



This is “Moore’s Law: Fast, Cheap Computing and What It Means for the Manager”, chapter 5 from the book [Getting the Most Out of Information Systems \(index.html\)](#) (v. 1.4).

This book is licensed under a [Creative Commons by-nc-sa 3.0](http://creativecommons.org/licenses/by-nc-sa/3.0/) license. See the license for more details, but that basically means you can share this book as long as you credit the author (but see below), don't make money from it, and do make it available to everyone else under the same terms.

This content was accessible as of December 29, 2012, and it was downloaded then by [Andy Schmitz](#) (<http://lardbucket.org>) in an effort to preserve the availability of this book.

Normally, the author and publisher would be credited here. However, the publisher has asked for the customary Creative Commons attribution to the original publisher, authors, title, and book URI to be removed. Additionally, per the publisher's request, their name has been removed in some passages. More information is available on this project's [attribution page](http://2012books.lardbucket.org/attribution.html?utm_source=header).

For more information on the source of this book, or why it is available for free, please see [the project's home page](#) (<http://2012books.lardbucket.org/>). You can browse or download additional books there.

Chapter 5

Moore's Law: Fast, Cheap Computing and What It Means for the Manager

5.1 Introduction

LEARNING OBJECTIVES

1. Define Moore's Law and understand the approximate rate of advancement for other technologies, including magnetic storage (disk drives) and telecommunications (fiber-optic transmission).
2. Understand how the price elasticity associated with faster and cheaper technologies opens new markets, creates new opportunities for firms and society, and can catalyze industry disruption.
3. Recognize and define various terms for measuring data capacity.
4. Consider the managerial implication of faster and cheaper computing on areas such as strategic planning, inventory, and accounting.

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**¹, after Intel cofounder, 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 more powerful chips to be manufactured at cheaper prices. G. Moore, “Cramming More Components onto Integrated Circuits,” *Electronics Magazine*, April 19, 1965.

Moore's friend, legendary chip entrepreneur and CalTech professor Carver Mead, later coined the “Moore's Law” moniker. That name sounded 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 eighteen months*. (Moore's original paper stated transistors per chip, a proxy for power, would double every two years, but many sources today refer to the *eighteen*-month figure, so we'll

1. Chip performance per dollar doubles every eighteen months.

stick with that—either way, we're still talking about ridiculously accelerating power and plummeting costs.)

Moore's Law applies to chips—broadly speaking, to *processors*, or the electronics stuff that's made out of silicon. Although other materials besides silicon are increasingly being used. The **microprocessor**² is the brain of a computing device. It's the part of the computer that executes the instructions of a computer program, allowing it to run 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 *eighteen* months, but cost the same as today's models (or from another perspective, in a year and a half, chips that are same speed as today's models should be available for half the price).

Random-access memory (RAM)³ is chip-based memory. The RAM inside your personal computer is **volatile memory**⁴, meaning that when the power goes out, all is lost that wasn't saved to **nonvolatile memory**⁵ (i.e., a more permanent storage media like a hard disk or flash memory). Think of RAM as temporary storage that provides fast access for executing computer programs and files. When you "load" or "launch" a program, it usually moves from your hard drive to those RAM chips, where it can be more quickly executed by the processor.


Cameras, MP3 players, USB drives, and mobile phones often use **flash memory**⁶ (sometimes called *flash RAM*). It's not as fast as the RAM used in most traditional PCs, but holds data even when the power is off (so flash memory is also nonvolatile memory). You can think of flash memory as the chip-based equivalent of a hard drive. In fact, flash memory prices are falling so rapidly that several manufactures including Apple and the One Laptop per Child initiative (see the "Tech for the Poor" sidebar later in this section) have begun offering chip-based, nonvolatile memory as an alternative to laptop hard drives. The big advantage? Chips are **solid state electronics**⁷ (meaning no moving parts), so they're less likely to fail, and they draw less power. The solid state advantage also means that chip-based MP3 players like the iPod nano make better jogging companions than hard drive players, which can skip if jostled. For RAM chips and flash memory, Moore's Law means that in *eighteen* months you'll pay the same price as today for twice as much storage.

Computer chips are sometimes also referred to as **semiconductors**⁸ (a substance such as silicon dioxide used inside most computer chips that is capable of enabling as well as inhibiting the flow of electricity). So if someone refers to the *semiconductor industry*, they're talking about the chip business. 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.

2. The part of the computer that executes the instructions of a computer program.
3. The fast, chip-based volatile storage in a computing device.
4. Storage (such as RAM chips) that is wiped clean when power is cut off from a device.
5. Storage that retains data even when powered down (such as flash memory, hard disk, or DVD storage).
6. Nonvolatile, chip-based storage, often used in mobile phones, cameras, and MP3 players. Sometimes called flash RAM, flash memory is slower than conventional RAM, but holds its charge even when the power goes out.
7. Semiconductor-based devices. Solid state components often suffer fewer failures and require less energy than mechanical counterparts because they have no moving parts. RAM, flash memory and microprocessors are solid state devices. Hard drives are not.
8. A substance such as silicon dioxide used inside most computer chips that is capable of enabling as well as inhibiting the flow of electricity. From a managerial perspective, when someone refers to semiconductors, they are talking about computer chips, and the semiconductor industry is the chip business.

