

CHAPTER 6

Evolutionary Trends in Cyberspace

Edward Skoudis

CYBERPOWER is built on the rapidly shifting terrain of cyberspace, which includes not only the Internet, but also the legacy telephony infrastructure, cellular phone technologies, and wireless data services. The technologies underlying all of these aspects of cyberspace—such as bandwidth, interconnectedness, processor speed, functionality, and security vulnerabilities—have evolved over decades. The purpose of this chapter is to identify these evolutionary trends and to extrapolate their implications, creating a view of possible and likely aspects of the future of cyberspace.

This chapter focuses on the accumulation of incremental evolutionary change over long periods. Even individually small quantitative changes, when compounded over time, can bring about great qualitative changes on their own. In addition to evolutionary trends, revolutions are also possible, but revolutionary discontinuities are difficult to predict, and so they are not discussed here (but see chapter 8 in this volume, “The Future of the Internet and Cyberpower”).

Trends that have long-term staying power and transformative implications will be examined, while fads and “flash-in-the-pan” issues will be ignored, even though it is not always easy to see the difference while in the midst of major transformations.¹ While fads can help set initial conditions for future evolution, they are hard to predict as they begin. However, the follow-on evolution can be observed and extrapolated into the future. This chapter attempts to identify real changes affecting cyberspace and to filter out passing fads.

With that goal in mind, three types of trends will be examined: computer and network, software, and social. Computer and network trends include:

- increases in computer and network power
- proliferation of broadband connectivity
- proliferation of wireless connectivity
- transition from Internet Protocol version 4 (IPv4) to IPv6

Software trends include:

- increases in software complexity
- enhanced capabilities for search both across local systems and Internet-wide
- widespread virtualization of operating systems
- convergence of technologies
- increased noise in most aspects of cyberspace
- increased vulnerability due to advancement of computer and network attack and exploit methodologies.

Social trends in the use and development of cyberspace include:

- worldwide technological development, with different local emphases
- rise in online communities, collaboration, and information-sharing.

While cyberspace evolution is proceeding on a multitude of fronts, these trends were identified by a team of authors associated with this book as the sources of a great deal of the momentum for the evolution of cyberspace. The trends in this list represent the evolution of the underlying elements of cyberspace itself as the basis for a variety of other technological, social, economic, and related changes.

Taken in their totality, these trends point to a future in which cyberspace becomes far more pervasive, touching most aspects of daily life in some way for a majority of the world. Two longstanding trends—significantly lower cost of processor performance and increases in flexible network connectivity—will facilitate the incorporation of cyberspace into more and more products. If these trends continue, some form of intelligence and network communication will eventually be embedded in most electrically powered devices: if you plug it into a wall socket today, it is likely to have some cyberspace functionality in the future. Indeed, with advances in wireless and battery technologies, even many objects that do not get plugged into the wall will also have cyberspace components. These trends mean that cyberspace is increasingly becoming an overlay technical infrastructure to our physical world, as it increasingly becomes involved in monitoring, analyzing, and altering the physical landscape.

Computer and Network Trends

Among the computer and network trends likely to have lasting effect are increases in power, proliferation of broadband and wireless connectivity, and upgrades in the fundamental protocols of the Internet.

Increases in Computer and Network Power

Moore's law describes a major component of the evolution of the information technology industry. Originally observed in 1965 by Gordon Moore, co-founder of Intel Corporation, Moore's law posits that industry's ability to produce integrated circuits continually improves, so that the number of microcomponents that can be etched on a chip will double at regular intervals.² There is some variation in the specific interval cited for the doubling timeframe observed by the law. Gordon Moore originally predicted doubling each year but later revised the timeframe to two years. Historically, the timeframe has varied between one and two years. Most current estimates focus on the two-year timeframe.

The doubling in the density of circuitry translates to increased processor performance and lower costs. Although the slowdown and even ultimate demise of Moore's law are predicted from time to time, the pace it describes has continued for over 40 years. It should be noted that, although Moore's observation is commonly referred to as "Moore's law," it is an observation and an industry goal, not a "law" in the physical sense.³

As individual machines grow more powerful, they are also increasingly interconnected in networks. The benefits of the increase in interconnectivity are addressed in Metcalfe's law, named after Robert Metcalfe, one of the inventors of Ethernet technology. Metcalfe posited that the value of a telecommunications network is proportional to the square of the number of its users. According to this hypothesis, as more users are brought onto a shared communications

network, the value of the network to the overall community grows not just at a linear rate, but as the square of the number of users. A related hypothesis, Reed's law, estimates the value even higher, saying that the utility of a network can scale exponentially with the number of participants in the network.

Unlike Moore's law, which has demonstrably reflected reality for the past 40 or more years, Metcalfe's and Reed's laws cannot be quantified: they are more a metaphorical statement of the value and power of networks than a quantifiable observation or prediction. Furthermore, Metcalfe's and Reed's laws have been challenged; some observers have said that they overstate the increase in value of a network when new users join.⁴ Still, it is generally agreed that the value of a network grows faster than at a linear rate of the number of users.

Well established technological and economic trends led to the observations known as Moore's law and Metcalfe's law; these trends suggest that the future of cyberspace will see faster computing devices interconnected in more powerful and valuable networks. With the compounding impact of these trends over time, cyberspace will continue to grow in importance and influence as more economic, military, and even social activity migrates to that realm.

Broadband Proliferation

Another major evolutionary trend in cyberspace is the widespread deployment of broadband Internet services. Cyberspace experienced at a speed of just 56 kilobytes per second (kbps)—allowing rudimentary Web surfing and data exchange—is very different from cyberspace at 400 kbps or more. Widespread access to faster connectivity by desktops, laptops, personal digital assistants, and cell phones enables new business models and new social interactions. Just as business and social models were transformed by the move from trains and ocean liners to jet airline flights in the past century, today the richer audio, video, and networked applications supported by higher speed Internet connections make cyberspace significantly more valuable to its users. Many face-to-face transactions and physical exchanges can be supplanted by much less expensive software-based interactions. New services not previously associated with the Internet, including telephony and television, are moving to this plentiful, cheap bandwidth. Widespread fiber optics, cable modems, digital subscriber loops/lines (DSL), and broadband wireless services are the technological underpinnings allowing broadband access by consumers or businesses throughout the world.⁵ These technologies have higher speeds (typically greater than 400 kbps) and are also always on, not requiring time-consuming dial-up “handshakes.” While broadband deployment is already a reality in many parts of the world, it is not yet universal. Some set of users probably will continue to use dial-up access for some time, but there are likely to be fewer and fewer.

