenide Valley?

Think about it – the triple threat of size, heat, and power means that Moore's Law, perhaps the greatest economic gravy train in history, will likely come to a grinding halt in your lifetime. Mutli-core and 3D semi-conductors are here today, but what else is happening to help stave off the death of Moore's Law?

Every once in a while a material breakthrough comes along that improves chip performance. A few years back researchers discovered that replacing a chip's aluminum components with copper could increase speeds up to thirty percent. Now scientists are concentrating on improving the very semiconductor material that chips are made of. While the silicon used in chips is wonderfully abundant (it has pretty much the same chemistry found in sand), researchers are investigating other materials that might allow for chips with even tighter component densities. Intel has demonstrated that chips made with the super-geeky-sounding semiconductor materials of indium gallium arsenide and indium aluminum arsenide can run faster and require less wattage than their silicon counterparts. Perhaps even more exotic (and downright bizarre), researchers at the University of Delaware have experimented with a faster-than-silicon material derived from chicken feathers! Hyper-efficient chips of the future may also be made out of carbon nanotubes, once the technology to assemble the tiny structures becomes commercially viable.

Other designs move away from electricity over silicon. Optical computing, where signals are sent via light rather than electricity, promises to be faster than conventional chips, if lasers can be mass-produced in miniature (silicon laser experiments show promise). Others are experimenting by crafting computing components using biological material (think a DNA-based storage device).

One yet-to-be-proven technology that could blow the lid off what's possible today is quantum computing. Conventional computing stores data as a combination of bits, where a bit is either a one or a zero. Quantum computers, leveraging principles of quantum physics, employ qubits that can be both one *and* zero at the same time. Add a bit to a conventional computer's memory and you double its capacity. Add a bit to a quantum computer and its capacity increases exponentially. For comparison, consider that a computer model of serotonin, a molecule vital to regulating the human central nervous system, would require 10 to the power of 94 bytes of information. Unfortunately there's not enough matter in the universe to build a computer that big. But modeling a serotonin molecule using quantum computing would take just 424 qubits[22].

Some speculate that quantum computers could one day allow pharmaceutical companies to create hyper-detailed representations of the human body that reveal drug side effects before they're even tested on humans. Quantum computing might also accurately predict the weather months in advance or offer unbreakable computer security. Ever have trouble placing a name with a face? A quantum computer linked to a camera, say, in your sunglasses could recognize the faces of anyone you've met and give you a heads up to their name and background[23]. Opportunities abound. Of course, before quantum computing can be commercialized, researchers need to harness the freaky properties of quantum physics wherein your answer may reside in another universe, or could disappear if observed (Einstein himself referred to certain behaviors in quantum physics as 'spooky action at a distance').

[22] Kaihla, 2004.
[23] Schwartz, Taylor and Koselka, 2006.

Pioneers in quantum computing include IBM, HP, NEC, and a Canadian startup named D-Wave. At a 2007 at the Computer History Museum in Mountain View, CA, D-Wave demonstrated Orion, a 16 qubit computer that could find a protein in a database, figure out the optimal wedding guest seating arrangements, and solve a Sudoku puzzle. Scientific opinion varied widely as to the significance of the D-Wave advance. The Orion was built using a chip cooled to minus 273 degrees Celsius in a bath of liquid helium and tasks were performed at speeds about 100 times slower than conventional PCs. Not exactly commercial stuff. But it was the most advanced quantum computing demonstration to date. If or when quantum computing becomes a reality is still unknown, but the promise exists that while Moore's Law may run into limits imposed by Mother Nature, a new way of computing may blow past anything we can do with silicon, continuing to make possible the once impossible.

LEARNING OBJECTIVE:
After studying this section you should be able to:
1. Give examples of the business use of supercomputing and grid computing.
2. Describe grid computing and discuss how grids transform the economics of supercomputing.
3. Understand the characteristics of problems that are and are not well suited for supercomputing and grid computing.