Strictly speaking, Moore's Law does not apply to other technology components. But other computing components are also seeing their price versus performance curves skyrocket exponentially. Data storage doubles every twelve 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 squirted over an **optical fiber line**⁹ can double every nine months. 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. These numbers should be taken as rough approximations and shouldn't be expected to be strictly precise over time. However, they are useful as rough guides regarding future computing price/performance trends. Despite any fluctuation, it's clear that the price/performance curve for many technologies is exponential, offering astonishing improvement over time.

Figure 5.1 Advancing Rates of Technology (Silicon, Storage, Telecom)



Source: Adopted from Shareholder Presentation by Jeff Bezos, Amazon.com, 2006.

- 9. A high-speed glass or plastic-lined networking cable used in telecommunications.
- 10. The rate at which the demand for a product or service fluctuates with price change. Goods and services that are highly price elastic (e.g., most consumer electronics) see demand spike as prices drop, whereas goods and services that are less price elastic are less responsive to price change (think heart surgery).

Get Out Your Crystal Ball

Faster and cheaper makes possible the once impossible. 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**¹⁰ kicks in. Tech

products are highly *price elastic*, meaning consumers buy more products as they become cheaper. As opposed to goods and services that are *price inelastic* (like health care and housing), which consumers will try their best to buy even if prices go up. 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. M. Copeland, "How to Ride the Fifth Wave," *Business 2.0*, July 1, 2005. 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 as 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 and 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 have become ubiquitous—woven into products in ways few imagined years before. Silicon is everywhere! It provides the smarts in the world's billion-plus mobile phones. It's in the throwaway radio frequency identification (RFID) tags that track your luggage at the airport. 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 giveaway as part of its primary product, mobile phones. This shift has occurred with such sweeping impact that former photography giants Kodak and the now-merged Konica-Minolta have exited the camera business.

Ambient Devices and the Fifth Wave

Pritesh Gandhi almost never gets caught in the rain without his umbrella. That's because Gandhi'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." Gandhi is CEO of "fifth wave" firm Ambient Devices, a Massachusetts start-up that's embedding computing and communications technology into everyday devices in an attempt to make them "smarter" and more useful (the weather-sensing umbrella was developed while he helmed the firm).

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 enabled by 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 your stock portfolio or the intensity of the local pollen count. In just six months, the ten refugees from MIT's Media Lab that founded Ambient Devices took the idea for the Orb, designed the device and its software, and licensed wireless spectrum from a pager firm that had both excess capacity and a footprint to cover over 90 percent of the United States. M. Copeland, "How to Ride the Fifth Wave," *Business 2.0*, July 1, 2005; and J. Miller, "Goodbye G.U.I? Ambient Orb a Computer 'Mood Ring,'" *Mass High Tech*, February 10, 2003.

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 until 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.

Figure 5.2



Products developed by "fifth wave" firm Ambient Devices include the weather-reading Ambient Umbrella, the Energy Joule, a seven-day forecaster, and the Orb lamp.

Source: Used with permission from Ambient Devices.

Moore's Law inside Your Medicine Cabinet

Moore's Law is about to hit your medicine cabinet, and several early-stage efforts show the potential for fifth-wave computing to improve healthcare quality while lowering costs. The GlowCap from Vitality, Inc., is a "smart" pill bottle that will flash when you're supposed to take your medicine. It will play a little tune if you're an hour late for your dose and will also squirt a signal to a night-light that flashes as a reminder (in case you're out of view of the cap). GlowCaps can also be set to call or send a text if you haven't responded past a set period of time. And the device will send a report to you, your doc, or whomever else you approve. The GlowCap can even alert your pharmacy when it's time for refills. The bottles sell for as little as \$10 but in some cases are likely to be free. The business case for that? The World Health Organization estimates drug adherence at just 50 percent, and analysts estimate that up to \$290 billion in increased medical costs are due to patients missing their meds. Vitality CEO David Rose (who incidentally also cofounded Ambient Devices) recently cited a test in which GlowCap users reported a 98 percent medication adherence rate. D. Rose, presentation as part of "From Disruption to Innovation" at the MIT Enterprise Forum, Cambridge, MA, June 23, 2010.

Figure 5.3



The GlowCap from Vitality, Inc., will flash, beep, call, and text you if you've skipped your meds. It can also send reports to you, your doctor, and your loved ones and even notify your pharmacy when it's time for a refill.

Source: Used with permission from Vitality, Inc.

And there might also be a chip inside the pills, too! Proteus, a Novartis-backed venture, has developed a sensor made of food and vitamin materials that can be swallowed in medicine. The sensor is activated and powered by the body's digestive acids (think of your stomach as a battery). Once inside you, the chip sends out a signal with vitals such as heart rate, body angle, temperature, sleep, and more. A waterproof skin patch picks up the signal and can wirelessly relay the pill's findings when the patient walks within twenty feet of their phone. Proteus will then compile a report from the data and send it to their mobile device or e-mail account. The gizmo's already in clinical trials for heart disease, hypertension, and tuberculosis and for monitoring psychiatric illnesses.E. Landau, "Tattletale Pills, Bottles Remind You to Take Your Meds," *CNN*, February 2, 2010. And a pill with built-in smarts can identify itself to help guard against taking counterfeit drugs, a serious worldwide concern. Pills that chat with mobile phones could help promote telemedicine, bringing health care to hard-to-reach rural populations. And games and social apps based on this information can provide motivating, fun ways to nudge patients into healthy habits. The CEO of Proteus Health says that soon you may be able to think of your body as "the ultimate game controller."K. Rozendal, "The Democratic, Digital Future of Healthcare," *Scope*, May 13, 2011.

