CHAPTER 8

The Future of the Internet and Cyberpower

Marjory S. Blumenthal and David D. Clark

THE PURPOSE of this chapter is to provide a forward-looking perspective on cyberspace that will support effective cyberpower analysis. The development and implementation of any cyberpower policy will take time. If the policy is to be relevant and not out-of-date before it is implemented, it must be based on a future view of cyberspace, not what cyberspace is today. Looking forward even 10 years, the shape of cyberspace will likely be much changed from today. There are many possible outcomes for the future of cyberspace, driven not only by the evolution of component technologies (see chapter 6 in this volume, "Evolutionary Trends in Cyberspace"), but also by the decisions and interactions of many actors whose interests may not be aligned: the research community, innovators in technology, investors in communications facilities and the services that use them, government entities at all levels, and the multitude of individual and organizational users. Indeed, the technical fact of information convergence—that all information can be digitized and carried as strings of bits—has driven industry sectors that used to be separate, such as cable and telephone providers, into vigorous competition and contention. It is the actions of these powerful players that will define the future, not the fact that bits are bits. Out of their decisions and actions will emerge new conditions for connectivity and performance, kinds of content, and contributions to cognition, as well as degrees of choice (for connectivity and content) and credibility (for content and hence cognition). Whichever of the many possible outcomes is realized will affect the level, distribution, and balance of cyberpower. We argue that a number of trends combine to point to a cyberspace that, 10 years from now, will more fundamentally blend technical, economic, and social elements in complex ways. However, we are only beginning to be able to analyze such trends, let alone predict them.

This chapter looks at eight factors that will influence the future of cyber-space:

- the present and future character of the Internet, since it is a central component of cyberspace today: it is an example of the *platform* character of cyberspace
- the future shape of computing devices, which will shape networking (and thus cyberspace in general) and define the basis for the user experience in cyberspace
- the changing nature of information and the tools and mechanisms to deal with it
- the emergence of network design principles at a layer higher than simple data transfer
- the nature of the future user experience in cyberspace, drawing on today's research
- the challenges associated with security and responses thereto
- the role of private sector investment in defining the future
- the health and character of research on the science and technology of cyberspace.

For each factor, our discussion tries to bring out issues and implications for cyberpower, such as the tools for control and the balance of power among different actors and stakeholders. We also note places where we can expect differences in the nature of cyberspace in different nations.

Platforms: The Centrality of the Internet Today

There are many aspects to cyberspace, from the computing and communications

infrastructure, through the information that is processed and transported, up to the users that operate in cyberspace. We consider a number of these aspects in this chapter, but we focus on the Internet because of the central role it plays in so many of these dimensions.

The importance of the Internet arises from its design and the purpose for which it was conceived. The Internet was designed not to support a specific application, but with the goal of generality. In contrast to the telephone network, for example, which was initially designed specifically to carry telephone calls, the Internet was designed to support a wide range of applications, even those not yet thought of. And indeed, new applications such as the World Wide Web were later conceived and deployed to extend the general service provided by the Internet.

The Internet provides a general platform for applications and services on top of a variety of network technologies, ranging from high-speed optical links and local area networks to wireless connections. This generality means that the application designer need not be concerned with the technical details of specific network technology. The use of different technologies can, however, affect the way different parts of the Internet perform, so one can design an application and later adjust parameters such as throughput or resilience by changing the technology over which the application is operating.

The Internet illustrates a critical characteristic of cyberspace: that it is built out of components that provide services, and these services are designed so that they can be composed and combined to form ever more complex services. Low-level services include program execution environments, mechanisms for data transport, and standards for data formats. From these are built applications, such as a word processor, a database, or the World Wide Web. By combining these, more complex services emerge. For example, by combining a database with the Web, we can get dynamic content generation and active Web objects. Cyberspace features the continuous and rapid evolution of new capabilities and services, based on the creation and combination of new logical constructs, all running on top of the physical foundations. Cyberspace is a sequence of platforms on which new capabilities are constructed, each of which in turn becomes a platform for the next innovation; it is thus very plastic. Cyberspace has emergent qualities: it is the consequence not of a coherent or centralized design, but of the independent contributions of many actors with various motivations and capabilities.

Technically, the essence of the Internet, defined by the Internet protocol (IP), is a set of standards that describe how data to be sent is broken up into *packets* (small units of data together with instructions for delivery). IP defines the format of those packets, how the destination of the packet is specified, and so on. The specification at this middle level, above the transmission technology and below the application, has sometimes been illustrated by drawing the Internet as an hourglass (see figure 8–1); diverse transmission technologies support the packet transport service of the Internet, and diverse applications sit on top of this service. Defining a level of standardization (the narrow waist in the illustration) supports a wide range of diversity and innovation in other parts of the system (the wide top and bottom in the illustration).

Figure 8-1: Internet Protocol

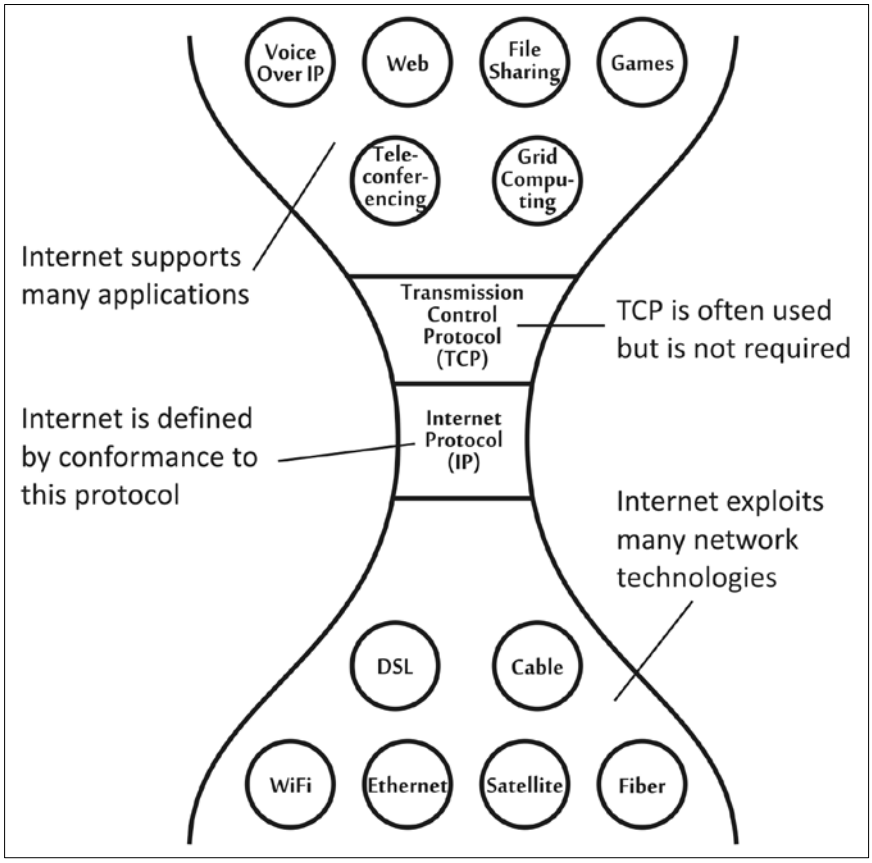

Source: Computer Science and Telecommunications Board, *Realizing the Information Future* (Washington, DC: National Academies Press, 1994), 53.

Today's wireless technology illustrates how different kinds of networks that came into being at different times may coexist within the Internet, with different implications. Cellular systems are centrally administered and have global reach; WiFi networks are deployed at the edge of, but are seen as part of, a global network; Bluetooth wireless systems are very small scale, operating without wide-area connectivity but capable of serving as edge systems if connected to a global network. All of these technologies, sitting "below" the narrow waist in figure 8–1, are being used today to carry packets as part of the Internet.

What Does** Internet **Mean?

The term *Internet* can mean a number of things, which we must tease apart. In most cases, the term is synonymous with the public Internet—the globally interconnected service so many of us use. The term can also refer to the standards—the specification of the technology of the Internet—and to the technology base itself, the set of products consistent with those standards. This distinction is important because many of the Internet products purchased today do not end up in the public Internet, but in private deployments, restricted to a corporation, government, or other actor. It is often carelessly said that "in the future, all

communications will be over the Internet." It is easy to see how this assertion might be made, as telephony and television begin to migrate to the Internet, and it is used to build control systems (such as supervisory control and data acquisition [SCADA] systems for the power grid and other critical infrastructure, as described in chapter 5 in this volume, "Cyberspace and Infrastructure"), and so on. However, this masks an important distinction: these applications are being redesigned to take advantage of Internet technology, but this does not mean that they will all be carried over the public Internet. Many private or closed networks, including networks for telephone and television as well as for large users in government and the private sector, are being built out of Internet technology, because it is general-purpose, inexpensive, and available.

An example of both other services being carried on an "Internet" base and IP networks separated from the public Internet is provided by Britain's telephone provider, BT Group, and its ongoing effort to transfer all its communications traffic to an IP backbone. This offers simplification and savings in the costs of future upgrades. In contrast to "Voice over IP" (VOIP) services that run on top of the public Internet, BT Group's telecommunications system will operate in essence as a separate, private Internet. This isolation is intended to enhance security and reliability.[1]

All uses and forms of the Internet are not equivalent: there is tremendous diversity in the variety of technology, the degree of interconnection, the range of use, and the controls on that use. Recognition of this point, with careful choices about technology, the degree of interconnection, and so on, may affect cyberpower.

In general, the architecture of the Internet helps to define what it is used for, how it works, and how it is composed. The interaction of use, operation, and structure points the way to potential paths for change as well as to levers or targets for cyberpower. Possible revolutionary change in cyberspace will influence both the future of the "public Internet" and the future penetration of Internet technology, potentially in different ways.