BRINGING BRAINS TOGETHER – SUPERCOMPUTING AND GRID COMPUTING

As Moore's Law makes possible the once impossible, businesses have begun to demand access to the world's most powerful computing technology. The term *supercomputer* is used to refer to computers that are among the fastest of any in the world at the time of their introduction[24]. Supercomputing used to be the domain of governments and high-end research labs, performing tasks such as simulating the explosion of nuclear devices, or analyzing large-scale weather and climate phenomena. But it turns out with a bit of tweaking, the algorithms used in this work are profoundly useful to business. Consider perhaps the world's most well-known supercomputer, IBM's Deep Blue. In May 1997, the computer rather controversially beat chess champion Gary Kasparov. While there is not a burning need for chess-playing computers in the world's corporations, it turns out that the computing algorithms to choose the best among multiple chess moves are similar to the math behind choosing the best combination of airline flights.

One of the first customers of Deep Blue technologies was United Airlines, which gained an ability to examine 350,000 flight path combinations for its scheduling systems – a figure well ahead of the previous limit of three thousand. Estimated savings through better yield management? Over $50 million! Finance found uses, too. An early adopter was CIBC (the Canadian Imperial Bank of Commerce), one of the largest banks in North America. Each morning CIBC uses a supercomputer to run its portfolio through Monte Carlo simulations that aren't all that different from the math used to simulate nuclear explosions. An early adopter of the technology, at the time of deployment, CIBC was the only bank that international regulators allowed to calculate its own capital needs rather than use boilerplate ratios. That cut capital-on-hand by hundreds of millions of dollars, a substantial percentage of the bank's capital, saving millions a year in funding costs.

Modern supercomputing is typically done via a technique called *massively parallel processing*. The fastest of these supercomputers are built using hundreds of microprocessors, all programmed

[24] A list of the current supercomputer performance champs can be found at Top500.org.

to work in unison as one big brain. While supercomputers use special electronics and software to handle the massive load, the processors themselves are often of the off-the-shelf variety that you'd find in a typical PC. Virginia Tech created what at the time was the world's third-fastest supercomputer by using chips from 1,100 Macintosh computers lashed together with off-the-shelf networking components. The total cost of the system was just $5.2 million, far less than the typical cost for such burly hardware. A similar effort by the University Illinois using Playstation 2's once ranked among the world's 500 fastest machines.

A new development known as *grid computing* is further transforming the economics of supercomputing. With grid computing, firms place special software on its existing PCs or servers that enables these computers to work together on a common problem. Large organizations may have thousands of PCs, but they're not necessarily being used all the time, or at full capacity. With grid software installed on them, these idle devices can be marshaled to attack portions of a complex task as if they collectively were one massively parallel supercomputer. This radically changes the economics of high-performance computing. *BusinessWeek* reports that while a middle-of-the-road supercomputer could run as much as $30 million, grid computing software and services to perform comparable tasks can cost as little as $25,000, assuming an organization already has PCs and servers in place.

An early pioneer in grid computing is the biotech firm Monsanto. Monsanto enlists computers to explore ways to manipulate genes to create crop strains that are resistant to cold, drought, bugs, pesticides, or that are more nutritious. Previously with even the largest computer Monsanto had in-house, gene analysis was taking 6 weeks and the firm was able to analyze only 10 to 50 genes a year. But by leveraging grid computing, Monsanto has reduced gene analysis to less than a day. The 50-fold time savings now lets the firm consider thousands of genetic combinations in a year[25]. Lower R&D time means faster time to market – critical to both the firm and its customers.

Grids are now everywhere. Movie studios use them to create special effects and animated films. Proctor & Gamble has used grids to redesign the manufacturing process for Pringles potato chips. GM and Ford use grids to simulate crash tests, saving millions in junked cars and speeding time-to-market. Pratt and Whitney test aircraft engine designs on a grid. And biotech firms including Aventis, GlaxoSmithKlein, and Pfizer push their research through a quicker pipeline by harnessing grid power. JP Morgan Chase even launched a grid effort that mimics CIBC's supercomputer, but at a fraction of the latter's cost. By the second year of operation, the JP Morgan Chase grid was saving the firm $5 million a year.

You can join a grid, too. SETI@Home turns your computer screen saver into a method to help "search for extraterrestrial intelligence", analyzing data from the Arecibo radio telescope system in Puerto Rico (no E.T. spotted yet). FightAids@Home will enlist your PC to explore AIDS treatments. Folding@Home is an effort by Stanford researchers to understanding the science of protein-folding within diseases such as Alzheimer's, cancer, and cystic fibrosis. A version of Folding@Home software for the PS3 had enlisted over half a million consoles by mid 2007. Having access to these free resources is an enormous advantage for researchers. Says the director

[25] Hamm, 2004

of Folding@Home "Even if we were given all of the NSF supercomputing centers combined for a couple of months, that is still fewer resources than we have now"[26].