One of the most agile surfers of this *fifth wave* is Apple, Inc.—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 keen insight on where trends in computing power and performance are headed is captured in this quotation by the firm's co-founder, the late Steve Jobs: "There's an old Wayne Gretzky quote that I love. 'I skate to where the puck is going to be, not where it has been.' And we've always tried to do that at Apple. Since the very very beginning. And we always will."H. Mitchell, "Steve Jobs Used Wayne Gretzky as Inspiration," *Los Angeles Times*, October 6, 2011. The curves above aren't perfect, but they can help managers see where the puck is headed, helping the savvy manager predict the future, plan for the impossible to become possible, and act as the disruptor rather than the disrupted.

Apple's breakout resurgence owes a great deal to the iPod. At launch, the original iPod sported a 5 GB 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 fifty dollars less than the original, yet held *forty times* the songs. By that time the firm had sold over one hundred fifty million iPods—an adoption rate faster than the original Sony Walkman. Apple’s high-end models have morphed into Internet browsing devices capable of showing maps, playing videos, and gulping down songs from Starbucks’ Wi-Fi 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, iPad, and iTunes. As an online store, iTunes is always open. iTunes regularly sells tens of millions of songs on Christmas Day alone, a date when virtually all of its offline competition is closed for the holiday. In a short five years after its introduction, iTunes has sold over 4 billion songs and has vaulted past retail giants Wal-Mart, Best Buy, and Target to become the number one music retailer in the world. 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 sixteenfold 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.

Table 5.1 Top U.S. 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 physical music CD now sells more than anyone else.			

Source: Michelle Quinn and Dawn C. Chmielewski, “Top Music Seller’s Store Has No Door,” *Los Angeles Times*, April 4, 2008.

Table 5.2 Tech's Price/Performance Trends in Action: Amazon Kindle and Apple Music Storage

Amazon Kindle		Apple	
First Generation	Fourth Generation	iPod	iCloud
250 MB	2 GB	5 GB	5 GB
November 2007	September 2011	October 2001	October 2011
\$399	\$79	\$399	Free

Amazon's first Kindle sold for nearly \$400. Less than four years later, Amazon was selling an updated version of the Kindle for one-fifth that price. Similarly, Apple offered 5 GB of music storage in the original iPod (also priced at roughly \$400). By the iPod's tenth birthday, Apple was giving away 5 GB of storage (for music or other media) for free via its iCloud service. Other factors influence price drops, such as being able to produce products and their components at scale, but Moore's Law and related price/performance trends are clearly behind the price decreases we see across a wide variety of tech products and services.

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 and 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 [Note 5.14 "Bits and Bytes"](#)) 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 7 percent sales increase over nonsearchable books. When "Search Inside the Book" launched, the database to support this effort was 20 TB in size. In just eight years, the firm found that it made good business sense to launch an effort that was a full *twenty times* larger than anything used by *any* firm less than a decade earlier. And of course, all of these capacities seem laughably small by today's standards. (See [Chapter 11 "The Data Asset: Databases, Business Intelligence, and Competitive Advantage"](#).) For Amazon, the impossible had not just become possible; it became good business. By 2009, digital books weren't just for search; they were for sale. Amazon's Kindle reader (a Moore's Law marvel sporting a microprocessor and flash storage) became the firm's top-selling product in terms of both unit sales and dollar volume. The real business opportunity for Amazon isn't Kindle as a consumer electronics device but as an ever-present, never-closing store, which also provides the firm with a migration path from atoms to bits. (For more on

that topic, see [Chapter 4 "Netflix in Two Acts: The Making of an E-commerce Giant and the Uncertain Future of Atoms to Bits"](#).) By 2011, Amazon (by then the largest book retailer in North America) reported that it was selling more electronic books than print ones. M. Hamblen, "Amazon: E-Books Now Outsell Print Books," *ComputerWorld*, May 19, 2011. Apple's introduction of the iPad, complete with an iBook store, shows how Moore's Law rewrites the boundaries of competition—bringing a firm that started as a computer retailer and a firm that started as an online bookstore in direct competition with one another.

Bits and Bytes

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 a thousand characters, megabyte = 1 million, gigabyte = 1 billion, terabyte = 1 trillion, petabyte = 1 quadrillion, and exabyte = 1 quintillion bytes.

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

These are managerial definitions, but technically, a kilobyte is 2^{10} or 1,024 bytes, mega = 2^{20} , giga = 2^{30} , tera = 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. E. Schuman, “At Wal-Mart, World’s Largest Retail Data Warehouse Gets Even Larger,” *eWeek*, October 13, 2004; and J. Huggins, “How Much Data Is That?” *Refrigerator Door*, August 19, 2008.

Table 5.3 Bytes Defined

	Managerial Definition	Exact Amount	To Put It in Perspective
1 Byte	One keyboard character	8 bits	1 letter or number = 1 byte
1 Kilobyte (KB)	One thousand bytes	2^{10} bytes	1 typewritten page = 2 KB
			1 digital book (Kindle) = approx. 500–800 KB
1 Megabyte (MB)	One million bytes	2^{20} bytes	1 digital photo (7 megapixels) = 1.3 MB
			1 MP3 song = approx. 3 MB
			1 CD = approx. 700 MB
1 Gigabyte (GB)	One billion bytes	2^{30} bytes	1 DVD movie = approx. 4.7 GB
			1 Blu-ray movie = approx. 25 GB