An Alternative to the Present Internet

One of the goals of this chapter is to illustrate the myriad possibilities for the future of cyberspace. *Virtualization* is an example of an alternative approach that might lead to a very different global network in 10 to 15 years. Virtualization has the objective of providing a greatly enhanced ability to exploit different network technologies. The Internet's service, characterized by the narrow waist of the hourglass in figure 8–1, can be implemented over a range of network technologies, since it only requires the technology to carry strings of bytes as packets. If these technologies have other capabilities, the applications sitting on top of the IP layer cannot exploit them. The simplicity of the IP layer is both powerful (in terms of generality) and limiting (in terms of exploiting technology features). For example, radio is intrinsically a broadcast medium, but the broadcast feature cannot be exploited in the Internet. Fiber is capable of carrying a very high bandwidth analog signal with very low distortion, but this quality cannot be exploited using the Internet design.

Virtualization, a very familiar concept in computer science and electrical engineering, attempts to make the advanced features of any network technology directly available to the application. It is a set of techniques that divide a physical set of technology resources among multiple uses in a way that gives each use (or user) the illusion of having a separate and distinct copy of these resources, essentially identical except that the copies run slower than the actual technology base that is being virtualized. We know how to take a computer and use

software that creates a number of *virtual machines.*[2] We can divide up and virtualize circuits, memory, and so on. Now research seeks to virtualize all the resources found in a global network and then give *virtual global networks* to different classes of users or uses.[3] There is no commitment to a common packet format (indeed, to packets at all) or to any of the other current conventions of the Internet. Different applications can use the resources in ways best suited to their needs: one set of users might run something like the Internet in one virtual global network, while another set of users might use a virtualized fiber optic path to carry analog signals, and a third set of users might use a virtual global network for global dissemination of bulk data, with a totally different approach to packetization and routing. User behavior, and consequences, such as traffic patterns, might become less predictable.

Compared to the illustration of the Internet, this approach moves the point of agreement—the narrow point of the hourglass—down, to the technology base. It demands of each technology not that it be able to carry packets, but that it be able to "virtualize itself."

So there is an interesting turbulence in the future of networking. As many industries (such as the cable television industry) are moving away from their current approach (transmission of analog signals) to a digital approach based on the existing Internet technology, in order to exploit the benefits of a low-cost integrated platform, the platform might evolve in an attempt to support what these industries already do.

One particular application area that might benefit from the ability to exploit the specific features of underlying technology is the use of wireless technology on the battlefield. This extremely demanding and hostile environment may require that we abandon the generality of the existing Internet design for a new kind of generality that allows applications to more directly exploit and benefit from the features of advanced wireless systems. But how best to do this, and what this means for a global network of the future, is still a question for the research community. Research along these lines has been a priority for the Defense Advanced Research Projects Agency (DARPA).[4]

Thus, a future global network might be based not on a common packet format, but on a set of conventions at a "lower layer," that is, the ability of the resource to be shared using virtualization. This approach, if successful, could lead to a network in which applications are able to exploit the native capabilities of current and future network technologies. If the virtualization approach provides more flexibility to the application in how it exploits underlying technology, it also raises fundamental architectural questions about the division of function between the virtualization layer and the layers above it—questions about security, management, the economics of the industry, and so on. And it invites the question of whether the associated generality is actually more valuable than the generality of the Internet of today. This concept, if it proves itself and moves out of the research community, could offer a very different future global network.

The Changing Face of Computing: Getting Fast and Getting Small

The shape of tomorrow's networks will be defined by the shape of tomorrow's computers. In 10 years, the range of networked computers will be much more diverse than what we see today; one trend leads toward massive server farms and other high-end computing platforms, and another toward dense universal deployment of very small and inexpensive computers embedded in everything. This changing nature of computing will influence many aspects of cyberspace. It will greatly expand the range of applications and functions that can

practically be supported, and it will change the nature of the Internet; since the Internet was designed to connect computers, as computers evolve, so does the Internet. The early Internet of the 1970s featured long-distance access to large, centralized mainframes by so-called dumb terminals; the last two decades, by contrast, have been dominated by personal computers (PCs), small systems that bring computing intelligence to the user. PCs popularized computers and the things that could be done with them, thanks to continued reductions in size and weight and the progressive linking of computing with communication, first through private networks and then the commercialization of the "public" Internet, while numerous innovations in software made more and more activities first possible and then easier. In an important sense, the Internet and the PC have co-evolved, and the apparent stability and maturity of the Internet is due to the maturity of the PC.

Now we are at a point where advances in computing are overturning the PC as the dominant paradigm. Thus, the next decade may be marked by rapid changes in the nature of computing and networking. PC-based computing as a dominant paradigm may well be displaced by embedded computing, sensors, and other applications of small, inexpensive processing. Proliferation rather than convergence of devices should be expected.

Continuing reductions in size and cost for equivalent performance will lead to the embedding of computing in other devices. This trend, in concept a decades-old phenomenon, is facilitated by the emergence of smaller and smaller complete platforms, which today can be the size of a stick of chewing gum or smaller.[5] Much of the value of such devices may come from their ability to communicate—they are likely to use networks, if intermittently—to share information or to give or receive directions.

The second trend in computing is at the high end, where Moore's law (see chapter 6 in this volume, "Evolutionary Trends in Cyberspace") brings ever more powerful processing. What is important is not so much the continued rapid improvements in the speed of single processors as the emergence of parallel processing at all scales. Even consumer PCs today have multiple processors ("cores") in them. Much high-end processing is based on various sorts of parallel processing. This trend will continue. What this implies for applications depends on the nature of the application itself. Some classic, computation-intensive applications, traditional drivers of supercomputing, are difficult to implement on a parallel processor system.[6] On the other hand, many important applications are naturally parallel, such as various sorts of search, pattern recognition, media formatting, and transformation. This sort of parallel processing can be seen today across processing platforms that are networked at all scales, from integrated multiprocessors, to grids (collections of computers hooked together by high-speed Internet functionality), to massively parallel processing performed on consumer PCs connected by residential broadband.[7] Such distributed, consumer-based problem solving can be seen as a positive and intentional version of the "botnets" used by criminal hackers (described in chapter 7 in this volume, "Information Security Issues in Cyberspace").[8] To the individual, this trend toward collaborative, massive, highly parallel computing can be seen as the next stage, after citizen empowerment by the PC, in the democratization of computing. All of these paradigms of parallel processing have implications for the performance and scale of the networks that connect them. The current trend is to push the function and performance of the Internet in directions that support this class of application.

The future Internet will thus be pulled in two directions by these trends in computing: toward higher speeds, by the need for high-performance parallel processing, and toward lower cost and ubiquitous access, by embedded processing. In 10 years, a single network architecture

might evolve that can serve this greatly diverse set of demands, or we might instead see a split in the technology and architecture, with high-end processing and embedded processing being served by different Internets.

How computers are used is changing with their physical form. There has already been movement, especially in developed countries, from the one-PC-per-person model of the early 1980s to multiple computing devices per person, including not only conventional PCs but also other devices that incorporate computing capabilities. We can assume that in another decade or so, the typical person, in the developed world at least, will be immersed in a pervasive matrix of computing and computers, including both the processors that the person carries and the processors embedded in the environment. All of these devices will interact, which will have many new implications for the architecture of cyberspace.

A third trend in computing is the rise of massive commodity server complexes.[9] The low cost of PCs has benefited not only the consumer, but also the application designer who can build a service that runs on a huge array of such devices. More than just a special case of parallel processing, the scale of current server complexes warrants special recognition.[10] The high-end server farms run by Google, Yahoo!, and Microsoft are so large that they are recognized in the industry as effective supercomputers.

Thus, we see a number of trends that signal a major change in the way applications are designed.[11] Google Search, for example, is highly parallel, running on servers that are physically centralized and under central control.[12] Projects such as SETI@home are also controlled centrally, are highly parallel, and run on commodity PC platforms, but they differ in that ordinary consumers contribute the processors as a volunteer gesture. Consumers also contribute their processing capacity to activities with little or no central control, such as peer-to-peer music sharing. Multiple approaches coexist and support different kinds of applications; there is no single preferred or more likely outcome (especially absent any steering or constraints from law or regulation). The broad trend, illustrated by evolving grid systems, is for increasing distribution—of data, software, and/or computational resources—and for benefit to individual entities from exploitation of distributed resources.

Another broad trend is a movement from people communicating with devices, or through devices to other people, toward devices communicating with each other. This trend raises questions about who or what is a "user" of the Internet (or cyberspace), an issue complicated by customary use of *user* to apply to both individuals and organizations, and to both producers and consumers of content and applications.

Embedded Sensor Networks

Many of the issues discussed above are illustrated by the development of networks of embedded sensors. Considerable work is under way in the research community to develop sensor networks embedded in a variety of physical environments. These networks will monitor such phenomena as ecological processes, microclimates, seismic activity, and pollution (for example, contaminant and fluid flow). Complementing remote sensing via satellite systems, these networks enable in situ exploration of an environment. Such sensor nets can feature from tens to thousands of devices of different sizes that connect to each other using wireless communication. They are constrained in terms of power, capacity to store information, and capacity to communicate, but technology development is improving all of these factors. These networks may be more or less dense and may be long- or short-lived; improvements in

technology and reductions in cost will even enable disposable or redeployable networks for some applications. Researchers' use of the term *smart dust* captures the potential.[13]

Current sensor networks use the public Internet to a limited degree, as systems that connect to it at its edge. For example, clusters of sensors in a given location, which process information collected locally at a local server before passing it on, may feed into networks serving a larger area; they, too, may undertake processing and aggregation of information before passing it to a higher level network.[14] The Internet of today, per se, may not be engaged directly by clusters of sensors, let alone individual sensors, but it does play a role in their interconnection.

There are many other kinds of sensors currently deployed that may have communications links but tend to be focused either on a local task or on the central monitoring of information from remote points. Examples include traffic control, home and industrial automation, industrial and infrastructure control (SCADA systems), and deployment of communicating identifiers (radio frequency identification [RFID] systems) for such applications as supply-chain management. It can be assumed that in the future, these sorts of applications will be fully networked and integrated into the larger cyberspace context. Cars and other motor vehicles, which have long had embedded sensors and computers that operate as parts of a closed system, present possibilities that range from peer-to- peer communication between vehicles (somewhat analogous to the peer-to-peer communication between airplanes that is contemplated for air traffic control) to communication between cars and other systems that could contribute to traffic monitoring and control.[15] Taken together, networks of embedded sensors all involve proliferating information technology that, to varying degrees, collects data, processes it, communicates, and automates those processes.