Multi-core, massively parallel, and grid computing are all related in that each attempts to lash together multiple computing devices so that they can work together to solve problems. Think of multi-core chips as having several processors in a single chip. Think of massively parallel supercomputers as having several chips in one computer, and think of grid computing as using existing computers to work together on a single task. While these technologies offer great promise, they're all subject to the same limitation – software must written to divide existing problems into smaller pieces that can be handled by each core, processor, or computer, respectively. Some problems, such as simulations, are easy to split up, but for problems that are linear (where, say, step 2 can't be completed until the results from step 1 are known), the multiple-brains approach doesn't offer much help.

Moore's Law will likely hit its physical limit in your lifetime, but no one really knows if this "Moore's Wall" is a decade away or more. What lies ahead is anyone's guess. Some technologies, such as still-experimental quantum computing, could make computers that are more powerful than all the world's conventional computers combined. Think strategically – new waves of innovation might soon be shouting "surf's up!"

LEARNING OBJECTIVE:
After studying this section you should be able to:

1. Understand the magnitude of the environmental issues caused by rapidly obsolete, faster / cheaper computing.
2. Explain the limitations of approaches attempting to tackle eWaste.
3. Understand the risks firms are exposed to when not fully considering the lifecycle of the products they sell or consume.
4. Ask questions that expose concerning ethical issues in a firm or partner's products and processes, and help the manager behave more responsibly.

E-WASTE: THE DARK SIDE OF MOORE'S LAW

We should celebrate the great bounty Moore's Law and the tech industry bestow on our lives. Costs fall, workers become more productive, innovations flourish, and we gorge at a buffet of digital entertainment that includes music, movies, and games. But there is a dark side to this faster / cheaper advancement. A PC has an expected lifetime of three to five years. A cell phone? Two years or less. Rapid obsolescence means the creation of every growing mountains of discarded tech junk, known as electronic waste or *e-waste*. According to the EPA, in 2005, the US alone generated over 2.6 million tons of e-waste, and the results aren't pretty. Consumer electronics and computing equipment can be a toxic cocktail that includes cadmium, mercury, lead, and other hazardous materials. Once called the 'effluent of the affluent', e-waste will only increase with the rise of living standards worldwide.

The quick answer would be to recycle this stuff. Not only does e-waste contain mainstream recyclable materials we're all familiar with, like plastics and aluminum, it also contains small

[26] Johnson, 2002

bits of increasingly valuable metals such as silver, platinum, and copper. In fact, there's more gold per pound of discarded tech equipment, than in mined ore[27]. But as the sordid record of e-waste management shows, there's often a disconnect between consumers and managers who *want* to do good and those efforts that are *actually* doing good. The complexities of the modern value chain, the vagaries of international law, and the nefarious actions of those willing to put profits above principle show how difficult addressing this problem will be.

The process of separating out the densely packed materials inside tech products so that the value in e-waste can be effectively harvested is extremely labor intensive, more akin to reverse manufacturing than any sort of curbside recycling efforts. Sending e-waste abroad can be ten times cheaper than dealing with it at home[28], so it's not surprising that up to 80 percent of the material dropped off for recycling is eventually exported[29]. Much of this waste ends up in China, South Asia, or Nigeria, where it is processed in dreadful conditions.

Consider the example of Guiyu, China, a region whose poisoning has been extensively chronicled by organizations such as the Silicon Valley Toxics Coalition, the Basel Action Network (BAN), and Greenpeace. Workers in and around Guiyu toil without protective equipment, breathing clouds of toxins generated as they burn the plastic skins off of wires to get at the copper inside. Others use buckets, pots, or wok-like pans (in many cases the same implements used for cooking) to sluice components in acid baths to release precious metals, recovery processes that create even more toxins. Waste sludge and the carcasses of what's left over are most often dumped in nearby fields and streams. Water samples taken in the region showed lead and heavy metal contamination levels some 400 to 600 times greater than what international standards deem safe[30]. The area is so polluted that drinking water must be trucked in from 18 miles away.