	Managerial Definition	Exact Amount	To Put It in Perspective
1 Terabyte (TB)	One trillion bytes	2^{40} bytes	Printed collection of the Library of Congress = 20 TB
1 Petabyte (PB)	One quadrillion bytes	2^{50} bytes	eBay data warehouse (2010) = 10 PBC. Monash, “eBay Followup—Greenplum Out, Teradata > 10 Petabytes, Hadoop Has Some Value, and More,” October 6, 2010. Note eBay plans to increase this value 2.5 times by the end of 2011.
1 Exabyte (EB)	One quintillion bytes	2^{60} bytes	
1 Zettabyte (ZB)	One sextillion bytes	2^{70} bytes	Amount of data consumed by U.S. households in 2008 = 3.6 ZB

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. B. Breen, “Living in Dell Time,” *Fast Company*, November 24, 2004. 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 and budgeting perspective, as a manager you’ll need to consider a number of questions: 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 (One Laptop per Child) project aims to deliver education to children in the world's poorest communities via ultralow-cost computing devices that the firm has developed. The first offering, the XO laptop, cost less than \$200, although a sub-\$100 tablet is in the works. The XO sports a rubberized keyboard and entirely solid-state design (flash RAM rather than hard drive) that helps make the machine durable. The XO's ultrabright 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. 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 nonexistent, several battery-charging power generation schemes are offered, including a hand crank and foldout flexible solar panels. The OLPC Foundation delivered over 2.4 million laptops to children in twenty-four countries. C. Lawton, "The X.O. Laptop Two Years Later," *Wired*, June 19, 2009, <http://laptop.org/map>. The XO is a product made possible by the rapidly falling price of computing.

Figure 5.4
The XO PC



Source: Used with permission from fuseproject.

While the success of the OLPC effort will reveal itself over time, 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 three 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 Vodafone and Indian telecom provider Spice sell for \$25 or less. While it took roughly twenty 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 percent of the world's population lives within cellular network range (double the 2000 level), and the vast majority of mobile subscriptions are in developing countries. S. Corbett, "Can the Cellphone Help End Global Poverty?" *New York Times Magazine*, April 13, 2008.

Why such demand? Mobiles change lives for the better. According to Columbia economist Jeffrey Sachs, "The cell phone is the single most transformative technology for world economic development." J. Ewing, "Upwardly Mobile in Africa," *BusinessWeek*, September 24, 2007, 64–71. 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 taking off work to make the costly trek to a remote clinic 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. S. Corbett, "Can the Cellphone Help End Global Poverty?" *New York Times Magazine*, April 13, 2008.

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 8 percent even though consumer prices for fish *dropped* 4 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. The phone-equipped fleet now see more consistent pricing, spreading their catch more evenly whereas previous fisherman often inefficiently clustered in one market, overserving one population while underserving another. A London Business School study found that for every ten mobile phones per one hundred people, a country's GDP bumps up 0.5 percent. J. Ewing, "Upwardly Mobile in Africa," *BusinessWeek*, September 24, 2007, 64–71.

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 two hundred and fifty thousand Bangladeshi "phone ladies" to start businesses that helped their communities become more productive. Phone ladies bought a phone and special antenna on microcredit, allowing them to become a sort of village operator and charge a small commission for sending and receiving calls. Through phone ladies, the power of the mobile reaches even those too poor to afford buying one outright. Grameen Phone grew, reaching revenues of over \$1 billion, and became Bangladesh's largest telecom provider.

In another ingenious scheme, phone minutes become a proxy for currency. The *New York Times* reports that a person "working in Kampala, for instance, who wishes to send the equivalent of five dollars back to his mother in a village will buy a five-dollar 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. [The operator] then uses the airtime for her phone and completes the transaction by giving the man's mother the money, minus a small commission." S. Corbett, "Can the Cellphone Help End Global Poverty?" *New York Times Magazine*, April 13, 2008.

When phones can be used as currency for purchases or payments, who needs Visa or, for that matter, a wallet and cash? The M-PESA mobile banking service run by Kenya's Safaricom, allows customers to transfer cash using text messages. The majority of Kenyan adults have M-PESA accounts, paying for everything from groceries to cab rides to school tuition. The service can also allow family members to quickly and securely send cash across the country. Only 4 million Kenyans have traditional bank accounts, but 17 million use M-PESA, and upwards of 11 percent of Kenya's GDP flows through the service. B. Greeley and E. Ombok, "In Kenya, Security Cash on a Cell Phone," *BusinessWeek*, September 8, 2011. M-PESA is spreading to other regions and is even used in Afghanistan, while similar schemes are offered by firms such as WIZZIT in South Africa and GCASH in the Philippines. The "mobile phone as bank" may bring banking to a billion unserved customers in a few years.

KEY TAKEAWAYS

- Moore's Law applies to the semiconductor industry. The widely accepted managerial interpretation of Moore's Law states that for the same money, roughly eighteen months from now you should be able to purchase computer chips that are twice as fast or store twice as much information. Or over that same time period, chips with the speed or storage of today's chips should cost half as much as they do now.
- Nonchip-based technology also advances rapidly. Disk drive storage doubles roughly every twelve months, while equipment to speed transmissions over fiber-optic lines has doubled every nine months. While these numbers are rough approximations, the price/performance curve of these technologies continues to advance exponentially.
- These trends influence inventory value, depreciation accounting, employee training, and other managerial functions. They also help improve productivity and keep interest rates low.
- From a strategic perspective, these trends suggest that what is impossible from a cost or performance perspective today may be possible in the future. This fact provides an opportunity to those who recognize and can capitalize on the capabilities of new technology. As technology advances, new industries, business models, and products are created, while established firms and ways of doing business can be destroyed.
- Managers must regularly study trends and trajectory in technology to recognize opportunity and avoid disruption.