Contemporary sensor network research has focused on a class of networks that are self-configuring (deployed sensors establish a network among themselves), self-monitoring, and, if a problem is detected, self-healing. These features contribute to making these systems of sensors relatively autonomous and easy to deploy. Automation intended to enhance management of the system also presents the prospect of added vulnerability, because it implies dependence on programming and therefore the potential to be compromised. On the other hand, automated sensor systems also increase people's ability to interact with sensors in order to monitor data quality and control their data collection or movement in real time; thus, systems may be more responsive both to circumstances and to the controller's interests. Again, however, a benefit for functionality may also generate vulnerability if the ability to interact with sensors cannot be controlled perfectly.

Research points to the potential for sensor networks to become more versatile. For example, some systems are gaining actuation capability, the ability to control other systems. Historically, many control systems (for example, for machine tools or building access) have been automated, but only locally; people had to intervene periodically. By contrast, removal of the need for people in or near the control process presents both opportunities and vulnerabilities. In addition to controlling other systems, some sensors may be able to actuate themselves: to move to where they can be more effective, or to explore different points in their general vicinity.

The introduction of sensor networks into social situations also presents possibilities. Leveraging lessons from natural environment applications, more research is focusing on urban situations that combine features of the built environment and human behavior. They might, for example, make use of worn or carried devices, perhaps cell phones, and a human

propensity for exhibitionism or voyeurism, such as Webcams, installed either by people recording themselves, or by government entities or other organizations. The ability to download software to cell phones for new functions allows cell phones to aggregate and communicate sensed data as well as to support interpersonal communication. Cell phones also illustrate how unpredictable the mix of automation and human agency can be. For example, a cell phone can be seen as a networked sensor that can operate independently (as opposed to being part of a sensor network by design), which raises questions about the potential to engage groups of cell phones as impromptu or ad hoc sensor networks.[16] The use of cell phones to trigger flashmobs (and even invoke vigilante justice), and the emerging use for relaying official emergency alerts, suggests their use as actuators, if indirect, for human behavior.

Implications of Sensor Networks

Sensors can be cheap, especially those without sophisticated instrumentation; technology development is likely to continue to decrease costs as well as size. Sensors are available to virtually anyone.[17] The proliferation of cell phones and of new uses for them indicates how ordinary citizens seek to monitor their environment and communicate the results with each other, illustrated by the sharing of photos via cell phones (and the further sharing of cell phone–captured images on the Web). Over time, more kinds of sensing modalities are likely to become available in consumer devices. Easy acquisition of sensors and their connection via open standards suggest possibilities for broadening commercial, public, and citizen use of sensor networks. We can anticipate that in the coming decade, any interested individual could deploy a set of sensors. More use of sensors may imply added societal vulnerability, for example, to electromagnetic interference (perhaps intentional); electromagnetic pulse, which could destroy large clusters of sensors; or tampering (as more and more sensors are scattered in the physical environment without supervision). More use of today's wireless technology itself implies added vulnerability, which sensor networks compound.[18]

Some sensor networks (for example, those the government uses in connection with weather) are deployed and managed top-down, while others are bottom-up, with sensors connecting according to choices made by users. A bottom-up system can operate in the same arena as a top-down system; for example, the Weather Underground system, using aggregated information from a wide range of local sensors,[19] complements government weather data. The ability to connect from the bottom up, peer to peer, can provide resilient connectivity: such networks may use a kind of cell structure, in which independent cells continue to operate even if one cell is knocked out. That independence can also increase the challenge to those who seek to monitor or control such networks. The emergence of botnets and the varied motives of their controllers suggests the analogous possibility of commercially useful sensor networks and their deployment and exploitation for productive rather than malignant purposes.

If sensor networks proliferate, how might they be interconnected? A future Internet might provide the technical means, but there are also organizational and even political issues. For example, current sensor networks are deployed by different owners, with different attitudes about the merits of interconnection. In the research community, a collaborative ethos and the incentive of funding drive efforts to interconnect different environmental sensor networks to promote greater understanding of the natural environment;[20] this also fosters greater understanding of sensor network interconnection per se. These projects suggest prospects for dual use, inasmuch as

data collected for one purpose might be used for others; indeed, a hallmark of cyberspace is abundant data used (eventually) for abundant purposes. Although current efforts to interconnect sensor networks seem to emphasize hierarchy, which implies some degree of coordination and control, broader trends make it easy to envision more peer-to-peer activity.

Information Ecology

Cyberspace is defined not just by technology and standards for communication: the purpose of information technology is the manipulation, storage, and communication of data. The nature of information is changing as fast as the technology that processes it, and the trends surrounding information will be a key determinant of a future cyberspace. We consider several major trends. First, there is an increase in the volume of sensor-based information, complementing the business data and the human-created content that dominate today's information. Second, the highly distributed and decentralized creation of content, from blogs and wikis to mashups,[21] challenges the traditional models of professional creation and editorship and, by encouraging a proliferation of information and sources, may impede any single party's ability to get its message out. Another implication is that the proliferation of information sources will allow users to ask to see only what they prefer to see, which may reinforce the polarized positions and opinions of special interest groups.[22]

In addition, the landscape of search is changing: search tools such as Google provide the power to search the whole Web, tools emerge for search of nontext content such as recorded speech, and metadata systems are more broadly implemented. These tools are available to everyone, not just to specialized high- end users, and their power and capabilities will increase. Another trend is that information as it is presented to each user is increasingly personalized: each person may see a version of the content expressly generated for that individual. These trends raise important concerns about control, quality and reliability, sharing, and options for manipulation and distortion.

New Sources of Information

The two traditional sources of online information have been human creativity— the collections of writings, utterances, and images produced by people for other people to consume directly[23]—and business data, such as the capture of online transactions, inventories, and so on. The proliferation of sensor networks predicted above would accelerate increases in the volume and variety of information that is available. Sensor data may call for different treatment than business data. With business transactions, it is important to capture and record each event reliably. With sensors, however, individual readings may not be so important, and together they are too voluminous to store individually; what is important are the trends, aggregates, and summaries of the raw data. Some researchers are exploring options for sensor output processing and aggregation, and others are exploring the designs of new sorts of databases.

A larger trend, accelerated by the emergence of embedded sensors, is geolocation: people (cell phones), things (RFID), and information (data tags) are increasingly being associated with their locations, both at a given point in time and over time, as they move around. Where sensor networks target people or social behavior, and even sometimes when they do not (as with routine utility monitoring), there can be privacy impacts, inasmuch as these systems collect personal information and permit inferences about behavior. Growth in sensor- delivered

information also puts a spotlight on its quality: sampling can affect the trustworthiness of sensor readings or interpretation, and sensors can have faults and might be tampered with.

That there is more information also implies that there is more potential for information to be manipulated and more incentive to do that manipulation. In one sense, manipulation is a goal in itself: collaborative generation of content implies, or makes explicit, an invitation by the originators of information to others to manipulate it and improve on it. This is the concept behind Wikipedia (and of wikis in general).[24] A related concept can be seen in blogging, where the blogger elicits comments and responses. At present, the phenomenon may be too new to predict where the revolution will end up, and among some "digerati," there is a vigorous debate about the relative wisdom of the collectivity versus the expert.[25] New streams of information also imply new possibilities for combining information from different sources.[26] New businesses may offer to sell that capability.

As more individuals take advantage of ease of access to generate and share content, more questions arise about the impact of the resulting volume of information, the ability to find useful elements, and the role of editorship. A related question for a national cyber strategy might be what it takes to get a nation's message out to a desired audience. At a national level, this may be like asking what the Radio Free Europe–equivalent of cyberspace could be. Given the uneven history of national broadcasting to foreign populations, putting a national brand on information dissemination may or may not be practical or desirable. A future analog of a Radio Free Europe might instead be more concerned with providing access to a wide range of information rather than packaging and disseminating selected information. This view is consistent with recent U.S. Department of State activities relating to information and communication policy.[27] (See chapter 14 in this volume, "Cyber Influence and International Security," for more on these issues.)

Search

Finding useful information is a perennial challenge. Two trends may combine to transform the way searches are carried out in cyberspace. The first is the growing power and sophistication of software for automated recognition and understanding. Today, text-based processing and analysis support tools ranging from Google to LexisNexis. Emerging from the research lab are programs for image recognition, speech recognition (now used commercially in a variety of contexts), and language processing. Tools such as these may make it possible to search an audio file, such as a recording of a conversation, a commercial broadcast, or a podcast, the same way that text is searched today. The other trend is the emergence of the Semantic Web, a set of standards and methods that allows the creator of a Web site, or a third party, to tag that site with information about its content (called metadata) in a form that is easily processed by a computer program.[28] How the information is labeled or described creates the potential for finding and combining information, but also for hiding from or redirecting searchers.[29]

We can expect individual elements of information to be embedded in a rich context of linkage, metadata, translation, transcription, search aids, and commentary. The ability to control this context—what the user can find or is prevented from finding—may be as important in an analysis of cyberpower as is access to the elemental information itself. If information systems are your "eyes," part of a cyberstruggle will be to blind you, lie to you, or cause hallucinations.

The immense processing power of Google, made available to all users of the Internet for

free, may be one of the most profound examples of the leveling of the playing field that is associated with the Internet. Nations desiring to control information access have found ways to moderate Google's impact, but this is only a change in degree.

Systems such as Google raise other questions. First, like other existing search systems, it can only search the Web pages it can find. There is also a so-called dark or deep Web that is not reached by the current crawlers, and perhaps because the information has no links that point to it so the crawlers cannot know it is there, it is protected from open access or crawler access, or it is created only on demand. If an increasing amount of information resides in this deep Web, the scope of information and the validity of search results will be affected. Second, the economic value of search and the complementary products generated by Google have inspired an attempt in France and Germany to mount an alternative.[30] Meanwhile, in China, a national effort to degrade Google service quality has made indigenous alternatives more competitive.[31] The ability to shape the user's view of search results is clearly seen as cyberpower on a national scale.

Dynamic Information and the Personalization of Content

Another important trend is the dynamic generation of information in response to queries. More Web pages are generated on the fly out of component parts when a user asks to retrieve them, so that each user gets a personalized version of specially selected content. This personalization can be done for many reasons; for example, an online newspaper might be tailored to the interests and priorities of each reader; different advertising might be presented to each user based on the profile maintained on that user by the advertiser; or different pricing might be offered at an e-commerce site.