Broadband connectivity to the endpoints of computer communication is only possible if the Internet backbone itself can carry all of the extra traffic generated by these end systems. Internet backbone providers have deployed more high-speed links, using new fiber technologies such as OC-48 and OC-192 (operating at speeds of 2.488 Gigabits per second [Gbps] and 10 Gbps, respectively).⁶ Higher speed satellite and microwave towers are interconnecting high-speed networks around the world. Interconnectivity between various backbone providers has increased to carry more traffic more quickly. As this interconnectivity grows, the topology of the Internet backbone becomes more complex from a design and

management perspective.

As client computers increasingly rely on broadband access, and as the Internet infrastructure is refined to carry all of those bits, the servers to which the communications are directed likewise need additional computational and network power to provide new services. Faster servers can help deal with some of these needs, but most large Internet companies are instead relying on larger numbers of less expensive servers, distributed across one or more campuses around the world. Large Internet companies, among them Google, Amazon, eBay, and Microsoft, are constructing vast “server farms.” By some estimates, Google’s server count is at over half a million computers in 2007 and rising quickly.

Business models are also evolving, as Internet service providers (ISPs) try to furnish value-added services and applications on top of their increasingly commoditized bandwidth business. To realize the business value in the telephony, video, and various other applications and content being distributed via their “pipes,” some ISPs are partnering with, buying, or building in-house application services and content. Such services and content are directly affiliated with that ISP, in contrast to other services that are not affiliated but that are accessed through that ISP by its customers. This evolution has led some ISPs to consider charging nonaffiliated application service providers and users a premium for their use of high bandwidth. Those that do not pay extra may face lower performance, with their traffic handled at a lower priority than higher paying affiliated users and application providers. This economic friction between ISPs and application service providers is a worldwide phenomenon and has been termed the “Net neutrality” issue. Some argue that governments should require the equal or “neutral” handling of affiliate and nonaffiliate traffic alike by ISPs, to foster interoperability and prevent the fragmentation of the Internet into various ISP enclaves. Others argue that allowing economic advantages for an ISP’s affiliated services will encourage investment in new and improved services by the ISPs.

Several countries are grappling with this complex and contentious issue. Net neutrality issues are being studied by Japan’s Ministry of Internal Affairs and Communications to determine a reasonable balance among the competing factors. The European Union has tentatively supported neutral networks, but companies such as Deutsche Telekom and Telecom Italia are beginning to lobby for changes to the existing European approach. A Net neutrality bill being debated in the U.S. Congress would require ISPs to handle traffic independently of their business relationships. Such legislation is hotly contested, and it is not yet clear how the issue will evolve.

In the future, bandwidth to the end system and on the backbone will likely grow even faster, with more widespread deployment of 10 megabits per second (Mbps) or higher rates to the home and a corresponding backbone to support it. Video applications will almost certainly increase, as today’s nascent Internet television business grows much larger and video conferencing is more widely used.

Wireless Proliferation

Cyberspace connectivity is increasingly moving to wireless communication, in the form of wireless local area networks (WLANs), wireless broadband, and similar technologies. One of the major vectors of this move is often overlooked in discussions of cyberspace: the cell phone. An estimated two billion people have access to cell phones. Over 150 million camera

phones have been sold to date, each supporting voice, text, and image transfer. Even with low bandwidth text messaging, their small size, decentralized communications capacity, and relatively low cost have made cell phones an increasingly important technology underlying social change. Rioters in France in late 2005 and early 2006, people involved in Ukraine's Orange Revolution in the winter of 2004/2005, and terrorist organizations have all relied on cheap and ubiquitous cell phone text-messaging to exercise command and control and to disseminate information.

Where today's cell phones have simple text messaging and still cameras, future cell phones with higher bandwidth applications and full video capabilities are likely to have even greater impact as they gain the capabilities of modern personal computers (PCs) and become, in effect, pocket-sized television studios. The hand-held video cameras of the late 1980s and early 1990s led to countless major news stories, such as the 1992 Rodney King riots in Los Angeles. However, that technology was limited: few people carried cameras with them most of the time, and the technology required the cumbersome handling of physical videotapes, typically delivered by hand or by courier. By contrast, when a major portion of the world's population carries cell phone-based video cameras wirelessly linked to the Internet at all times, cell phones will likely have a much larger impact on society, as when camera-equipped cell phones disseminated graphic pictures and video, albeit of low quality, of the execution of Saddam Hussein in 2006 within minutes of the event.

Today's grainy cell phone photos and simple videos will be replaced by much better multi-megapixel video cameras integrated into cell phones as video capture technology becomes smaller, cheaper, and more widely disseminated. Already television news outlets in some countries ask members of the public to submit cell phone video of events immediately after they happen. Web sites offer to broker the sale of videos of hot news events taken by private individuals. These trends decrease the time between an event's occurrence and its coverage in the news, further shrinking the news reporting cycle and the time available for the public and policymakers to analyze events.

Numerous other wireless technologies are transforming the nature of cyberspace. WLANs, in their most popular form, are implemented according to a set of standards denominated "802.11." Such "WiFi" networks allow nearby systems (within perhaps 100 meters) to communicate with one another or gain access to the Internet. Untethering computers from wired access makes it possible to use them for a wider variety of applications, ranging from industrial use to home appliances. By minimizing the need for costly wiring installations, WLANs allow for more rapid network deployment at much lower costs. A variety of organizations have taken advantage of this: stores in shopping malls use wireless for point-of-sales terminals and inventory control, and military deployments can rapidly deploy computer networks. WLAN technology is pushing the boundaries of the definition of "local" as well: originally designed for shorter distances, 802.11 signals have been successfully carried several miles and have been detected at over 200 miles under ideal circumstances (atop mountains on a clear day). Wireless signal transmission for computer devices often outpaces the distances it was designed to span. While WLAN technologies were created for distances up to 100 meters, their propagation across a mile or more has significant implications. Most consumers would be surprised to hear that the WLANs they have deployed in their homes can be detected and even accessed many blocks or miles away. Wireless data communications opens up computer

network access over longer distances, making computers more accessible to both their users and would-be attackers.

With widespread deployment of WLANs, numerous wireless networks often occupy overlapping physical spaces. Activating a wireless detection device in any major city typically reveals at least a half-dozen nearby WLANs ready for a connection. Systems may appear on a WLAN for a short time, use it to transmit some vital information, and quickly disappear; this makes it hard to determine which assets are part of a given organization's network and which are not. For enterprises, maintaining a list of computing assets in such environments is difficult. And if an enterprise does not know which machines are part of its network and which are not, managing and securing those assets becomes impossible.

Another rapidly rising wireless technology is evolution data optimized (EVDO) service, by which cellular networks make possible high-speed data transmission from PCs and cell phones. EVDO is an evolutionary descendant of code division multiple access (CDMA) technology used by some cell phones for wireless data transmission. In the United States, a handful of carriers have deployed networks that use EVDO and CDMA technology, allowing business users and consumers in most major cities wireless connectivity at 128 kbps to two Mbps over distances of a mile or more. With a simple PC card, cell phone and PC users can gain wireless broadband access to the Internet from most major population centers in the United States. Such services could supplant cable modems and DSL, allowing more mobility at high bandwidths over longer distances.