Photos from Guiyu, China from the Basel Action Network and posted on CrunchGear[31]

China cares about its environment. The nation has banned the importing of e-waste since 2000[32]. But corruption ensures that e-waste continues to flow into the country. According to one exporter, all that's required to get e-waste past customs authorities is to tap a $100 bill on the side of the container[33]. Well-meaning US recyclers, as well as those attempting to collect technology for reuse in poorer countries, are often in the dark as to where their products end up. The trade is often brokered by middlemen who mask the eventual destination and fate of the

[27] Kovessey, 2008.
[28] Bodeen, 2007.
[29] Royte, 2006.
[30] Grossman, 2006.
[31] Biggs, 2008
[32] Grossman, 2006.
[33] Bodeen, 2007.

products purchased. BAN investigators in Lagos, Nigeria, documented mountains of e-waste with labels from schools, US government agencies, and even some of the world's largest corporations. And despite Europe's prohibition on exporting e-waste, many products originally labeled for repair and reuse end up in toxic recycling efforts. Even among those products that gain a second or third life in developing nations, the inevitable is simply postponed, with e-waste almost certain to end up in landfills that lack the protective groundwater barriers and environmental engineering of their industrial-nation counterparts. The reality is e-waste management is extraordinarily difficult to monitor and track, and loopholes are rampant.

Thinking deeply about the ethical consequences of a firm's business is an imperative for the modern manager. A slip-up (intentional or not) can, in seconds, be captured by someone with a cell phone, uploaded to YouTube, or offered in a blog posting for the world to see. When Dell was caught using Chinese prison labor as part of its recycling efforts, one blogger chastised the firm with a tweak of its marketing tagline, posting "Dude, you're getting a cell[34]". The worst cases expose firms to legal action and can tarnish a brand for years. Big firms are big targets, and environmentalists have been quick to push the best known tech firms to take back their products for responsible recycling, and to eliminate the worst toxins from their offerings.

Consider that even Apple, Inc. (where Al Gore sits on the firm's Board of Directors), has been pushed by a coalition of environmental groups on all of these fronts. Critics have shot back that signaling out Apple is unfair. The firm was one of the first computer companies to eliminate lead-lined glass monitors from its product line, and has been a pioneer of reduced-sized packaging that leverage recyclable materials. But if the firm that counts Al Gore among its advisors can get tripped up on green issues, all firms are vulnerable.

Environmentalists see this pressure to deal with e-waste as yielding results: Apple and most other tech firms have continually moved to eliminate major toxins from their manufacturing processes. All this demonstrates that today's business leaders have to be far more attuned to the impact not only of their own actions, but also to those of their suppliers and partners. How were products manufactured? Using which materials? Under what conditions? What happens to items when they're discarded? Who provides collection and disposal? It also shows the futility of legislative efforts that don't fully consider and address the problems they are meant to target.

Which brings us back to Gordon Moore. To his credit, Moore is not just the founder of the world's largest microprocessor firm and first to identify the properties we've come to know as Moore's Law, he has also emerged as one of the world's leading supporters of environmental causes. The generosity of the Gordon and Betty Moore foundation includes, among other major contributions, the largest single gift to a private conservation organization. Indeed, Silicon Valley, while being the birthplace of products that become e-waste, also promises to be at the forefront of finding solutions to modern environmental challenges. The Valley's leading venture capitalists, including Sequoia and Kleiner Perkins (where Al Gore is now a partner), have started multi-million-dollar green investment funds, targeted at funding the next generation of sustainable, environmental initiatives.

• • •

[34] Russell, 2003 and http://laughingmeme.org/2003/03/23/dell-recycling-a-ways-to-go-still/

About the Author:

John Gallaugher is a member of the Dept. of Information Systems in Boston College's Carroll School of Management. Prof. Gallaugher teaches courses and conducts research at the intersection of technology and strategy. He leads the School's TechTrek programs, co-leads the Asian field study program, and has consulted to and taught executive seminars for several organizations including Accenture, Alcoa, Brattle Group, ING Group, Patni Computer Systems, Staples, State Street, and the U.S. Information Agency. Writings, podcasts, course material, and research by Prof. Gallaugher can be found online at www.gallaugher.com.

This reading was specifically designed for faculty to use in their classes. Enjoy! If you do use it, please send an e-mail to john.gallaugher@bc.edu. More chapters and cases will follow in Professor Gallaugher's forthcoming book "Information Systems: A Manager's Guide to Harnessing Technology", to published (in both free online and low-cost print version) by Flat World Knowledge (FlatWorldKnowledge.com). Thanks!