QUESTIONS AND EXERCISES

1. What is Moore's Law? What does it apply to?
2. Are other aspects of computing advancing as well? At what rates?
3. What is a microprocessor? What devices do you or your family own that contain microprocessors (and hence are impacted by Moore's Law)?
4. What is a semiconductor? What is the substance from which most semiconductors are made?
5. How does flash memory differ from the memory in a PC? Are both solid state?
6. Which of the following are solid state devices: an iPod shuffle, a TiVo DVR, a typical laptop PC?
7. Why is Moore's Law important for managers? How does it influence managerial thinking?
8. What is price elasticity? How does Moore's Law relate to this concept? What's special about falling chip prices compared to price drops for products like clothing or food?
9. Give examples of firms that have effectively leveraged the advancement of processing, storage, and networking technology.
10. What are the five waves of computing? Give examples of firms and industries impacted by the fifth wave.
11. As Moore's Law advances, technology becomes increasingly accessible to the poor. Give examples of how tech has benefited those who likely would not have been able to afford the technology of a prior generation.
12. How have cheaper, faster chips impacted the camera industry? Give an example of the leadership shifts that have occurred in this industry.
13. What has been the impact of "faster, cheaper" on Apple's business lines?
14. How did Amazon utilize the steep decline in magnetic storage costs to its advantage?
15. How does Moore's Law impact production and inventory decisions?

5.2 The Death of Moore's Law?

LEARNING OBJECTIVES

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.

Moore simply observed that we're getting better over time at squeezing more stuff into tinier spaces. Moore's Law is possible because the distance between the pathways inside silicon chips gets smaller with each successive generation. While chip plants (semiconductor fabrication facilities, or **fabs**¹¹) are incredibly expensive to build, each new generation of fabs can crank out more chips per **silicon wafer**¹². 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 Moore's Law's advance. 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, many 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 meltdown. A trip through the Facebook data center during its recent rise would show that the firm was a "hot" start-up in more ways than one. The firm's servers ran so hot that the Plexiglas sides of the firm's server racks were warped and melting! E. McGirt, "Hacker, Dropout, C.E.O.," *Fast Company*, May 2007. The need to cool modern data centers draws a lot of power and that costs a lot of money.

The chief eco officer at Sun Microsystems has claimed that computers draw 4 to 5 percent of the world's power. Google's chief technology officer has said that the firm spends more to power its servers than the cost of the servers themselves. D. Kirkpatrick, "The Greenest Computer Company under the Sun," April 13, 2007. Apple, Facebook, 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

11. Semiconductor fabrication facilities; the multibillion dollar plants used to manufacture semiconductors.

12. A thin, circular slice of material used to create semiconductor devices. Hundreds of chips may be etched on a single wafer, where they are eventually cut out for individual packaging.

in The Dalles, Oregon, is charged a cost per kilowatt hour of two cents by the local power provider, less than one-fifth of the eleven-cent rate the firm pays in Silicon Valley. S. Mehta, "Behold the Server Farm," *Fortune*, August 1, 2006. Also see [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#) in this book. This difference 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!

Buying Time

One way to overcome this problem is with **multicore microprocessors**¹³, made by 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 three hundred-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." A. Ashton, "More Life for Moore's Law," *BusinessWeek*, June 20, 2005. For many applications, the multicore chips will outperform a single speedy chip, while running cooler and drawing less power. Multicore processors are now mainstream.

Today, most smartphones, PCs, and laptops sold have 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. Intel has even demonstrated chips with upwards of fifty cores.

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. Multicore operating systems can help achieve some performance gains. Versions of Windows or the Mac OS that are aware of multicore processors can assign one program to run on one core, while a second application is assigned to the next core. But in order to take full advantage of multicore chips, applications need to be rewritten to split up tasks so that smaller portions of a problem are executed simultaneously inside each core.

13. Microprocessors with two or more (typically lower power) calculating processor cores on the same piece of silicon.

Writing code for this “divide and conquer” approach is not trivial. In fact, developing software for multicore 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.” A. Ashton, “More Life for Moore's Law,” *BusinessWeek*, June 20, 2005. Microsoft's chief research and strategy officer has called coding for these chips “the most conceptually different [change] in the history of modern computing.” M. Copeland, “A Chip Too Far?” *Fortune*, September 1, 2008. Despite this challenge, some of the most aggressive adaptors of multicore 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 multiplayer games.

Another approach that's breathing more life into Moore's Law moves chips from being paper-flat devices to built-up 3-D affairs. By building up as well as out, firms are radically boosting speed and efficiency of chips. Intel has flipped upward the basic component of chips—the transistor. Transistors are the supertiny on-off switches in a chip that work collectively to calculate or store things in memory (a high-end microprocessor might include over two billion transistors). While you won't notice that chips are much thicker, Intel says that on the miniscule scale of modern chip manufacturing, the new designs will be 37 percent faster and half as power hungry as conventional chips. K. Bourzac, “How Three-Dimensional Transistors Went from Lab to Fab,” *Technology Review*, May 6, 2011.