There are a number of implications of dynamic, personalized information. For librarians and archivists, this trend implies that there is no definitive version of a piece of information—no reference version that can be universally referenced or cited. Any producer of information may, literally, say something different to each person.

Future Architectural Concepts at a Higher Level

Current research may lead to higher level services that might define an Internet of tomorrow. These services would involve not just simple point-to-point data carriage, but more complex functions closer to what the user is trying to do and further away from moving bytes. These future services will better match the needs of sophisticated applications. We describe three examples of such research: support for an information architecture; support for the creation and management of services that depend on functions placed on servers distributed across the network; and support for communication among nodes and regions of the network that are connected only intermittently.

A Future Architecture for Information

Some current research aims at an architecture for information as opposed to an architecture for byte carriage. Some examples of what such an architecture might capture are outlined here.

On the Internet today, the authenticity and provenance of an information object are validated by where it comes from. For example, in today's Internet, the only way a user can be

sure that a Web page that appears to be from the Cable News Network (CNN) is legitimate is to try to retrieve it from CNN. Once it is retrieved, there is no trustworthy information associated with the page itself that captures and conveys the provenance of the information. If one user downloads a page and then sends it on to a second person, there is no way for that second person to be sure it is not a forgery. If the validity of the information were associated with the information itself, perhaps using some sort of encryption scheme to sign the information, then we could use a much richer set of mechanisms for information dissemination. These might include peer-to-peer systems, third-party caching, archiving, and forwarding, or casual user-to-user transmission via email. This richness would provide a network that was more diverse and efficient.

An architecture for information could include a means to name information. Current Internet names (domain names) and addresses do not name information objects; rather, they name physical endpoints—that is, computers attached to the edge of the network. The Web provides one scheme for naming objects, the uniform resource locator. There has been much research on other schemes for naming objects, ranging from Internet Engineering Task Force research on universal resource names[32] and centrally managed schemes such as the Handle scheme proposed by the Corporation for National Research Initiatives,[33] to highly decentralized schemes that allow users to pick local names and share them.[34] All of these schemes imply some solution to the location problem: given a name for an object, can I find the object itself? Some solutions to the location problem would involve a centrally managed catalog of objects; others would use a more decentralized approach, with location services run by different operators in different regions of the network; some schemes would be completely decentralized, based on a user broadcasting a request for an object until a copy is found. These different schemes have very different implications for control, resilience, performance, and utility. Our concern here is the implications for control, a core aspect of cyberpower.

The problem of *search* is distinct from the problem of *location.* The location problem is to find a copy of an object once I know exactly what I am looking for. The analog in the real world is to ask what library or bookstore has the closest copy of a specific book. The *search* problem is to look for objects that may match certain selection criteria, analogous to seeking a book on the history of baseball that costs less than $20. Google's approach for finding a good match to given search criteria, for example, is quite complex and sophisticated, but Google facilitates the search process by keeping a centralized copy of every candidate page from the Web that it can discover. This massive database is very complex and costly to maintain, but it simplifies the subsequent search process. An important search question is whether it is possible to perform useful and effective searches without centralizing all the candidate material.[35] Another problem that has a critical influence on searches is the emergence of dynamic information, described above, that is generated on demand for each user. It is not clear that current search systems reliably or consistently index information that only exists when a particular user asks to retrieve it. There is ongoing research in the design of advanced search engines that could search or otherwise deal with dynamic content.

The importance of this research to the future of cyberspace is that, depending on how such efforts succeed, we may be able in the future to place either much more or much less confidence in the validity, authenticity, completeness, and persistence of information in cyberspace. Fraud and forgery could destroy all important uses of online information or, at the other possible extreme, cyberspace could become the only important and relevant source of information.

An Architecture for Services and Service Construction

The architecture of the current Internet really recognizes only two sorts of devices: the devices that make up the Internet itself—the routers that forward packets—and everything else, including all devices attached to the edges of the network, between and among which the Internet packets are transported. The dominant examples of edge devices now are the PCs that belong to the users and the servers that provide content to users. In the Internet architecture, there is no consideration of what these servers do, what patterns of communication connect them, or how applications actually exploit servers to build up their services.

Most applications today display a somewhat complex structure of servers and services to implement their functions. Content may be prepositioned on servers near the ultimate recipient to improve performance and resilience.[36] Communication among end nodes is often relayed through intermediate servers to realize such functions as mutual assurance of identity, selective hiding of certain information, logging of transactions, or insertion of advertising. Some servers, such as email servers, hold information for recipients until they choose to connect to the network to retrieve it. This allows nodes to communicate even if they are not simultaneously connected and active.

Once we recognize this rich structure, it is natural to ask whether the design of the Internet should take account of it in some way. For example, today's Internet does not provide a direct way to ask the question, "Which of these servers has the shortest path to a specified set of recipients?" The routing system knows the answer to this question, but there is no easy way for an application to retrieve the information. In another example, the Internet has no way to take account of or respond to the different traffic patterns from different clients and different sorts of servers. A particular server may, for example, typically receive content from only a certain set of sources but send information on to any number of clients. If the network could incorporate this knowledge, it might be able to detect nonstandard patterns of communication as a security attack.

The emergence (and recognition) of systems that incorporate servers and services may signal a move of processing and control from the edge of the network—the computers of the end users—toward the center, locations operated by providers. Providers might include those who supply the basic packet carriage of the Internet, as well as specialty providers of specific application services and providers of commodity platform computing on which many services can be hosted.[37] These types of players may become important parts of the Internet experience of tomorrow. At the moment, there is no evidence that the movement of function toward the center implies the emergence of new monopolists, but this is a factor to analyze and monitor, along with potential government exploitation of such centralization. Players at many layers of the Internet—including providers of residential broadband service, providers of core software products such as Microsoft, or higher layer providers that deal in information, such as Google—may have significant market power now or in the future. Areas that will warrant attention may include advertising, marketing, and the collection and management of information about consumers.

A basic aspect of such an analysis is to identify all the components that make up a system, and ask who controls these components and what options exist to influence that control. The emergence of servers as a distinct element raises the question of who controls the servers, who has the power to select which servers are used to realize a given service, and so on. This adds

richness to the dimensions of power and control in cyberspace.

At the same time, we can see the possibility of a future with increasing regional differences and a possible balkanization of the network at these higher layers. Even at the basic packet-carriage layer of the Internet, which is defined by a very simple universal standard, there have been attempts to control the flow of information across regions of the Internet. These higher layer services are more complex, with richer structure and more visible signals that may reveal the intentions of the users, and they provide many more options for points of control. The experience may differ depending on the user's location; many actors may try to shape and limit the user's experience.[38] Differentiation and fragmentation of users' experiences have many implications for cyberpower.

One of the specific issues that arises in these more complex, server-based applications is that the physical location of the servers is relevant, and indeed sometimes critical, to the effective operation of the application. If the purpose of the server is to preposition content close to the user to increase performance and reliability, then the server, because it is physically close to the user, may be more likely to be under the same legal or governmental jurisdiction as the user. It may be that services that are designed for performance and for resilience in the face of network failures will have a quite different design than services that position the content at a distance from the user in order to avoid the imposition of local controls and balkanization.[39] Research could help to enable such geographic independence. A continuing tension between technical mechanisms and legal or other nontechnical mechanisms is likely, given competing interests in disseminating or restricting specific kinds of information in specific locales.

An Architecture for Relayed Delivery of Content

A central assumption of the current Internet architecture is that communication between pairs of nodes is interactive in real time, and since any pair of nodes may want to establish such interactive communication, the network provides this capability universally. The dominant protocol for data transfer on the Internet, the transmission control protocol (TCP), is based on the immediate confirmation to the sender of the delivery of each packet. But if applications can define restricted patterns of communication involving intermediate servers, delivery from the original source to the ultimate destination is not necessarily immediate but may instead be delayed, staged, or relayed. This pattern of communication seems to be common to many sensor applications and information dissemination applications, so it is posited that this pattern should also be recognized as a part of the core Internet architecture. A current line of research involves the investigation of delay/disruption tolerant networking (DTN), which generalizes this concept of "store and forward" networking (in contrast to "end-to-end interactive" networking).[40] DTN raises many important questions about security (how much we must trust the intermediate nodes), resilience (how the service might recover from failed intermediates), routing (who controls which intermediates are used), and assurance of delivery (whether the design provides confirmation of delivery that can be relayed from recipient to sender). At the same time, they allow the network designer great freedom in dealing with challenging contexts such as poor and intermittent connectivity or hosts with intermittent duty cycles, and they allow the application designer to hand off part of the design problem to the network.

A Long-term Outcome: Revolutionary Integration of New Architecture Ideas

Above we described an alternative network design based on virtualization, in which the basic network service is not packet carriage, but access to the network technology at a lower level and more technology-specific way. This lower layer idea gains importance when it is combined with a higher level architecture for services and delay-tolerant delivery. If information is being delivered in stages, not directly from source to destination, then the details of how the data is transmitted can be different in each stage. The requirement for a uniform and universal commitment to a single packet modality, for example, is much reduced if different parts of the network talk to each other only via a higher level server. The result, in 10 or 15 years, might blend all of these ideas in a global network that is based, at a lower level, on virtualized network technology and, at a higher level, on application-aware servers that connect parts of the network in ways that are matched to the features of the technology in that region of the network.

Compared to the Internet of today, future architecture may focus on the specification of features that are "higher level"—closer to what the application designer and the user are trying to accomplish. Future networks may focus on architecture for information-handling services built out of distributed servers and staged delivery. Attention to architecture at these higher levels may provide an alternative to today's focus on common packet formats and allow the lower layers of a future network to more directly exploit features of diverse technology.

Some basic features of the present Internet, such as global end-node packet-level connectivity, will not be as important in a future where most machines only communicate via application-level servers.