Another emerging wireless technology is Worldwide Interoperability for Microwave Access (WiMAX), designed to obviate expensive copper and fiber solutions for the "last mile" to consumers' homes. Although it has a maximum throughput of 70 Mbps and a maximum distance of 70 miles, WiMAX in real-world circumstances achieves approximately 10 Mbps over about 2 miles. WiMAX deployment is beginning to be used in urban and suburban areas to link WiFi LANs and to connect users to their ISPs.

Other wireless technologies, such as Bluetooth, are interconnecting the components of an individual computer wirelessly over distances of up to 10 meters. Bluetooth capabilities are built into many modern laptops, keyboards, computer mouse devices, cell phones, headsets, and music players. While designed for distances of only 10 meters, hobbyists have discovered that Bluetooth signals can sometimes be detected over a mile away. With numerous Bluetooth-enabled devices in close proximity, it can be hard to determine which assets are part of a given computer and which belong to another, again complicating the management and security of assets that are increasingly ephemerally tied to the network via wireless.

Another rapidly rising wireless technology is Radio Frequency Identifier (RFID) tags, very small, simple computer chips that send a unique identifier number. They are being used to support inventory activities, augmenting and possibly someday supplanting the familiar universal product code barcodes found on nearly all products today. With RFID tags, which are about the size of a grain of rice, a given product's code can be read without line-of-site viewing or direct physical access, as long as the radio frequency transmission can be read. RFIDs were designed to communicate over short distances (originally one to ten meters, but hobbyists have demonstrated reading such tags over 100 meters away). Codes identifying equipment can be read without physical contact over such distances, with possible privacy implications as RFID applications spread.

Several organizations have expressed interest in using RFID technology for large-scale inventory management, including Wal-Mart Corporation, the U.S. military, and the Chinese government. The U.S. State Department has begun using RFID tags in electronic passports. As RFID deployment becomes more prominent, implementation vulnerabilities are likely to be discovered and scrutinized. Researchers have begun to devise methods for attacking RFID infrastructures, devising hypothetical worms that could spread from tag to tag, infecting large numbers of systems with very simple code. Research on attacks against RFID readers, including the transmission of malicious code to such readers, is also under way. RFID spoofing, whereby an attacker makes a bogus RFID tag impersonate another legitimate tag, is an active area of research today with implications on cloning passports. Skimming is the process of surreptitiously reading an RFID tag to extract its vital information, which may later be used in a spoofing attack to clone the tag. With RFID information embedded into consumer products, sensors deployed in a city could instantly determine the products carried by citizens who walk within 100 meters of the sensors, allowing monitors to determine the make and model of various devices carried by the user—an issue with significant privacy implications. Very invasive remote search of pedestrians or houses by government and law enforcement officials (as well as thieves) becomes possible with the technology.

Transition from IPv4 to IPv6

The current Internet infrastructure is based on the widely deployed IPv4, a specification originally created in the late 1970s that spread widely in the early 1980s and throughout the 1990s as the Internet grew. This protocol far exceeded its original expectations, becoming the common language for communication across the Internet and large numbers of private networks, and allowing a huge variety of devices—from mainframe systems to cell phones—to communicate. Despite its unprecedented success as a protocol, the original IPv4 design had significant drawbacks: a limited number of network addresses that were distributed inefficiently, no built-in support for security, a lack of quality-of-service features, and limited support for mobile devices. To address these shortcomings, the Internet Engineering Task Force set out in the mid-1990s to define a next-generation Internet protocol, termed IPv6.⁷

While the IPv6 specifications were completed some time ago, full deployment has been slow. Most modern operating systems have IPv6 software, but few use it. Pockets of IPv6 networks exist in specialized laboratory and educational environments. Small IPv6 networks have been overlaid on the existing IPv4 Internet, with network nodes run by academics, researchers, and hobbyists around the world.

One of the major reasons IPv6 deployment has moved slowly involves the innovative retrofitting of its various concepts into the existing IPv4. For example, the Internet Protocol Security (IPsec) specification was designed to be mandatory in IPv6. However, the cryptographic protections supported by IPsec—such as confidentiality protection, integrity checks, and system authentication—have also been included in many IPv4 implementations at least since 1997, reducing the incentive to move to IPv6. The limited 32-bit network address space of IPv4 was to be supplanted by the vast 128-bit address space of IPv6 to allow for more systems on the global Internet, increasing the number of IP addresses from about 4 billion (4×10^9) to

3.4×10^{38} . However, most organizations have deployed various network address translation devices to shuffle and reuse private network addresses, somewhat alleviating the original constraints of IPv4's 32-bit addresses. Likewise, quality-of-service and mobility options have been implemented in IPv4. These adaptations to IPv4's limits have eased many of the "pain points" driving the demand for IPv6.

Although IPv6 deployment has started slowly, it is expected to ramp up; both the Chinese government and the U.S. military have announced intentions to move to IPv6 by 2012 to support the modernization of their large networks. Even so, some Internet experts have viewed IPv6 deployment as a perpetual "five years in the future"—always predicted, but never actually occurring. However, over the next decade, IPv6 deployment seems very likely, given the momentum of decisions by large buyers, large vendors (including Cisco and Microsoft, whose products support the protocol), and the Internet Engineering Task Force, which crafts the specifications for the protocols used on the Internet.

What are the implications of true widespread IPv6 deployment? Although IPsec is mandatory in IPv6, that does not necessarily mean that the newer protocol will immediately boost security. To speed and simplify deployment, users sometimes implement IPv6 with IPsec without the necessary trusted encryption keys, in effect blindly trusting any system on the network. Some IPv6 implementations use blank ciphers, leaving data unencrypted. Such deployments nullify any authentication and confidentiality benefits of IPsec within IPv6. Even with the careful use of trustworthy keys and ciphers, systems supporting IPv6 may still have a large number of security flaws, at least initially, in their protocol stacks. These could allow for remote denial-of-service attacks that cause a system crash or that exhaust all processing or memory resources, or they could permit system compromise and control by an attacker. Building and maintaining IP stacks is very difficult, even using the far simpler IPv4 protocol. The software development community has required a full 20 years to scrub similar problems due to faulty code out of their IPv4 implementations.⁸ IPv6 software is likely to go through a similar process as vulnerabilities are discovered and fixed. While it may not take another 20 years to get IPv6 right, it will certainly require significant effort to discern flaws in the numerous implementations of this vastly more complex protocol. The Internet and its users may be exposed to attacks for some time.

IPv6 also raises other security implications. The very large address space can make it easier for systems to hide: an attacker who modulates a network address across a very large address space can hide systems more easily than within the smaller and simpler IPv4 landscape.