New Materials & Quantum Leaps? Thinking Beyond Moore's Law—Constraining Silicon

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. Multicore and 3-D transistors 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 30 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. New processors made with hafnium-based high-k components can pack circuits tighter, leading to faster chips that require less wattage than traditional silicon alone—and even more efficient materials may be on the horizon. Y. L. Chen, J. G. Analytis, J.-H. Chu, Z. K. Liu, S.-K. Mo, X. L. Qi, H. J. Zhang, et al., "Experimental Realization of a Three-Dimensional Topological Insulator, Bi_2Te_3 ," *Science* 325, no. 5937 (July 10, 2009): 178–81; K. Greene, "Intel Looks Beyond Silicon," *Technology Review*, December 11, 2007; and I. Thomson, "AMD: New Trinity Laptop Chips Out-Juice Intel Graphics," *The Register*, May 15, 2012. Hyperefficient 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^{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. P. Kaihla, "Quantum Leap," *Business 2.0*, August 1, 2004.

Some speculate that quantum computers could one day allow pharmaceutical companies to create hyperdetailed 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 (in your sunglasses, for example) could recognize the faces of anyone you've met and give you a heads-up to their name and background. P. Schwartz, C. Taylor, and R. Koselka, "The Future of Computing: Quantum Leap," *Fortune*, August 2, 2006. 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").

Pioneers in quantum computing include IBM, HP, NEC, and a Canadian start-up named D-Wave. 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.

KEY TAKEAWAYS

- As chips get smaller and more powerful, they get hotter and present power-management challenges. And at some point Moore's Law will stop because we will no longer be able to shrink the spaces between components on a chip.
- Multicore chips use two or more low-power calculating "cores" to work together in unison, but to take optimal advantage of multicore chips, software must be rewritten to "divide" a task among multiple cores.
- 3-D transistors are also helping extend Moore's Law by producing chips that require less power and run faster.
- New materials may extend the life of Moore's Law, allowing chips to get smaller, still. Entirely new methods for calculating, such as quantum computing, may also dramatically increase computing capabilities far beyond what is available today.

QUESTIONS AND EXERCISES

1. What three interrelated forces threaten to slow the advancement of Moore's Law?
2. Which commercial solutions, described in the section above, are currently being used to counteract the forces mentioned above? How do these solutions work? What are the limitations of each?
3. Will multicore chips run software designed for single-core processors?
4. As chips grow smaller they generate increasing amounts of heat that needs to be dissipated. Why is keeping systems cool such a challenge? What are the implications for a firm like Yahoo! or Google? For a firm like Apple or Dell?
5. What are some of the materials that may replace the silicon that current chips are made of?
6. What kinds of problems might be solved if the promise of quantum computing is achieved? How might individuals and organizations leverage quantum computing? What sorts of challenges could arise from the widespread availability of such powerful computing technology?

5.3 Bringing Brains Together: Supercomputing, Grid Computing, and Putting Smarts in the Cloud

LEARNING OBJECTIVES

1. Give examples of the business use of supercomputing, grid computing, and cloud 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 parallel processing found in modern supercomputing and grid computing efforts as well as multi-core processors.

As Moore's Law makes possible the once impossible, businesses have begun to demand access to the world's most powerful computing technology.

Supercomputers¹⁴ are computers that are among the fastest of any in the world at the time of their introduction. A list of the current supercomputer performance champs can be found at <http://www.top500.org>. Supercomputing was once 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, the machine that rather controversially beat chess champion Garry 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.

14. Computers that are among the fastest of any in the world at the time of their introduction.

Paging Doctor Watson

In spring 2011, the world was introduced to another IBM supercomputer—the Jeopardy-playing Watson. Built to quickly answer questions posed in natural language, by the end of a televised three-day tournament Watson had put the hurt on prior Jeopardy champs Ken Jennings and Brad Rutter, trouncing the human rivals and winning one million dollars (donated to a children's charity). Watson's accomplishment represented a four-year project that involved some twenty-five people across eight IBM research labs, creating algorithms, in a system with ninety servers, "many, many" processors, terabytes of storage, and "tens of millions of dollars" in investment. L. Sumagaysay, "After Man vs. Machine on 'Jeopardy,' What's Next for IBM's Watson?" *Good Morning Silicon Valley*, February 17, 2011. Winning Jeopardy makes for a few nights of interesting TV, but what else can it do? Well, the "Deep QA" technology behind Watson might end up in your doctor's office, and docs could likely use that kind of "exobrain." On average "primary care physicians spend less than twenty minutes face-to-face with each patient per visit, and average little more than an hour each week reading medical journals." C. Nickisch, "IBM to Roll Out Watson, M.D.," *WBUR*, February 18, 2011. Now imagine a physician assistant Watson that could leverage massive diagnosis databases while scanning hundreds of pages in a person's medical history, surfacing a best guess at what docs should be paying attention to. A *JAMA* study suggested that medical errors may be the third leading cause of death in the United States, B. Starfield, "Is U.S. Health Really the Best in the World?" *Journal of the American Medical Association*, July 26, 2000. so there's apparently an enormous and mighty troubling opportunity in health care alone. IBM is partnering with Massachusetts voice-rec leader Nuance Communications (the Dragon people) to bring Watson to the doc's office. Med schools at Columbia and the University of Maryland will help with the research effort. Of course, you wouldn't want to completely trust a Watson recommendation. While Watson was good enough to be tournament champ, IBM's baby missed a final Jeopardy answer of "Chicago" because it answered that Toronto was a U.S. city.



Source: <http://www.ibm.com/press>.

One of the first customers of Deep Blue technologies was United Airlines, which gained an ability to examine three hundred and fifty thousand 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. Also noteworthy: the supercomputer-enabled, risk-savvy CIBC was relatively unscathed by the subprime crisis.