The Changing User Experience in Cyberspace

Our view of cyberspace is that it is best understood as a series of layers: technology, platform, information, and human. In this section, we complete our look at these layers by an assessment of the user experience in cyberspace. We identify several important trends:

- the increasing importance of physical location as a factor influencing what happens in cyberspace
- the increasing interconnection of real and cyberspace, with the ability of the user to move fluidly back and forth between the two viewpoints
- the importance of massively multiplayer online role-playing games (MMORPGs) and other shared virtual experiences as early signals of the rich range of options for people to interact in cyberspace
- the ability to earn money and make a living by working in cyberspace.

Users are engaging cyberspace not only more overall, but also in different ways. For example, instant messaging (and its equivalents, such as chat features embedded in different kinds of applications) introduces a new kind of immediacy to online interactions; it implies an ability of the system to monitor and register presence (whether a user is online at any given moment) and reachability. Augmented reality, combining real space and cyberspace in different ways, is becoming more common. For example, navigation systems (portable or in vehicles) can now link one's location to information about nearby businesses and other sites. Other products similarly blend the virtual and the real, based on a user's location.[41] More generally, easier access combines with different kinds of activities to make the Internet more than a place to find things to look at, read, or buy; for many, it is becoming a true alternative

venue for daily life.

MMORPGs, which draw on virtual reality technology, demonstrate how cyberspace can become more engaging and immersive, hence more broadly integrated into people's lives. Evidence is the use of real currency to acquire items found and used only in games, and the use of virtual product placement as a kind of advertising. In some contexts, MMORPGs can range from being a team sport to a cause for social concern about Internet addiction; the latter is often reported in South Korea, where comparatively early widespread access to broadband facilitated extensive involvement with games. The capacity to involve people to the point of addiction or self-neglect is a kind of cyberpower, albeit presumably not one intended by the application providers.

MMORPGs embody the development of synthetic worlds, which are also being considered for a range of real-world activity such as training.[42] Some synthetic worlds and many other areas of cyberspace—from eBay to gambling sites and even some blogs and other sites that host advertisements—allow more and more people to "go to work" and make "real" money in cyberspace, sometimes obviating the need for conventional jobs. MMORPGs are also an instance of social networking, itself an evolving genre. The Internet is making possible much larger social networks, and their scale (which can range to the tens of millions of users) is already motivating researchers to understand the kinds of social links being formed, the different roles people play in diffusing information through networks (as discoverers, amplifiers, reshapers), and the potential outcomes of the interactions of large numbers of self-interested parties competing for resources in online environments. In contrast to the physical world, the software basis of the online world helps to constrain and capture the interactions, making explicit some of what might have been implicit and generating detailed records of behavior.[43]

The example of MMORPGs illustrates an important distinction between cyberspace (and the resulting features of cyberpower) and other realms of power. While realms such as sea and air are physical and are strongly constrained by physical reality, cyberspace is much more a synthetic realm, constrained less by physics than by imagination.

The potential for people to spend more time (and money) in cyberspace raises questions about distinctions between who can and who cannot do so—the so-called digital divide. In some locales, individuals have their own computing systems and personal access to cyberspace. In others, such as many locations in Asia, cyberspace is accessed communally in cybercafés and other public access places. In yet others, such as some locations in Africa, cyberspace can be accessed only in a public place from which most people are remote, and it is therefore experienced infrequently. Various "appropriate technology" projects aim to overcome limitations of local power and communications infrastructures, holding out the possibility of greater universality of cyberspace access and use over time.[44]

Frequency and continuity of access may vary by choice as well as by ease of access. For example, if the stereotype that, other things being equal, younger people enjoy being online more than older people is accurate, at least in the short term, that may indicate differences in social impact and potential susceptibility to cyberpower among nations with different age structures (most developing nations have younger age structures). More generally, it is reasonable to expect that groups of people will vary in their willingness and ability to enter cyberspace, and that may affect considerations of alternative modalities for conveying, or for interfering with, information and influence.[45]

Although the overall trend with much of the technology that makes the Internet work

is toward greater ease of use and even invisibility, disparities in people's ability to use technology imply differences in their ability to wield and to counter cyberpower. At one level, information can be hidden in cyberspace in so many ways that no special expertise may be required. At another, cat-and-mouse games between dissidents and censoring governments suggest that expertise may be important in some contexts. For example, sophisticated users may be able to establish a path (called a *tunnel*) to a different part of the Internet and send their packets across it so they appear to enter the Internet from this other location. This allows a user to experience what it would be like to use the Internet from a different region where there might be fewer controls on activity.

Implications of the Security Challenge and Responses

Success at improving the security properties of the Internet will determine the range of purposes for which it can be employed. If security problems continue to grow, the Internet's utility may be limited, or it may fracture into smaller, closed user groups focused on specific activities and interactions with specific parties. If we can materially improve the security of the Internet, however, it may become the primary platform for the delivery of critical information and services.

The need to improve the security of the Internet may change its nature in important ways. For example, the basic mechanisms of the Internet must take explicit account of the level of trust among the communicating parties and either permit them unrestricted efficient communication, if they trust each other, or provide protections and constraints where they lack mutual trust. We must allow trust to form part of the basis of collective action for defense against attack. To permit users to make trust decisions, they must know the identities of the parties with whom they are communicating; this means there must be explicit attention to the integration of identity mechanisms into the network. Security mechanisms must focus on the security of the information, not just the security of the conversation. We must, therefore, design a balanced set of mechanisms in the network and in the end node to deal with the reality that large operating systems will never be free of exploitable flaws.

Such security issues can be addressed in various ways that would have profoundly different effects on the balance between, at one end of the spectrum, freedom of action and unfettered exploitation of cyberspace by the users, and at the other, better ability to control hostile or unwelcome actions by other actors. One of the most fundamental questions in an analysis of cyberpower is to determine which point along this spectrum will be in the best interest of a nation.

The security of the Internet will be very different in 10 years, not because we get better at what we understand best today—mechanisms such as encryption, key management, and user sign-on—but because we become better able to understand and address the key paradox of securing the Internet: it is the act of communication that is risky, but communication is the whole goal of the Internet.

If we only communicated with those we fully trust, as did the Internet's developers and first users, the risk would be minimal. But today, much communication is with those we do not trust: we agree to accept mail from strangers, and must deal with spam and viruses; we connect to unfamiliar Web sites, and must deal with spyware and other malicious code. The Internet provides cost- effective, high-performance paths, which are great among trusting parties but also are great at delivering attacks and other unwelcome traffic. The security

community has traditionally focused on the endpoints of this spectrum: either complete trust or none. The middle condition—communication despite lack of full trust—represents the everyday norm on the Internet, as in real life, but has been the least studied.[46] If security is better in 10 years, it will be because the security community has developed a framework to understand how to deal with this middle case. Failure could lead to the realization of perils discussed for years—curtailment of participation by individuals and organizations, and balkanization as the Internet breaks into zones, each having its separate rules and mechanisms for safe interaction. Potential directions to avoid this outcome are outlined below; each would have its own implications for points of control and for cyberpower.

Trust-modulated Transparency

In the Internet as in real life, when we lack trust, we turn to constraints and trusted third parties to make the interaction safe. We use firewalls as a crude means to limit the modes of communication with "outsiders," we employ a trusted service to screen our mail for spam and viruses, and so on. But these approaches are not part of the Internet—they are afterthoughts, added piecemeal without an overall conception of what the result should be. In a future Internet, we should expect an architected solution. A *trust-modulated transparency* solution may involve a basic transport service that can morph between totally transparent and very constrained, depending on the degree of trust the endpoints have for each other. Unlike current approaches, this capability would involve support at the network level, as well as new mechanisms at the application level. It would require a relayering and reorganization of basic network functions.[47]

How fundamental the impact of trust-modulated transparency may be is illustrated by a reconsideration of connection initiation. The TCP was designed to establish a connection between two endpoints as efficiently as possible, where efficiency is largely measured in the number of round-trip delays that are required before data can be sent. But no endpoint identity validation is associated with this protocol: the TCP connection must be fully established before any identity credentials can be exchanged. Trust-modulated transparency argues for redesigning these sorts of protocols to achieve a preliminary phase for establishing identity and level of trust before the connection is completed. This redesign could allow for the "outsourcing" of that first phase to a remote machine whose sole job would be protection of the end node from untrusted or unwelcome connections.[48]

Identity Mechanisms

A future Internet would benefit from a rich set of identity mechanisms, because identity is a necessary basis for trust: if you do not know to whom you are talking, there is no way to determine a level of trust. At one end of the spectrum, a trusted third party provides the confirmation of identity and perhaps credentials of trustworthy status. This role might be played by government or reliable private sector players, as credit card companies do today for e-commerce. At the other end of the spectrum are identity schemes that are private to the parties in question: each party has a model of the identity of the others, but there are no third parties involved in confirming the nature of the identity.[49] These schemes are very different in their implications for the balancing of privacy and accountability and in the balance of power and control among the various actors.

Collective Action

A traditional mode of thinking about computer security considers each computer responsible for its own defense. But many actions, such as virus detection, spam filtering, or trust assessment, are best carried out in a collaborative or collective way; this implies creating and building on trust among the collaborating parties. Another motivation to make collective action trustworthy is the expectation of increasing collaboration for other purposes, such as use of grids for data storage or computation, or use of synthetic worlds. Collaboration will be part of the context for more and more activities and functionality.

Dealing with the Insecure End Node

Many of the immediate security problems that plague users today—viruses, spyware, zombies, botnets—are the result not of an insecure network, but of an insecure end node in an open network. A material improvement in overall security requires a holistic view of how end node and network security can be integrated. We can never assume that a system of the complexity of today's PC could be designed and implemented so it is free of flaws; we must assume that the end node will always offer some opportunities for a clever attacker. Those opportunities must be foreclosed through a combination of network controls and end node controls, which might include detecting unexpected and unfamiliar patterns of communication. Another approach to making end nodes more secure is to augment them with trusted components that provide some basic assured functions even if the larger system is compromised; this "innovation" actually harks back to the concept of a trusted computing base, developed by the Department of Defense in the 1970s.[50] A third approach is to virtualize the end node; this is similar to the approach of virtualizing the network discussed above. If the end node is operated as a number of virtual machines, each used for different purposes—the distinction might be as simple as between high security concerns and low security concerns—then a combination of functions in the end node and in trusted regions of the network might be able to materially improve the expected level of security of the end node.[51]

Securing the Information

Internet technology today focuses on protecting the integrity of information during its transfer. For example, the secure sockets layer encryption protocol provides disclosure control and integrity protection for data in transit, and also provides some assurance to the user that the server offering up the data is legitimate. But once the information has been received and stored on the receiving machine, there is no way to confirm where it came from or that it is legitimate. This pattern, adequate when the mode is predominantly direct transfer from source to destination, is unsustainable when the predominant modes are staged delivery, peer-to-peer redistribution, and retrieval of copies from archival storage (such as with Google). Trust in any of these modes to deliver legitimate information implies the ability to test the integrity of information, independently of how it was received.