Increases in Software Complexity

Although underlying hardware and network speeds have increased, most new computers do not seem to their users to be significantly faster than their predecessors for very long after their introduction. This is largely due to increases in software complexity, as additional features, increased error-handling capabilities, and more complex security facilities sap the processing gains reflected in Moore's law and in higher bandwidth. This phenomenon is sometimes referred to as Wirth's law, named after Niklaus Wirth, a Swiss computer scientist and inventor of Pascal and several other programming languages. Wirth's law states that software is decelerating faster than hardware is accelerating.

Modern software includes a proliferation of features, some important and useful to large

numbers of users, and others providing utility to only a small fraction of the user base. Users demand that new programs do more than their old software, and vendors cater to this expectation. Software vendors introduce these features to entice new customers to purchase their products, as well as to inspire existing customers to continue on the treadmill of constant software upgrades. Unfortunately, some software vendors do not spend the resources necessary for thorough development, integration, and testing of these features and modifications.

More complex software is more likely to have flaws, which may manifest themselves in broken features, software crashes, or security vulnerabilities. When a software flaw is discovered, especially one with major security implications, the software vendor typically releases a “patch” to alleviate the condition. Microsoft, for example, releases patches once per month, each typically including between five and a dozen major fixes, often requiring upwards of 10 megabytes of new code. With such massive changes pushed to over 100 million systems, the patching process for Windows alone is a monumental worldwide undertaking on a monthly basis, involving not just Microsoft, but also hundreds of thousands of enterprise users and consumers, not all of whom test such patches carefully before installing. Moreover, multiple patches applied over time could introduce additional flaws. The constant accumulation of patches can make systems more “brittle,” requiring even more complexity to patch adequately without breaking functionality.

Unfortunately, complexity is often the enemy of security, as subtle vulnerabilities linger in highly complex, perhaps poorly understood code. To address such vulnerabilities, security tools—antivirus tools, antispymware software, and intrusion prevention systems—are common defenses for systems today. Many of these tools operate in a reactive fashion, to clean up after an infection has occurred, and most defend against specific vulnerabilities that have already been found, not against as-yet-undiscovered security flaws. Compounding the problem, these security tools themselves often have flaws, so they need patches as well. This, again, increases the overall complexity of computer systems and makes them even more brittle.

Antivirus, antispymware, and other anti-malicious code technologies use a mixture of techniques to detect such “malware,” analyzing protected computers on which they are installed at a granular level to police the system for infection and attacks. The need for such defenses is turning into a significant security tax on the increases described by Moore’s law and is boosting complexity.

Enhanced Search Capabilities

With increasingly complex software used for a greater number of applications, more and more vital data is accumulating in databases and file systems. On a local system, data are typically stored in multi-Gigabyte or even Terabyte (1,000 Gigabyte) file systems. On large servers or even networked groups of systems, databases often exceed 100 Terabytes. These data are only useful if users and applications can search for information; high-quality search functionality is therefore vital, and the data must be organized, stored, and presented in useful structures. The metadata that describe the data, tagging, and visualization technologies are thus increasingly critical.

Because of their involvement with these crucial functions, Internet search engines are

currently at the center of activity in cyberspace evolution. As search engines acquire more data sources, including phone books, highly detailed satellite imagery, and maps, users are presented with more powerful search directives and operators.⁹ Such options let users hone in on specific items they seek, using a complex array of search directives instead of merely grabbing data with a specific set of search terms located in it. In addition, simple text-based searches are expanding to searches for images, sound, or videos.

With so much information on the Internet, many users need help to find what they need. Users might not even know precisely what to search for and would benefit from a ranking system of interesting or useful sources of information. To address this need, other sites on the Internet act as aggregating front-end portals that organize data from multiple data sources and process it to provide users with extra value. Sites such as digg.com and del.icio.us contain lists of popular articles and sites that are voted on by users, giving other users a guide to information in a variety of categories.

The need for search capabilities is not limited to the Internet. Internal network searching is increasingly important for locating useful information from an organization's internal servers and desktop computers. To address this need, Google offers an appliance that explores an internal enterprise network and creates a "mini-Google" for that organization, searchable in the same manner as the Internet-wide Google itself. Many other players are also moving into the internal enterprise network search market.

Local system search tools are also being deployed, including Google's Desktop Search software, Apple's Spotlight for Macintosh, and Microsoft's enhanced search capabilities for Windows machines. These tools let users formulate queries to find important information stored on their local hard drives rapidly. Current search capabilities of local products have only a limited syntax of search directives and operators, but these technologies will improve.

Search capabilities are particularly vital to analysis of very large centralized data repositories. Services such as LexisNexis (for searches in published news sources and legal documents), credit reporting agencies, and fraud detection tools for financial services organizations require extremely rapid search of large databases maintained by a small number of companies and government agencies. These data sources are used for data mining, correlation, and detailed analysis to discern trends and important outliers. Many of the operators of such databases sell search services to their clients; major economic decisions may be based on the results of searches of these data repositories. System downtime of a credit-reporting database, for example, could have major economic impact.

Novel search strategies made Google what it is today; its Page Rank algorithm for associating pages together provides a fast and reliable method for searches. In the future, as data sources and the amount of information stored grow, searching and prioritizing information are likely to become even more important, and new search strategies and new companies will arise to help people use data.

Widespread Operating System Virtualization

Virtual machine environments (VMEs) such as VMware, Microsoft's Virtual Server, and Xen let a user or administrator run one or more guest operating systems on top of a single host operating system. With such VME tools, for example, three or four instances of the

Microsoft Windows operating system can run as guest systems on top of a single Linux host operating system on a single PC or server. The concepts of virtualization were pioneered in the mainframe world but are now migrating to standard PCs and even to cell phone systems. Such virtualized environments are used for clients as well as servers in a variety of commercial, government, and military organizations, and their deployment is increasing very rapidly for several reasons. First, VMEs improve server operations by helping to cut hardware costs, simplify maintenance, and improve reliability by consolidating multiple servers onto a single hardware platform. Second, they may lower the cost of providing user access to multiple networks having different sensitivity levels. By means of VMEs, some government and military agencies and departments may use one PC, with different guest operating systems associated with each separate network a user may need.

Third, numerous “honeypot” defensive technologies rely on VMEs because they can be more easily monitored and reset after a compromise occurs. Honeypots are used to detect attacks, research attackers’ motives and methods, and provide a limited environment, isolated from critical facilities, in which to engage attackers. Given VMEs’ ability to reset an infected system quickly, most malicious code researchers utilize them to analyze the capabilities of the latest malware and to construct defenses. If a malware specimen under analysis infects and damages a guest virtual machine, the VME lets a researcher revert to the last good virtual machine image, quickly and easily removing all effects of the malware without having to reinstall the operating system.