Modern supercomputing is typically done via a technique called **massively parallel**¹⁵ processing (computers designed with many microprocessors that work together, simultaneously, to solve problems). The fastest of these supercomputers are built using hundreds of microprocessors, all programmed 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. The Air Force recently issued a request-for-proposal to purchase 2,200

15. Computers designed with many microprocessors that work together, simultaneously, to solve problems.

PlayStation 3 systems in hopes of crafting a supercheap, superpowerful machine using off-the-shelf parts.

Another technology, known as **grid computing**¹⁶, 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 technique 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 twenty-five thousand dollars, 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 six weeks and the firm was able to analyze only ten to fifty genes a year. But by leveraging grid computing, Monsanto has reduced gene analysis to less than a day. The fiftyfold time savings now lets the firm consider thousands of genetic combinations in a year. P. Schwartz, C. Taylor, and R. Koselka, "The Future of Computing: Quantum Leap," *Fortune*, August 2, 2006. 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, GlaxoSmithKline, 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 JPMorgan Chase grid was saving the firm \$5 million per year.

16. A type of computing that uses special software to enable several computers to work together on a common problem as if they were a massively parallel supercomputer.

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. And 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 PlayStation 3 had enlisted over half a million consoles within months of release. Having access to these free resources is an enormous advantage for researchers. Says the director 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." G. Johnson, "Supercomputing '@Home' Is Paying Off," *New York Times*, April 23, 2002.

Multicore, 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 multicore 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 (essentially a computer made up of multiple computers). While these technologies offer great promise, they're all subject to the same limitation: software must be 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, for example, step two can't be started until the results from step one are known), the multiple-brain approach doesn't offer much help.

Massive clusters of computers running software that allows them to operate as a unified service also enable new service-based computing models, such as **software as a service (SaaS)**¹⁷ and **cloud computing**¹⁸. In these models, organizations replace traditional software and hardware that they would run in-house with services that are delivered online. Google, Microsoft, Salesforce.com, and Amazon are among the firms that have sunk billions into these Moore's Law-enabled **server farms**¹⁹, creating entirely new businesses that promise to radically redraw the software and hardware landscape while bringing gargantuan computing power to the little guy. (See [Chapter 10 "Software in Flux: Partly Cloudy and Sometimes Free"](#).) Moving 'brains' to the cloud can help increase calculating performance even when we can't pack more processing brawn into our devices, but the cloud requires a long-distance connection that's a lot slower and perhaps less reliable than the quick hop from storage to processor that occurs inside most consumer electronics.

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!"

17. A form of cloud computing where a firm subscribes to a third-party software and receives a service that is delivered online.
18. Replacing computing resources—either an organization's or individual's hardware or software—with services provided over the Internet.
19. A massive network of computer servers running software to coordinate their collective use. Server farms provide the infrastructure backbone to SaaS and hardware cloud efforts, as well as many large-scale Internet services.

KEY TAKEAWAYS

- Most modern supercomputers use massive sets of microprocessors working in parallel.
- The microprocessors used in most modern supercomputers are often the same commodity chips that can be found in conventional PCs and servers.
- Moore's Law means that businesses as diverse as financial services firms, industrial manufacturers, consumer goods firms, and film studios can now afford access to supercomputers.
- Grid computing software uses existing computer hardware to work together and mimic a massively parallel supercomputer. Using existing hardware for a grid can save a firm the millions of dollars it might otherwise cost to buy a conventional supercomputer, further bringing massive computing capabilities to organizations that would otherwise never benefit from this kind of power.
- Massively parallel computing also enables the vast server farms that power online businesses like Google and Facebook, and which create new computing models, like software as a service (SaaS) and cloud computing.
- The characteristics of problems best suited for solving via multicore systems, parallel supercomputers, or grid computers are those that can be divided up so that multiple calculating components can simultaneously work on a portion of the problem. Problems that are linear—where one part must be solved before moving to the next and the next—may have difficulty benefiting from these kinds of “divide and conquer” computing. Fortunately many problems such as financial risk modeling, animation, manufacturing simulation, and gene analysis are all suited for parallel systems.

QUESTIONS AND EXERCISES

1. What is the difference between supercomputing and grid computing? How is each phenomenon empowered by Moore's Law?
2. How does grid computing change the economics of supercomputing?
3. Which businesses are using supercomputing and grid computing? Describe these uses and the advantages they offer their adopting firms. Are they a source of competitive advantage? Why or why not?
4. What are the characteristics of problems that are most easily solved using the types of parallel computing found in grids and modern day supercomputers? What are the characteristics of the sorts of problems not well suited for this type of computing?
5. Visit the SETI@Home Web site (<http://setiathome.ssl.berkeley.edu/>). What is the purpose of the SETI@Home project? How do you participate? Is there any possible danger to your computer if you choose to participate? (Read their rules and policies.)
6. Search online to identify the five fastest supercomputers currently in operation. Who sponsors these machines? What are they used for? How many processors do they have?
7. What is "Moore's Wall"?
8. What is the advantage of using grid computing to simulate an automobile crash test as opposed to actually staging a crash?

5.4 E-waste: The Dark Side of Moore's Law

LEARNING OBJECTIVES

1. Understand the magnitude of the environmental issues caused by rapidly obsolete, faster and cheaper computing.
2. Explain the limitations of approaches attempting to tackle e-waste.
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 that help the manager behave more responsibly.