Implications of Security Responses

At a superficial level, a more secure network is more stable and predictable, and when this works to the benefit of all parties, it enhances cyberpower. However, there are at least two considerations that provide a more nuanced view of the situation. First, not all parts of a design need to be standardized. Some parts must be standardized for the system to fulfill the basic need for interoperation and assurance, but beyond that point, the designers and standards writers often have latitude as to what they choose to standardize—what is always done the same way everywhere—and what they leave unspecified, so that it can be implemented differently in different places. A design that leaves more unspecified might open more opportunities for different countries to intervene in order to, for example, shape how their country's networks are designed and deployed.[52] Designs might vary in how much a country could exploit such variations.

Second, different designs will create different opportunities for various actors to exercise control and power over the system. The creation of new mechanisms often creates new opportunities for control, as well as new opportunities to fight over control.[53] Different designs can shift the balance of power among the citizen, the state, the Internet service provider (ISP), third-party actors, and others.

A crucial policy consideration is whether a nation's cyberpower is best served by giving power to its citizens—which may lead to more widespread use, cyber literacy, and innovation—or by shifting power to the state, which may foster some kinds of robustness, along with greater powers of policing and enforcement. The history of the Internet has generally featured shifts away from the state, but more recently there have been attempts to shift it back, combining basic technology elements with conditions imposed on providers or users.

Incentives and Investment

Some aspects of cyberspace, particularly networks, are capital-intensive. Since the U.S. view is that construction of network facilities is a private sector activity, the future shape of much of U.S. cyberspace will be a direct consequence of investment decisions by the private sector. In this section, we discuss factors that relate to the landscape of investment, starting with the pressures that might move the Internet and other key components of cyberspace from an open platform for innovation, toward a more closed and vertically integrated service platform. Then we note the young and highly dynamic nature of some key industry sectors, such as ISPs, which have existed for only about a decade. We examine the importance of advertising as a (perhaps underappreciated) driver of investment and capital and the interplay of regulation and private sector investment. Variations in these factors among different countries contribute to variations in their cyberpower.

One of the mantras of computing comes from technologist Alan Kay: "The best way to predict the future is to invent it." That sentiment has driven advances in computing devices and in applications for computing and networking. But invention is only one key to the future of networks: because they require the construction of physical facilities such as fiber optic transmission lines and radio towers, physical networks are capital-intensive. For the Internet, we might paraphrase Kay: "The best way to predict the future is to invest in it." The critical questions that will shape what we might see in 10 years include which players will build the future communications infrastructure; why they will invest and what they will build; and what will influence those decisions. A first pass at answers points to the obvious: in developed nations, the critical players are in the private sector; firms will invest to obtain a financial return; and

they will prefer the investment that yields the greatest return without excessive risk. How they calculate the balance of potential return and risk, and what returns and risks are realized, will define the infrastructure of the future.

We review here several questions that can help to frame the alternative futures. Will the Internet be open? Who is in the Internet industry? What is the role of advertising? Will the Internet be free of regulation? What are the implications of economic factors?

Will the Network Be Open?

One of the defining characteristics of the Internet is that it is *open* in several key ways. It is an open platform: any innovator can offer a product or service that runs on top of it (see figure 8–1). It is based on open standards that are free of intellectual property limitations or licensing requirements. It therefore sustains, other things being equal, an open market: new entrants can offer Internet services alongside established players. This openness, fostered by the software that defines the Internet, contrasts with the traditional, vertically integrated model of communication networks in which the investor in infrastructure recovered its investment by selling services over that infrastructure. For example, the telephone company recovered its costs in the physical assets by selling telephone service; the cable television industry recovered its costs by selling entertainment programs. An ISP charges for access to the Internet but does not currently make additional money when a user accesses the Web, watches a video, or makes a telephone call over the Internet.

There is power in the Internet's open nature: it has stimulated all sorts of third-party innovation by the likes of Yahoo!, Amazon, eBay, Google, Facebook, and YouTube, which increase both the total value of the Internet to its users and the number of users that may be interested in exploring the Internet. Inevitably, however, the demand for the Internet will mature and growth will slow, reflecting the typical *S*-shaped adoption curve for new technologies.[54] The open paradigm might continue to prevail, but another outcome—signaled today by the debate over "network neutrality"—is that ISPs might move in a direction of greater control of the services offered over their networks, in an attempt to make more money by treating different kinds of users and uses differently. This prospect has aroused concern about chilling effects: on innovation in the uses of the Internet, on speech and culture, and on the economy, which has fed off the Internet in many ways.[55]

If ISPs raise costs for intensive and extensive users, such as major content providers, creation of new applications by smaller players unable to afford to pay for good access might be inhibited; this is likely to reduce innovation by such actors and decrease the attraction of the Internet for those who would use their potential products, diminishing growth in Internet use. If ISPs move toward favoring their own content and applications over those of third parties, resultant reductions in competition in the supply of content and applications—a kind of return to the pre-Internet status quo—might reduce innovation in those areas and on the part of those who would use them; this, too, could stifle growth in Internet use.[56] If ISPs eschew involvement with content and applications but, facing a commodity business with low margins, do not see a case for investment in upgrades to their facilities, there might be stagnation in the capabilities or scale of the Internet, again diminishing growth in its use. Avoiding these perilous paths implies that ISPs must be able to make enough money to continue investing.

We do not address these issues in detail, but as the kinds of networks associated with the Internet change, questions about openness can be raised more broadly. Sensor networks, for

example, will present opportunities to promote or constrain openness, as will networks for such specific environments as motor vehicles.

Who Is in the Internet Industry?

Understanding Internet industry economic scenarios is confounded by the dynamic nature of the ISP industry. We associate delivery of the Internet with ISPs, but that industry is not much over a decade old (it came together with the commercialization of the National Science Foundation [NSF] backbone in 1995). Within the last decade, there has been both consolidation and realignment, not only among ISPs proper, but between ISPs and more conventional telecommunications providers such as telephone and cable companies. These moves have been aided by transitions in the technology bases of those companies, such as the diffusion of IP within the telephony infrastructure and of optical fiber in telephone distribution facilities. Among other elements of convergence, there has even been consolidation between ISPs and content providers.[57] Meanwhile, a new category of player has emerged: the overlay network of servers that support a class of applications. Interposing a service between users and ISPs, they do not conform to expectations for the operation of Internet applications and the kinds of businesses that were originally differentiated by the Internet architecture layer in which they operate.

Akamai, a third-party provider of content distribution, provides an example of an overlay network.[58] It manages a set of caches that distribute Web content globally and expedite its delivery to dispersed users: users receive content from caches that are closer to them than the servers operated directly by the source may be. Akamai recognized early on that major content producers are a special kind of user that the company could serve by centralizing and packaging the task of coordinating with multiple ISPs on cache location and service quality, allowing the content producers to offer better and more uniform content delivery to their ultimate users. In identifying this service to offer to potential customers, Akamai created a new way to make money; such a model may be influencing current ISP interest in new terms for doing business with content providers, because it amplified the business potential of Internet content. It also illustrates the need for caution about defining the industry.

What Is the Role of Advertising?

The importance of advertising in the future of the Internet may not yet be adequately recognized. Today, advertising funds most of the "free" Web sites on the Internet. It is the engine that funds Google, among others. The shift of advertising revenues away from print and other traditional channels to the Internet is causing major disruption to those channels. Advertising is also embedded in new ways, such as the mounting of attractive sites directly by advertisers.[59]

The increasing importance of advertising, in terms of both revenues and clout, will become obvious as the television experience moves to the Internet. The production of almost all television content (with the exception of content for premium channels such as Home Box Office) is funded by ad revenues. If advertisers lose confidence that television advertising is effective, it could destabilize the whole supply chain for entertainment content. Today, the TiVo digital video recorder (DVR) is a threat, but the Internet could be an even more serious threat, as content becomes something the consumer can watch at will and even modify.[60] Content producers are likely to respond with technical changes to try to produce a more stable

"advertising contract;" current disputes over the protection of digital rights will pale in comparison to the fights over whether the consumer can be forced to watch an ad.

This issue is potentially of broader concern, because it may influence such basic questions as who owns the servers that provide content, and how that content is managed. TiVo provides a specific example drawing on today's experience. It is a popular device that many consumers choose to install. They pay for it, and then they own and control it (to some extent). But in 10 years, who will control DVRs? It might be the advertisers, who would give the consumer the DVR for free, and then give that user access to any content desired under a wholesale agreement (somewhat like that offered by today's cable providers), assuming they can insert the ads targeted to that viewer. Technically, one way to enforce this contract would be to provide a remote control with the DVR that has been modified so that even if you switch channels, you still have to watch the ad.[61]

The question of who owns and controls the device in your home that acts as the source of the content you can see is a question that may be critical to an analysis of cyberpower. This is a variation of the classical set of questions raised in media policy discussions: media and telecommunications analysts worry about concentration of media ownership and control of information to the consumer. Will the advertisers end up in charge of what we can see to a greater extent than today? If so, how can they contribute to defining cyberpower?

Will the Internet Be Free of Regulation?

As framed, the question of whether the Internet will be free of regulation has an obvious answer: it is not free of regulation today, never has been, and never will be. But regulatory constraints directly on the Internet have been minimal compared to those on traditional telephone companies, broadcasters, and cable television providers; and they have tended to focus on applications and consumer protection rather than the structure or competitive conduct of the industry.[62] As more traditional players have gotten more involved in the Internet, they have brought some of their regulatory burdens with them.[63]

In 2006, Congress considered telecommunications reform legislation; the proposals point to the potential for new rules. In particular, proposals for some form of "network neutrality" protection could constrain ISP pricing schemes and investment behavior. In the wake of 1996 legislation aimed at reforming telecommunications regulation,[64] the regulatory battle in the United States focused on "unbundling," requiring telephone facilities owners to lease local- loop capacity for resale to service competitors in order to increase the number of competing service providers. This approach was challenged legally and abandoned in the United States, although it remains popular in other developed countries, especially in Europe. Where there is regulation, there typically is some expectation of formal observation or monitoring by the government (and by competitors and civil society organizations); this in itself is one reason why regulation of the Internet has been resisted.