Finally, systems that are directly accessible from the Internet have a high risk of compromise; in multi-tiered e-commerce environments, it can be expected that the front-end system will be compromised. Increasingly, therefore, these exposed hosts are installed on VMEs to minimize downtime, increase security, and simplify forensic procedures.

Computer attackers are, accordingly, becoming very interested in detecting the presence of VMEs, both locally on a potential VME and across the network. If malicious code (such technologies as spyware or keystroke loggers) detects a VME, it might shut off some of its functionality to keep researchers from observing it and devising defenses. Researchers might not notice its deeper and more insidious functionality or may have to work harder to determine what the code would do when not in the presence of a VME. Either way, the attacker buys time, and with it, additional profit. VME detection is useful to attackers who seek to avoid wasting time on honeypots. Attackers also have other motivations for discovering whether a given system is running in a VME. If, for example, an attacker could find out that a group of five systems were guest virtual machines all on a single host, launching a denial-of-service attack against the host machine would be a far easier way to cause more harm to the target organization.

As VMEs are deployed more widely, even perhaps to a majority of machines on the Internet, their detection may become a less significant issue, as attackers may come to assume they are always in a guest machine. However, other security implications of VMEs would come to the forefront. VME detection could become a precursor to VME escape, whereby an attacker might leak classified information from a more sensitive guest machine to a more easily compromised guest, undermining isolation and exposing sensitive data. An attacker or malicious code that detects a VME might try to move from a guest machine to the host machine or to other guests, compromising security, infecting other guest systems, or breaking into higher levels of security classification.

Technological Convergence

Digital convergence, which has been predicted at least since the late 1970s, is starting to happen rapidly.¹⁰ Previously disparate technologies, such as telephones, radio, television, and desktop computers, increasingly use a common set of underlying technologies and communications networks. Convergence has recently been manifested with the migration of PC-based technology—hardware such as processors and hard drives as well as software such as operating systems and browsers—to non-PC products. Internet telephony, for example, is growing rapidly in companies, government agencies, and households, and new phone companies are being formed to provide related services. Many radio stations stream their audio across the Internet to augment their broadcast service and reach new audiences around the world. Music sales on the Internet have been a small share of music sales (six percent in 2007), but the percentage is increasing rapidly.¹¹ Some music playback and video equipment incorporates hard drives to store digitized music files or video footage. Apple, Google, and others now sell or distribute television shows on the Internet, both for streaming (ABC's most popular shows, for example) and for download (YouTube, Google, and others). YouTube announced that it handled over 100 million video downloads per day as of mid-2006.¹² These are typically homemade videos or captured broadcast or cable television snippets ranging from two to five minutes in length. Through services such as Google Earth and other satellite and mapping services, cyberspace is influencing perception and use of the physical world. High-quality, easily available maps allow people to understand their surroundings and even to manipulate them better.

Convergence can help to lower transmission costs for content because a single network—the Internet—can deliver disparate services. Consumers can move content among different types of systems—from a portable music player to a computer to a television, for example. However, from a security and reliability perspective, convergence brings new risks. An attack or an accidental disruption can have a more significant impact because a larger population of users and organizations is relying on a common infrastructure using common technologies. Such disruptions may affect services that users may not realize are related. For example, if an enterprise's Internet connection goes down, most users will expect to be unable to get email, but in some enterprises, critical business functionality might also become inaccessible if it depends on Web applications residing on third-party servers on the Internet. Major upgrades of converged infrastructure may be more complex because their implications for multiple services must be considered but perhaps cannot all even be anticipated. Most organizations, and the ISPs they rely on, strive not to have potential single points of failure in their physical deployments, but their use of common operating system software (typically Windows, variations of UNIX and Linux, or Cisco's Internetwork Operating System) may raise the risk of downtime and attack.

As significant services such as telephony, television, and business transactions move to broadband Internet connectivity, modern economies increasingly depend on the Internet infrastructure and on the ISPs. In many countries, competing ISPs have taken on a role filled by monopolistic telephone companies in the past as stewards of the nation's communications infrastructure. In the United States, the older system was dominated by the Regional Bell

Operating Companies and their parent company, AT&T, which provided this stewardship under significant Federal, state, and even local regulatory control that constrained the rates they could charge for service and set standards for service reliability. In most countries, such regulations do not exist for the collection of ISPs and Internet backbone providers, who have nonetheless provided the United States at least with relatively high reliability.

There are even bigger implications as the networks used to manage various critical infrastructures converge with the Internet itself. The supervisory control and data acquisition (SCADA) systems associated with control of electrical power distribution, water distribution, pipelines, and manufacturing are increasingly managed using PC- and Internet-related technologies. Although most SCADA systems are not directly connected to the Internet itself, they are managed via maintenance ports that use the same technologies as the Internet, and even isolated SCADA systems sometimes communicate with laptop PCs that also connect to the Internet from time to time. These maintenance ports could offer a backdoor avenue for attack, exposing SCADA systems. Once convergence occurs, attacks, infections, and disruptions launched from other aspects of cyberspace could have amplified economic effects. Convergence of technologies also implies convergence of threats and vulnerabilities that thus amplify risk, perhaps in unexpected ways. Even if the SCADA systems themselves could withstand such an attack, the management systems controlling them might be disabled or impaired, affecting control of critical infrastructure facilities and thus preventing alerts and corrective actions in response to an emergency.

Another area of network convergence is the rise of voice over Internet protocol (VoIP) services. The most familiar aspect of VoIP involves end-user telephony; for example, cable companies and others advertise phone service carried over broadband Internet connections. VoIP, however, is not limited to carrying calls from individuals to their local phone companies: many long-distance companies are already transporting at least some of their long-distance traffic by means of IP-based networks to take advantage of the lower costs associated with transporting such calls over broadband pipes that mix Internet data and voice. Even users relying on traditional phone lines—the so-called Plain Old Telephone Service (POTS)—may thus have their calls carried in part over the Internet, unaware that VoIP was associated with the call made POTS-line to POTS-line. A major concern is that the existing non-VoIP long-distance network may not be able to handle all of the long-distance traffic if the Internet and its VoIP systems were impaired. Because design of local and long-distance telephony capacity now presumes that a certain call volume will be handled via VoIP, the telephony infrastructure might not be able to handle the entire load if all VoIP became unavailable due to an Internet attack or outage. This is of particular concern for emergency organizations that must rely on public facilities to communicate, including law enforcement, government, and military groups. The converse could also be a problem, if a cellular outage caused a surge in VoIP calls, overloading the carrying capacity of the Internet.

As convergence continues, many more services will be delivered to homes, government agencies, and military operations via combined networks. Even when multiple networks exist, gateways may shuttle information between one network and another, resulting in a network of interconnected networks. The resulting converged systems and networks will offer some highly useful synergies, but at the cost of increased risk.