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 and 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 ever-growing mountains of discarded tech junk, known as electronic waste or **e-waste**²⁰. According to the U.S. Environmental Protection Agency (EPA), in 2007 the United States alone generated over 2.5 million tons of e-waste, U.S. Environmental Protection Agency, *General Information on E-waste*, February 5, 2009. 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 bits of increasingly valuable metals such as silver, platinum, and copper. In fact, there's more gold in one pound of discarded tech equipment than in one pound of mined ore. P. Kovessy, "How to Trash Toxic Tech," *Ottawa Business Journal*, May 12, 2008. 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.

20. Discarded, often obsolete technology; also known as electronic waste.

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, C. Bodeen, "In 'E-waste' Heartland, a Toxic China," *International Herald Tribune*, November 18, 2007. so it's not surprising that up to 80 percent of the material dropped off for recycling is eventually exported. E. Royte, "E-waste@Large," *New York Times*, January 27, 2006. Much of this waste ends up in China, South Asia, or sub-Saharan Africa, 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 four hundred to six hundred times greater than what international standards deem safe. E. Grossman, "Where Computers Go to Die—and Kill," *Salon.com*, April 10, 2006, <http://www.salon.com/news/feature/2006/04/10/ewaste>. The area is so polluted that drinking water must be trucked in from eighteen miles away. Pregnancies are six times more likely to end in miscarriage, and 70 percent of the kids in the region have too much lead in their blood. *60 Minutes*, "Following the Trail of Toxic E-waste," November 9, 2008.

Figure 5.5 Photos from Guiyu, China. J. Biggs, "Guiyu, the E-waste Capital of China," *CrunchGear*, April 4, 2008.



Source: © 2006 Basel Action Network (BAN).

China cares about its environment. The nation has banned the importing of e-waste since 2000. E. Grossman, "Where Computers Go to Die—and Kill," *Salon.com*, April 10, 2006. 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 tape a one-hundred-dollar bill on the side of the container. C. Bodeen, "In 'E-waste' Heartland, a Toxic China," *International Herald Tribune*, November 18, 2007. Well-meaning U.S. 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 products purchased. BAN investigators in Lagos, Nigeria, documented mountains of e-waste with labels from schools, U.S. 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 industrialized counterparts. The reality is that 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." J. Russell, "Dell under Attack over Using Prison Labour," *Inquirer*, January 10, 2003. See also <http://laughingmeme.org/2003/03/23/dell-recycling-a-ways-to-go-still>. 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 and retailers to take back their products for responsible recycling and to eliminate the worst toxins from their offerings.

Consider that even Apple (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. And Apple eventually claimed the top in Greenpeace's "Greener Electronics" rankings. J. Dalrymple, "Apple Ranks Highest among Greenpeace's Top

Tech Companies," *The Loop*, January 7, 2010. 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.

Finding Responsible E-waste Disposers

A recent sting operation led by the U.S. Government Accountability Office (U.S. GAO) found that forty-three American recyclers were willing to sell e-waste illegally to foreign countries, without gaining EPA or foreign country approval. Appallingly, at least three of them held Earth Day electronics-recycling events. U.S. Government Accountability Office (U.S. GAO), *Report to the Chairman: Committee on Foreign Affairs, House of Representatives: Electronic Waste*, August 2008.

So how can firms and individuals choose proper disposal partners? Several certification mechanisms can help shed light on whether the partner you're dealing with is a responsible player. The Basel Action Network e-Stewards program certifies firms via a third-party audit, with compliant participants committing to eliminating e-waste export, land dumping, incineration, and toxic recycling via prison labor. The International Association of Electronics Recyclers (IAER) also offers audited electronics recycler certification. And firms certified as ISO 9001 and ISO 14001 compliant attest to quality management and environmental processes. Standards, techniques, and auditing practices are constantly in flux, so consult these organizations for the latest partner lists, guidelines, and audit practices. Basel Action Network e-Stewards program accessed via http://www.e-stewards.org/esteward_certification.html; International Standards Organization accessed via <http://www.iso.org/iso/home.htm>; the IAER accessed via <http://www.iaer.org/search>; and G. MacDonald, "Don't Recycle 'E-waste' with Haste, Activists Warn," *USA Today*, July 6, 2008.

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 multimillion-dollar green investment funds, targeted at funding the next generation of sustainable, environmental initiatives.

KEY TAKEAWAYS

- E-waste may be particularly toxic since many components contain harmful materials such as lead, cadmium, and mercury.
- Managers must consider and plan for the waste created by their products, services, and technology used by the organization. Consumers and governments are increasingly demanding that firms offer responsible methods for the disposal of their manufactured goods and the technology used in their operations.
- Managers must audit disposal and recycling partners with the same vigor as their suppliers and other corporate partners. If not, an organization's equipment may end up in environmentally harmful disposal operations.

QUESTIONS AND EXERCISES

1. What is e-waste? What is so dangerous about e-waste?
2. What sorts of materials might be harvested from e-waste recycling?
3. Many well-meaning individuals thought that recycling was the answer to the e-waste problem. But why hasn't e-waste recycling yielded the results hoped for?
4. What lessons do the challenges of e-waste offer the manager? What issues will your firm need to consider as it consumes or offers products that contain computing components?
5. Why is it difficult to recycle e-waste?
6. Why is e-waste exported abroad for recycling rather than processed domestically?
7. What part does corruption play in the recycling and disposal of e-waste?
8. What part might product design and production engineering play in the reduction of the impact of technology waste on the environment?
9. What are the possible consequences should a U.S. firm be deemed "environmentally irresponsible"?
10. Name two companies that have incurred the wrath of environmental advocates. What might these firms have done to avoid such criticism?