CHAPTER REFERENCES

Andrews, Robert, "A Decade After Kasparov's Defeat, Deep Blue Coder Relives Victory," *Wired*, May 11, 2007.

Ashton, Adam, "More Life for Moore's Law," *BusinessWeek*, June 20, 2005.

Biggs, John, "Guiyu, the e-Waste Capital of China", CrunchGear, April 4, 2008.

Bodeen, Christopher, "In 'e-waste' heartland, a toxic China", *International Herald Tribune*, Nov. 18, 2007.

Breen, Bill, "Living in Dell Time," *Fast Company*, November 24, 2004.

Bulkeley, William, "IBM Touts Breakthrough in 3-D Chips," *Wall Street Journal*, April 12, 2007, p. B3.

Caulfield, Brian, "Apple Promises to Get Greener", Forbes, May 2, 2007.

Copeland, Michael, "How to Ride the Fifth Wave," *Business 2.0*, July 1, 2005.

Copeland, M., "A Chip Too Far?", *Fortune*, Sept. 1, 2008.

Corbett, Sara, "Can the Cellphone Help End Global Poverty?", *The New York Times Magazine*, April 13, 2008.

Ewing, Jack, "Upwardly Mobile In Africa", *BusinessWeek*, Sept. 24, 2007, pp. 64-71.

"Faster, Cheaper, Better," *The Economist*, July 22, 2004.

"Goodbye to Analog TVs - the Coming Toxic Tidal Wave", Business & the Environment with ISO 14000, Vol. 19, No. 5., May 2008, pp. 1-3.

Greene, Kate, "Intel Looks Beyond Silicon", *Technology Review*, December 11, 2007.

Grossman, Elizabeth, "Where Computers Go to Die – and Kill", Salon.com, April 10, 2006.

Hamblen, Matthew, " J.P. Morgan Harnesses Power with Grid Computing System," *ComputerWorld*, March 15, 2004.

Hamm, Steve, "Getting A Grip On Grid Computing," *BusinessWeek*, Oct. 18, 2004.

Huggins, J., "How Much Data is That?", *Refrigerator Door*, Aug. 19, 2008.

Jacobson, Louis, "Can Computers Fly on the Wings of a Chicken?", *The Washington Post*, July 8, 2002, p. A07.

Johnson, G., “Supercomputing ‘@Home’ is Paying Off”, *The New York Times*, April 23, 2002.

Kaihla, Paul, “Quantum Leap”, *Business 2.0*, August 1, 2004.

Kirkpatrick, D., “The greenest computer company under the Sun”, April 13, 2007.

Kovessy, Peter, “How to Trash Toxic Tech”, *Ottawa Business Journal*, May 12, 2008.

Little, Darnell, " Who Needs Supercomputers?", *BusinessWeek*, June 3, 2002.

McGirt, Ellen, "Hacker, Dropout, CEO," *Fast Company*, May 2007.

Mehta, Stephanie, "Behold the Server Farm," *Fortune*, August 1, 2006.

Miller, J., “Goodbye GUI? Ambient orb a computer 'mood ring'”, *Mass High Tech*, Feb. 10, 2003.

Mooallem, Jon, “The Afterlife of Cellphones”, *The New York Times*, Jan. 13, 2008.

Pontin, Jason, “A Giant Leap Forward in Computing? Maybe Not”, The New York Times, April 8, 2007

Pullen, Dean, " Folding@Home embraces Playstation 3," *The Inquirer*, August 24, 2006.

Quinn, Michelle, and Chmielewski, Dawn, “Top music seller’s store has no door,” *The Los Angeles Times*, April 04, 2008.

Royte, Elizabeth, “E-Waste@Large”, *The New York Times*, Jan. 27, 2006.

Russell, J., “Dell under attack over using prison labour”, The Inquirer, Jan. 10, 2003.

Schuman, E., “At Wal-Mart, Worlds Largest Retail Data Warehouse Gets Even Larger”, eWeek, Oct. 13, 2004.

Schwartz, Peter, Taylor, Chris, and Koselka, Rita, “The Future of Computing: Quantum Leap”, *Fortune*, Aug. 2, 2006.