Increased Noise in Most Aspects of Cyberspace

As cyberspace has grown, various aspects of it have become more full of *noise*, or apparently random data without meaning to most users or applications. Consider spam, or unsolicited email that typically has a commercial message: spam comprised about 30 percent of all email in 2003, has gone to over 80 percent today, and continues to rise. Newsgroups often host messages full of apparent nonsense that just takes up space. Another form of cyberspace noise is the clutter of advertisements on major Web sites today, including pop-up ads. Search engine results often include noise, either from mistaken matches returned by search engine software or Web sites that deliberately try to fool search engines in order to be included inappropriately in search results. Uniform resource locators, which point to Web pages, are increasingly cluttered with complex symbols and in some cases even small software snippets designed to run in browsers. Roving, automated Web crawlers search the Internet looking for Web sites with forms to fill out, which they then populate with advertisements or political messages. The Internet even sees raw packet noise. “Sniffing” software that monitors an unfiltered Internet connection with no regular use will see a significant amount of nonsense traffic, as users around the world inadvertently type incorrect IP addresses, attackers scan for weak target sites, and backscatter is generated by spoofed attacks against other sites. The Internet is indeed a noisy place, and it is growing noisier.

Noise is helpful to those wanting to hide: a noisy environment can let covert channels blend in, so attackers can communicate and coordinate. An attacker might send spam messages, newsgroup postings, or even raw packets to millions of targets on the Internet just to obscure delivery of a single encoded message meant for a single individual or group. Such a message is not likely to be noticed in the day-to-day noise distributed via these same mechanisms. What’s more, the location-laundering mechanisms pioneered by the spammers would make it very difficult to find the source of such a message and, given that the message is spewed to millions of destinations, locating the true intended recipient is likewise a major hurdle. Finding an information needle in the haystack of noise data is difficult, and as noise increases, it becomes harder.

In the future, the Internet will likely become even noisier, as many new services are deployed that are filled with advertising and nonsensical information. Computer attackers and political dissidents will likely increasingly use this noise to camouflage their plans and communications with each other.

Advancement of Computer and Network Attack and Exploitation Methodologies

A major trend fueling the evolution of computer attacks and exploits involves the rising profit motive associated with malicious code.¹³ Some attackers sell to the highest bidder customized malicious code to control victim machines. They may rent out armies of infected systems useful for spam delivery, phishing schemes, denial-of-service attacks, or identity theft. Spyware companies and overly aggressive advertisers buy such code to infiltrate and control victim machines. A single infected machine displaying pop-up ads, customizing search engine results, and intercepting keystrokes for financial accounts could net an attacker \$1 per month or more. A keystroke logger on an infected machine could help the attacker gather credit card

numbers and make \$1,000 or more from that victim before the fraud is discovered. With control of 10,000 machines, an attacker could set up a solid profit flow from cyber crime. Organized crime groups may assemble collectives of such attackers to create a business, giving rise to a malicious code industry. In the late 1990s, most malicious code publicly released was the work of determined hobbyists, but today, attackers have monetized their malicious code; their profit centers throw off funds that can be channeled into research and development to create more powerful malicious software and refined business models, as well as to fund other crimes.

When criminals figure out a reliable way to make money from a given kind of crime, incidents of that kind of crime inevitably rise. Computer attackers have devised various business models that are low risk, in that the attackers' chances of being apprehended are very small when they carefully cover their tracks in cyberspace. They can make hundreds of thousands or even many millions of dollars.

A factor fueling the growth of cyber attacks is bot software. Named after an abbreviated form of the word *robot*, this software allows an attacker to control a system across the Internet. A single attacker or group may set up vast botnets— groups of infected machines—scattered around the world.¹⁴ Bot-controlled machines give attackers economies of scale in launching attacks and allow them to set up virtual supercomputers that could rival the computer power of a nation- state. They can use that resource to conduct a massive flood, to crack crypto keys or passwords, or to mine for sensitive financial data used in identity theft.

Bots and other computer attack tools have become highly modular, using interchangeable software components that allow attackers to alter functionality quickly to launch new kinds of attacks. Common bots today include 50 to 100 different functional modules; an attacker could shut off or remove those modules not needed for a given attack, while more easily integrating new code features. Other modular attack tools include exploitation frameworks, which create packaged exploitation code that can infiltrate a target machine that is vulnerable (because it is misconfigured or unpatched). Just as interchangeable parts revolutionized military equipment in the early 19th century and consumer manufacturing in the early 20th century, interchangeable software components today offer computer attackers and exploiters significant advantages in flexibility and speed of evolution.

Speeding up evolution further, attackers increasingly rely on exploit and bot code that morphs itself, dynamically creating a functionally equivalent version with different sets of underlying code. Such polymorphic code helps attackers evade the signature-based detection tools used by the dominant antivirus and antispyware technology of today. This dynamically self-altering code is also harder to filter, given that it constantly modulates its underlying software. This “moving target” of code also makes analysis by defenders more difficult. Polymorphic code furthers attackers' goals because the longer the attackers have control of a botnet by evading filters and signature-based detection, the more money they can realize from the infected systems.

Another trend in attacks involves “phishing” email: an attacker pretending to be a trusted organization or individual sends email that aims to dupe a user into revealing sensitive information or installing malicious software. Attackers often spoof or mimic email from legitimate e-commerce and financial services companies to try to trick a user into surfing to a bogus Web site that appears to be a retailer or bank. When the unsuspecting user enters account information, the attacker harvests this data, using it for identity theft. More recent phishing attacks include so-called spear phishing attacks that target a particular organization or even

individuals. Such phishing email may appear to come from a trusted individual, such as a corporate executive, government manager, or military officer, and exhorts the recipient to take some action. Simply clicking on a link in a spear-phishing email could allow the attacker to exploit the victim's browser, installing a bot on that machine that would act as the attacker's agent inside of the victim enterprise.

In phone phishing, an attacker sends email with a phone number for the victim to call or even leaves POTS voice mail with a recording asking the user to call back. The calls appear to go to a major U.S. or European bank, perhaps by using the area code of 212 associated with New York City. Attackers use VoIP technology with standard voice mail software to transfer the calls outside of the United States or Europe to a voice mail system located elsewhere; a friendly recorded voice then asks the user for confidential account information. Phone phishing is automated, telephone-based criminal social engineering on a worldwide scale.

The attack and exploitation issues described in this section reinforce one another, allowing the attackers to dominate more "ground" for longer times in cyberspace, evading authorities and making money while doing so.

Social Trends

Finally, we turn to the social trends that may result from changes in cyberspace.