There are other categories of regulation that bear more directly on government interests and may shape the Internet and its uses in different ways. One is the provision for law enforcement access to communications. Government can gain access (under subpoena in the United States) to usage data from ISPs; in addition, there has been movement to build in to core Internet systems— routers—the conceptual equivalent of wiretapping. This was mandated in the United States by the Federal Communications Assistance to Law Enforcement Act (CALEA), which was designed originally for telephone networks and is being extended to the

Internet and to wireless communications systems.[65] An unintended consequence of adding this kind of capability is that, once available, it could be exploited by others—criminals, for example—and not just by the law enforcement agencies that were its intended beneficiaries.[66]

Another government-focused category of regulation relates to critical infrastructure and involves both government-industry cooperation and investment. Arrangements for national security and emergency preparedness (NS/EP) in telecommunications arose historically in the context of the telephone industry, and the application of these arrangements to the Internet has been the subject of exploration in the United States for about a decade. A related and parallel exploration has focused on critical infrastructure protection, which involves monitoring for potential problems, sharing of related information, and implementing protection mechanisms.[67]

The ISPs are unpredictable and potentially inconsistent players in critical infrastructure protection. At issue is their willingness and ability to invest in such approaches to protection as physical facility diversity (redundant investments).[68] Also problematic is how to coordinate and allocate planning and investment within the industry and between industry and government. In the 1960s, when NS/EP mechanisms were introduced in the United States, the telecommunications industry centered squarely on one company, American Telephone and Telegraph (AT&T). By contrast, the Internet involves a wider variety of players, most of which have little history of cooperation with the government; it also involves more instability in the market, which complicates understanding of the elements of the infrastructure and identifying who is responsible for them. The controversy beginning in mid-2006 over secret U.S. Government interception of communications and the associated backlash do not bode well for future government-industry cooperation. Were there to be mandated investment in physical robustness (as CALEA compliance was mandated), any increase in costs would affect all players in a geographic region; if not offset by other factors, such an increase could have an impact on the level or growth of use in that region, and therefore on how Internet conditions in that region might compare to those elsewhere.

A final category of regulation may emerge from the international debate about Internet governance that was prominent at the 2005 World Summit on the Information Society and the convening of the Internet Governance Forum beginning in fall 2006. So far, these activities have provided forums for discussion and debate, but there has been no movement to decision or action, in part because of the difficulty of achieving consensus among large numbers of nations and different components of national governments.[69]

Implications of Economic Factors

The Internet marketplace supports much variation, even though the providers implement common standards. Issues include the degree of openness to innovation, the degree of penetration of access (both dial-up and broadband), the extent of planning and investment in resilience and diversity of infrastructure, and perhaps the degree of state and private sector observation of ongoing activities as part of regulation. Many of these issues play out differently in different countries, leading to variation in the issues that feed cyberpower.

Several factors can shift the balance. More investment in technology and operation for resilience might increase cost and reduce penetration. A shift toward more control and monitoring might lead to a network that is easier to police but might chill innovation and emerging uses. The United States must try to determine which would be a better contributor to cyberpower. It must also decide whether to do more to encourage certain sorts of investment

overseas (for example, in the developing world) or to drive faster penetration of services such as the Internet, and if doing either, it must figure out what sort of network to encourage and what its features should be in order to protect U.S. interests.

Research and Reinvention

We have looked at technical issues that might be influenced by attention from the research community. This section looks more generally at the institution of research: sources of funding, the health of the research establishment, and differences in policy and practice between the United States and other countries. At the moment, research in the United States seems disorganized compared to some other countries, which could mean a loss of U.S. economic advantage, a loss of control over the shape of cyberspace, or even loss of a seat at the table to influence the factors that shape cyberpower.

Research remains essential for the vitality of the Internet, and renewed commitment by the U.S. Government may be important to future developments. The expectation that research on the Internet will continue adds to the unpredictability of the Internet's future: research generates new options for users, for ISPs, and for a wide range of third parties in both the public and private sectors. Maturation of some aspects of the Internet makes it more challenging both to frame research and to assimilate research results. The increasing presence of industry in Internet research highlights the important role that has been played by government-funded academic basic research: such research is free to address issues not likely to be pursued by applied research in industry. An obvious example is options for maintaining openness.[70]

It is well known that the Internet and other underpinnings of cyberspace arose from research programs. However, the conditions for funding, framing, and conducting relevant research have changed considerably. The U.S. situation is characterized by significant limitations. Historic industrial research giants such as AT&T and International Business Machines have become less research-active at best, without being replaced by healthier firms (Cisco might be a candidate). The historic lead funding entity in government for large-scale computing systems—DARPA—has changed its focus. The academic research community has experienced a series of losses in personnel due to funding fluctuations and the lure of industry and focus away from basic research toward research that complements, even if it does not lead, a fast-moving marketplace.[71] Other countries, meanwhile, have increased levels of activity, support, and coordination, responding to the challenge posed by the commercial success of the United States and also to concerns about its cultural hegemony, fostered by the global reach of the Internet.[72] However, there is no evidence that activities overseas are contributing to revolutionary change in ways different from what can be seen in the United States (whose researchers, of course, also participate in international exchanges).

The U.S. National Science Foundation is mounting a pair of initiatives that hold the potential to influence Internet technology in an organized way. One, called Future Internet Design (FIND), has begun to fund research that is explicitly motivated by a vision of a future network with features radically different than ones that would result simply from incremental modification or improvement. FIND was framed with an eye toward addressing some of the problems outlined in this chapter, from deficiencies in security and trustworthiness to the economic viability concerns of ISPs. Many of the revolutionary architectural ideas described in this chapter are being explored within the NSF FIND initiative.[73]

The Global Environment for Network Innovation (GENI)[74] is an infrastructure project to complement the research funded in FIND. GENI is intended to apply ideas gleaned from FIND and elsewhere in real-world tests. It will support multiple simultaneous, independent experiments with new architectural, service, and application concepts at a large scale and with live user communities. The GENI concept recognizes that practical testing is needed not only to achieve the ambitious goals of significant changes in architecture but also to demonstrate to maturing markets what new ideas can really mean in practice. An issue being discussed that must be worked out, given that the Internet has been international from its early days, is how to achieve international participation in GENI and FIND.

Goal-directed Research: Priorities for the Future

Different funding agencies have different philosophies of research—fundamental or mission oriented, long term or short term, and so on. Currently, the only U.S. funding agency that has declared its intention to give priority to research on the future of cyberspace in general and on the Internet in particular is the NSF, an agency that traditionally has focused on long-term, fundamental research. Exploration in a number of other countries is shaped with the specific objective of giving the host nation's industry a strategic advantage in the cyberspace marketplace or creating a set of standards that will shape the future of cyberspace in a particular way. We have mentioned a number of issues ripe for research, some in detail and some in passing. There are many examples:

- *Security*. Will we have an "architecture for security" in a future Internet, or will we instead continue with a mix of point solutions to problems?
- *Object provenance*. Will we have some sort of trustworthy metadata attached to an object that provides assurance as to its origin and provenance?
- *Identity*. Will we have an architected framework to deal with the issues of identity that are going to arise in cyberspace?
- *Location-aware computing*. Will we have a coherent and unified framework to associate location information with people, devices, and objects, to facilitate location-aware search and notification?
- *Location-sensing*. The global positioning system now allows any suitably equipped device to know where it is, but only if it can receive the radio signals from the satellites. Should we work toward a next-generation location-sensing framework that allows devices to determine their location inside buildings and other places where the satellite signal does not reach?
- *Open sensor network*. Will we have an available, open, ubiquitous network suited to the needs of inexpensive sensors and embedded computers, or will sensors require the installation of a custom network?
- *Open vehicle network*. Will future automobiles offer an open network, along the lines of the Internet, or will they be closed networks, with manufacturer- provided services only?
- *Networks in times of crisis*. Will we make the necessary provisions to assure that the Internet will be available and useful even in times of disaster?

These questions and others represent forks in the road to the future. Industry may or may not decide to concentrate on some of these issues. Public sector funding agencies will also make decisions as to which of these objectives are important enough for funding and attention,

and which path we should take when we reach the fork. Different countries may set different priorities and work toward different outcomes around these issues. Whatever choices are made, these issues will have a profound effect on the relative balance of cyberpower in 10 years.

---

[1] See Steven Cherry, "Winner: Nothing but Net," *IEEE Spectrum* (January 2007), available at <www.spectrum.ieee.org/jan07/4831>; and Paul V. Mockapetris, "Telephony's Next Act," *IEEE Spectrum* (April 2006), available at <www.spectrum.ieee.org/apr06/3204>.

[2] See the discussion in chapter 6, "Evolutionary Trends in Cyberspace."

[3] Two proposals currently funded by the National Science Foundation that explore this concept are Jon Turner et al., "An Architecture for a Diversified Internet," National Science Foundation grant CNS-0626661, accessed at <www.nets-find.net/ DiversifiedInternet.php>; and Nick Feamster, Jennifer Rexford, and Lixin Gao, "CABO, Concurrent Architectures are Better than One," National Science Foundation Networking Technology and System grant CNS-0626771, accessed at <www.nets-find. net/Cabo.php>.

[4] See, for example, the Defense Advanced Research Projects Agency (DARPA) Networking in Extreme Environments Program, available at <www.darpa.mil/sto/strategic/netex. html>; and the DARPA Next Generation Program, available at <www.darpa.mil/sto/ smallunitops/xg.html>.

[5] See, for example, the Gumstix computers at <http://gumstix.com/platforms.html> or the iButton computer at <www.maxim-ic.com/auto_info.cfm>.

[6] Classical supercomputing combines specialized and typically custom-built hardware and software, including algorithms that may be associated with complex models of the physical world. See Computer Science and Telecommunications Board, *Getting Up to Speed: The Future of Supercomputing* (Washington, DC: National Academies Press, 2005).