Worldwide Technological Development, with Different Localized Emphases

Throughout the 1980s and much of the 1990s, the locus of cyberspace evolution was the United States and Europe. There, for example, the Internet originated with the Defense Advanced Research Projects Agency, many of its standards were developed by the Internet Engineering Task Force, and the World Wide Web standards were created at CERN in Geneva. More recently, however, the trend is toward more internationalization of cyberspace deployment and technological development. A large fraction of the planet's population is now online; over one billion people had at least rudimentary access to the Internet by 2007.¹⁵ The Internet is, however, only one aspect of cyberspace: today, cell phones give more than two billion people the ability to tap into the world's telephony network. This lowers barriers to entry and allowed players from around the world to participate in cyberspace activities, for good or ill. While this trend is broad-based, various countries have carved out particular niches of their focus in cyberspace. For example, China has aggressively moved into manufacturing computers, network equipment, and telecommunications infrastructure and devices. India offers various services using the distribution media of the Internet, including call center support, software development, tax preparation, and other knowledge-based services. Europe, South Korea, and the United States have widespread broad-band access and major software development, both by commercial companies and open-source initiatives. Europe has been particularly strong in cell phone development and Japan in hand-held gadgetry such as games, digital cameras, and video players.

A typical computer, network router, or operating system involves hardware and software assembled from several countries; the true source of given components may be difficult or

impossible to track down. Thus, the involvement of more countries in the advancement of global high-tech infrastructures means that covert monitoring and control capabilities for exploitation and disruption could be added at numerous points in the supply chain outside of the United States.

These overall national areas of technological dominance are blurring with time. Some countries are making major investments in underlying bandwidth and tweaking incentives so they can become havens for high technology and new corporate development. Such activities have diminished the overall U.S. dominance of cyberspace as more and more significant contributions are made on a worldwide basis, not just by U.S.-based or European companies. Even U.S.-based companies, such as Microsoft and Google, are investing in research and development operations outside of the United States, particularly in China. Such a shift has significant intellectual property implications, as innovations devised outside of the United States by international corporations offer fewer avenues for U.S. control and increased control by other countries.

From an economic perspective, the U.S. tendency has been to rely generally on a broad free-market approach to technological development. Some other countries have relied on targeted incentives for specific technologies in an attempt to leap ahead of other players. These differing approaches have contributed to the trend of different emphases in various countries' technological development.

The trend toward internationalization of cyberspace technological change is especially visible in the realm of computer and network attacks. In the 1990s, most attacks, and indeed almost all publicly released computer attack tools (both free and commercial), came from the United States or Europe. In the early 2000s, however, computer attacks went international. Several widespread worms have been released by citizens of countries not commonly associated with high technology: in 2000, the Love Bug computer virus was released by a student in the Philippines, and in 2005, the Zotob bot was released by a developer from Morocco, funded by an individual from Turkey.¹⁶ There have been plausible allegations of Chinese probing of cyberspace, including the highly publicized Titan Rain series of attacks against U.S. military facilities.¹⁷ North Korea, not a typical bastion of computer technology, has boasted of its cyberwar hacking abilities and is rumored to run a hacking training program.¹⁸

Not just attack tools but attacks themselves have taken on a more pronounced international flavor. Ten years ago, the origin of most attacks across the public Internet was within the United States and Europe, usually a single spot. Attacks typically now come simultaneously from multiple countries, often a dozen or more. Some of these attacks are conducted by one individual in one location using bot-infected machines in other countries to mask the source of the attack, while other attacks are launched by coordinated attackers located in multiple countries. Attackers sometimes log in to U.S. systems from outside the country and use them as a base for attacks against other targets or for hosting propaganda. Some attackers motivated by geopolitical or nationalism issues launch an attack from one country against another hostile country, while others choose to launch an attack between friendly countries, hoping it will escape scrutiny. Thus, attacks sometimes come from Canada to the United States when the attackers themselves, located elsewhere, use the Canadian machines in an effort to blend in with normal traffic between the two allies.

These trends are likely to continue, as broadband access, technical training and expertise, and high-technology industries spread. For decades, students from overseas received training in high technology at U.S. academic institutions, and many took their expertise back to their home countries. More recently, world-class high-technology universities have been established in India (including the Indian Institute of Technology) and China (with its University of Science and Technology of China); thus, many students worldwide now receive high-tech training indigenously. The gap between levels of technological skill in the United States or in Europe and those in the rest of the world is shrinking and will continue to do so. Of course, the United States and Europe will keep pushing ahead to new technologies, and their existing base is an advantage, but yesterday's gaps have significantly narrowed.

Rise in Online Communities, Collaboration, and Information-sharing

Another major cyberspace trend is the rise in online communities made up of people or organizations with common interests who coordinate in cyberspace to share information and achieve other goals. As the term is most commonly used, an *online community* refers to a social setting in cyberspace, such as MySpace, LinkedIn, and Orkut, where consumers with common interests communicate and share personal profiles, business contacts, and so forth. Such sites have flourished recently; MySpace had over 300 million accounts for users around the world as of 2007.¹⁹ Such social online communities also yield information- and data-mining opportunities to law enforcement, as users provide detailed information about their lives and their network of acquaintances. Unfortunately, there have also been high-profile cases of stalkers misusing this information to target children. While today's specific most popular online communities might be a mere fad (the Friendster site saw its popularity plummet with the rise of MySpace, only to rebound later), the concept of online communities in the form of social networking sites is likely to be an enduring trend.

Another type of community is the blog, an online diary where a writer shares information and commentary about politics, hobbies, or other interests. Most bloggers allow others to provide comments on their blog, resulting in a community of sorts. The blogosphere (as all of the blogs on the Internet are collectively known) is witnessing very rapid growth: there were over 20 million blogs in January 2006.²⁰ Some blogs have become quite popular and have helped shape political debates and news stories. Blogs will likely become more consequential as the distinction between the blogosphere and traditional news media blurs. Many major newspapers have started their own blogging operations, for example, or have hired bloggers to write content for the Internet and for printed media.

Making use of social networking and blogging for a very different objective, terrorist organizations have also formulated online communities, to aid in fundraising, recruitment, propaganda, and command and control of their operations.²¹ Some of these sites are available to the public, especially those associated with propaganda, fundraising, and communiqués from terrorist leadership, while other sites, containing more sensitive information about the organization, are available only to those specifically invited.

Another use of online communities involves integrated supply chains, in which a given manufacturer relies on a host of suppliers, who in turn rely on their own suppliers, distributed in countries around the world, the whole controlled through cyberspace. With an integrated

supply chain, messages regarding inventory, capacity, and payment can be transferred quickly, allowing the manufacturer to cope more efficiently with changes in demand and possible disruptions to supply. Dell Computer famously relies on integrated supply chains using Internet technology; many other companies are also heavy users of these technologies, including United Parcel Service and Wal-Mart.²² Given the transnational nature of such supply chains, each individual country through which the chain passes has some form of cyberpower over that chain, with the ability to tax, slow down, or even shut off a vital component of the chain. However, with the economic importance of these chains, and their corporate owners' ability to use cyberspace to reroute capacity and demand or to set up new chains rapidly, most countries will probably use some restraint in their exercise of such power.

Online communities also encompass sites associated with consumer commerce, such as Amazon.com and eBay; these have created a lively interchange of buyers and sellers, with complex ranking, preference, and voting systems for products and providers. Some online communities involve participants in even deeper immersion in the cyber world. An example is Second Life, a site run by Linden Research, which describes their offering as a "3D online digital world imagined, created, and owned by its residents." Users of this community create their own avatars, or online representatives, to explore and alter the virtual reality space inside of the community and to create objects to use or buildings to inhabit. People meet, have relationships, and conduct business transactions inside of Second Life, which even has its own currency.²³ While Second Life is targeted at adults, a special area within the Second Life world is geared to teenagers. Another example, for children 5 to 10 years old, is an online world of game-playing and avatars called Webkinz created by the toy company Ganz; it, too, has its own digital economy. There are many other such communities.

Conclusion

Where are these evolutionary trends heading? The future of cyberspace is likely to involve the embedding of Internet and PC technology deeply into many everyday objects, including not only technical items such as computers, telephones, radios, and televisions, but also items not now associated with cyberspace, such as home appliances, consumer products, clothing, and more. These everyday objects will incorporate software intelligence, processing information delivered wirelessly by a broadband connection. Global positioning systems and RFID tags will allow these everyday objects to locate and interact with each other in the physical world. Cars, ships, airplanes, weapons systems, and their components will all become more intelligent and interactive, increasingly without direct human control. The economies of developed nations will rely on this interconnected grid of objects.

Unfortunately, however, the grid is built on technologies not consciously designed to handle information of such an extent and value. Security flaws will let attackers establish widespread collections of infected objects they can use to exploit other objects, manipulate target organizations, and possibly disrupt countries and economies. Such actions are not at all far-fetched, based on the accumulated evolutionary trends of the past decade. Revolutionary changes, such as those described elsewhere in this book, could have an even bigger impact on cyberspace.

¹ Two examples come to mind: the Pointcast fad of 1996 and the rise and fall of Napster in 2001. Pointcast was a much-hyped screen saver that delivered news updates across the World Wide Web. The Pointcast company eventually collapsed and delivery of news via screen savers never caught on, but the underlying trend of gathering, researching, and disseminating news via the Internet itself has grown rapidly. Similarly, Napster and free file sharing were much ballyhooed concepts in 2000, but the original Napster disappeared in a flurry of lawsuits over music copyright issues. However, both file sharing and online music sales have flourished since.

² Intel Corporation, "Moore's Law," available at <www.intel.com/technology/mooreslaw/index.htm>

³ This practice of naming certain technological industry observations or principles as "laws" is common in the information technology industry, probably modeled on Moore's law.

⁴ Bob Briscoe, Andrew Odlyzko, and Benjamin Tilly, "Metcalfe's Law Is Wrong," *IEEE Spectrum* (July 2006).

⁵ International Telecommunication Union, "Broadband Penetration by Technology, Top

20 Countries Worldwide, 2004," available at <www.itu.int/ITU-ict/statistics/at_glance/top20_broad_2004.html>.

⁶ "What Is the Speed of Standard Data Rates?" Whatis.com, available at <http://whatis.techtarget.com/definition/0,,sid9_gci214198,00.html>.

⁷ The role of the Internet Engineering Task Force in the development of Internet technology is explained in chapter 21 in this volume, "Internet Governance."

⁸ Bugs in IPv4 stacks, which allowed an attacker to crash a target system by sending a large ping packet, led to the Ping of Death attack in 1996. Similarly, the Land attack of 1997 let an attacker drive a target system's central processing unit to 100 percent across the network by sending a spoofed packet with unusual settings to the target system. This vulnerability was originally discovered in most major operating systems in 1997 and was quickly patched. Yet in 2005, the vulnerability reappeared in a patch update to Microsoft Windows, forcing Microsoft to issue yet another patch to fix this recurrent flaw.

⁹ In an example of search directives, Google's "filetype:" allows a search for specific types of files: Microsoft Excel spreadsheets, where "filetype:xls" is a search term, or MS Word documents, if a search includes filetype:doc, while "site:" limits search results to a given Web site. An operator such as "-" (NOT) filters out all Web pages with a given term; using the operator "AND" allows a search limited to results containing both of the terms on either side of the operator.

¹⁰ Rich Gordon, "Convergence Defined," *USC Annenberg Online Journalism Review*, available at <www.ojr.org/ojr/business/1068686368.php>.

¹¹ John Borland, "iTunes Outsell Traditional Music Stores," CNET News, November 2005, available at <http://news.com.com/iTunes+outsells+traditional+music+stores/2100-1027_3-5965314.html>.

¹² Marshall Kirkpatrick, "YouTube Serves 100m Videos Each Day," TechCrunch, July 2006, available at <www.techcrunch.com/2006/07/17/youtube-serves-100m-videos-each-day/>.

¹³ See chapter 18, "Cyber Crime," in this volume.

¹⁴ Bots and their associated botnets are described in detail at <<http://en.wikipedia.org/wiki/Botnet>>.

¹⁵ Miniwatts Marketing Group, "World Internet Usage and Population Stats," available at <www.internetworldstats.com/stats.htm>.

¹⁶ Robert Lemos, "Zotob Suspects Arrested in Turkey and Morocco," *Security Focus*, August 2005, available at <www.securityfocus.com/news/11297>.

¹⁷ Nathan Thornburgh, "Inside the Chinese Hack Attack," *Time*, August 2005, available at <www.time.com/time/nation/article/0,8599,1098371,00.html>.

¹⁸ Brian McWilliams, "North Korea's School for Hackers," *Wired*, June 2003, available at <www.wired.com/news/politics/0,59043-0.html>.

¹⁹ "MySpace," Wikipedia, available at <<http://en.wikipedia.org/wiki/Myspace>>.

²⁰ "Weblogs Cumulative," *Technorati*, January 2006, available at <www.sifry.com/alerts/archives/000419.html>.

²¹ United States Institute of Peace, "www.terror.net: How Modern Terrorists Use the Internet," March 2004, available at <www.usip.org/pubs/specialreports/sr116.html>. Also see chapter 19 in this volume, "Cyber Terrorism: Menace or Myth?"

²² Thomas L. Friedman, *The World is Flat: A Brief History of the 21st Century* (New York: Farrar, Straus and Giroux, 2005).

²³ "Linden Dollars" can be bought and sold in the online community using U.S. dollars.

Trends in evolution. An evolutionary trend can be either directional change within a single lineage or parallel change across lineages, in other words, several lineages undergoing the same sort of change. However, not just any change counts as a trend. After all, if the weather gets warmer one day, you wouldn't call it a warming trend; warming would have to go on for some length of time before you'd call it a trend. Biologists think about evolutionary trends in the same way — there has to be something about the change that suggests that it's not just a random fluctuation before