[7] Examples of volunteer activities include SETI@home (searching for extraterrestrial radio signals), available at <http://setiathome.berkeley.edu/>, and Folding@home (research into protein folding), available at <http://folding.stanford.edu/>.

[8] Botnets involve the surreptitious placement of malicious software into personal computers (PCs), which are then mobilized into "nets" to perform attacks. The "@home" research projects suggest the potential for legitimate commercial botnets. This also raises questions about what might happen if PC owners willingly rent out their excess computational capacity. Such a practice might have implications for network capacity planning as well as liability and other practical issues.

[9] A commodity server in this context refers to a mass-produced commercial computer, in contrast to a special-purpose machine.

[10] Google is rumored to have almost a half-million commodity servers in operation. Rick Rashid of Microsoft has suggested that 75 percent of all servers in the United States are operated by just three companies: Google, Yahoo!, and Microsoft. Rick Rashid, remarks at 20th Anniversary Symposium of the Computer Science and Telecommunications Board, October 17, 2006.

[11] There are also longer term trends, such as quantum computing, but we believe these will not be relevant to a discussion in the next decade.

[12] Google has multiple server sites, but any one query is processed at one physical location. In that respect, the processing is centralized.

[13] See <http://robotics.eecs.berkeley.edu/~pister/SmartDust/>.

[14] See, for example, UCLA's Center for Embedded Networked Sensing, a National Science Foundation Science and Technology Center, at <http://research.cens.ucla.edu/>.

[15] See, for example, U.S. Department of Transportation Intelligent Transport Systems Joint Project, available at <www.its.dot.gov/index.htm>. An overview can be found in Jonathan Fahey, "Car Talk," *Forbes*, January 29, 2007, 52–54.

[16] The Center for Embedded Network Sensing at UCLA has a new program on "participatory sensing" as part of its efforts in urban sensing. See <http://research.cens. ucla.edu/projects/2007/Urban_Sensing/Applications/>.

[17] Today, except for a few consumer items such as motion-sensing lights, sensors are in the hobbyist category.

[18] For example, cell phone encryption is known to have limitations; WiFi deployments often lack any encryption, leaving them open to "wardriving" as people in cars search out networks to exploit (see <http://www.wardriving.com/>); and Bluetooth interception has also become recreation for the mischievous; see <http://seclists.org/lists/isn/2005/ Feb/0085.html>.

[19] See Weather Underground at <www.wunderground.com/>.
[20] See, for example, the National Ecological Observatory Network at <www.neoninc.org/> and Earthscope at <www.earthscope.org/> in the United States.
[21] A mashup combines data or functions from multiple sources into a single integrated tool; a mashup might add location information from Google Maps (<http://maps.google. com/maps?hl=en&tab=wl>) to real estate data from Craigslist (<www.craigslist.org/ about/sites.html>) to create a service not available from either by itself.
[22] This issue has been discussed and debated by observers of the social context of information. See Computer Science and Telecommunications Board, *Global Networks and Local Values* (Washington, DC: National Academies Press, 2001).
[23] Today's proliferation of blogs and communal sites for sharing video clips, photos, and so on points to a seeming flowering of creative output led by individuals; this activity, in turn, has generated additional growth in associated business data.
[24] See Wikipedia at <http://wikipedia.org/>.
[25] See, for example, Jaron Lanier, "Digital Maoism: The Hazards of the New Online Collectivism," and the responses the article generated, available at <www.edge.org/3rd_ culture/lanier06/lanier06_index.html>.
[26] For example, the OptIPuter project, "named for its use of optical networking, computer storage, processing and visualization technologies," is meant "to enable collaborating scientists to interactively explore massive amounts of previously uncorrelated data." The hope is that: the OptIPuter, when linked with remote "data generators," whether the TeraGrid, instrumentation, or data storage devices, will prove to be an enabling technology for large-scale networked science facilities, as well as for broader societal needs, including emergency response, homeland security, health services and science education. The TeraGrid comprises "resources at eleven partner sites to create an integrated, persistent computational resource. . . . Using high-performance network connections, the TeraGrid integrates high-performance computers, data resources and tools, and high-end experimental facilities around the country." See Faith Singer-Villalobos, "The OptIPuter: 21st-century E-science," available at <www.teragrid.org/news/news06/tg06_opti.html>.
[27] See, for example, Department of State press release, "Secretary of State Establishes New Global Internet Freedom Task Force," available at <www.state.gov/r/pa/prs/ ps/2006/61156.htm>.
[28] See, for example, W3C Semantic Web site, available at <www.w3.org/2001/sw/>.
[29] It is possible to tag a site in a way that diverts the crawlers used by search systems, and it is possible for crawlers to be programmed to ignore those tags. These possibilities create a space of contention for control surrounding search.
[30] See "EU: Quaero—a European multimedia search engine project to rival world leaders in Internet search," accessed at <http://ec.europa.eu/idabc/en/document/5554/194>. For a skeptical assessment of its prospects, see Philip E. Ross, "Loser: What's the Latin for 'Delusional'?" *IEEE Spectrum* (January 2007), available at <www.spectrum.ieee.org/ jan07/4842>.
[31] A prominent example is Baidu, which in early 2008 expanded its operations to Japan, signaling international ambitions. See <http://ir.baidu.com/phoenix.zhtml?c=1884 88&p=irol-homeprofile>.
[32] See K. Sollins, ed., "Architectural Principles of Uniform Resource Name Resolution," RFC 2276, Internet Engineering Task Force, 1998. For a discussion of confusion over terms, see "URIs, URLs, and URNs: Clarifications and Recommendations 1.0," available at <www.w3.org/TR/uri-clarification/>.
[33] See Corporation for National Research Initiatives, "Handle System: Unique Persistent Identifiers for Internet Resources," available at <www.handle.net/>.
[34] See, for example, Bryan Ford et al., "User-relative Names for Globally Connected Personal Devices," available at <http://publications.csail.mit.edu/abstracts/abstracts06/baford/ baford.html>.
[35] The TIA program (Total [or Terrorism] Information Awareness) begun at the Defense Advanced Research Projects Agency in 2002 explored ways to draw from multiple information stores on an ad hoc basis; some of that work continues under various auspices.
[36] This function is often provided by third-party providers such as Akamai, discussed below.
[37] The term *cloud computing* is being used to describe the idea of commodity computing service available within the Internet.
[38] The Open Net Initiative at <www.opennetinitiative.org/> is one of several organizations tracking current content controls on the Internet.
[39] See examples noted in Reporters without Borders, *Handbook for Bloggers and Cyber-Dissidents*, September 14, 2005, available at <www.rsf.org/rubrique.php3?id_rubrique=542>.

[40] Applications being explored range from outer space to the battlefield. For pointers to current research, see <www.dtnrg.org/wiki> and <www.darpa.mil/sto/solicitations/ DTN/>.
[41] One illustration is "a camera phone visual tag reader" that acts as "the glue that connects the two halves of your application, the physical world deployment and the virtual world application software." See <http://semacode.com>.
[42] See, for example, Second Life at <http://secondlife.com/>; Edward Castronova, *Synthetic Worlds: The Business and Culture of Online Games* (Chicago: The University of Chicago Press, 2006). For commentary and critique, see Jim Giles, "Life's a Game," *Nature*, 445 (January 4, 2007), 18–20.
[43] These issues were discussed by Jon Kleinberg and, to some extent, Richard Karp at the 20th Anniversary Symposium of the Computer Science and Telecommunications Board, October 17, 2006.
[44] Much attention has been given to the $100 laptop "one laptop per child" project (see <www.laptop.org/>) and, recently, Intel's similar effort, although some argue that for the foreseeable future, approaches that capitalize on the fact that in some cultures there is a greater level of comfort with shared technology, such as cell phones, may be more realistic.
[45] Of course, it can be argued that not becoming overly dependent on cyberspace might convey a different kind of resourcefulness, one that may be important in the event of disruption of access.
[46] Recent research programs have focused on making information technology more resilient: assuming failure of prevention (itself also a research focus), improved technology would better detect security problems, resist damage from them, continue to operate at some level despite damage, and even recover from them or regenerate damaged components.
[47] See the final report of the Future-Generation Internet Architecture Project (NewArch) at <www.isi.edu/newarch/>.
[48] This mechanism, once designed and deployed, might be usable by many actors for many purposes.
[49] See, for example, P. Gutmann, "Simplifying public key management," *IEEE Computer* 37, no. 2 (February 2004).
[50] U.S. Department of Defense, *5200.28-STD Trusted Computer System Evaluation Criteria* (December 1985), also known as the Orange Book.
[51] For one speculation on this approach, see Butler Lampson, "Accountability and Freedom," available at <http://research.microsoft.com/~risaacs/blampson.ppt>.
[52] See the discussion of pressures toward such outcomes in chapter 21 of this volume, "Internet Governance."
[53] The standards-development process displays such rivalry regularly. Disputes over the open systems interconnection versus the transmission control protocol/Internet protocol suites and over standards for wireless communication, such as China's recent attempt to promote indigenous technology, have had international dimensions reflecting concerns ranging from classical competitiveness to cyberpower.
[54] Korea has already experienced the challenge of slowing revenue flow due to market saturation, leading providers to consider usage-based pricing or other steps to preserve revenue flow. See Broadband Working Group, MIT Communications Futures Program, "The Broadband Incentive Problem," September 2005, available at <http://cfb.mit.edu/ groups/broadband/docs/2005/Incentive_Whitepaper_09-28-05.pdf>.
[55] There are different degrees of openness in cell phone systems and conflicting trends. Different countries have more differences in cell phone service than in Internet service.
[56] Constraints on innovation may promote monocultures, themselves a source of vulnerability.
[57] The apparent failure of the much-heralded AOL–Time Warner merger provides but one cautionary tale about the uncertainties of predicting industry trends, which reflect, among other things, technology change, consequent shifts in buyer behavior and the costs of doing business, and what the government does and does not do.
[58] See <www.akamai.com/>.
[59] An example of such a site is <http://mycoke.com>.
[60] See <www.tivo.com>.
[61] The DVR situation is an instance of a larger battle associated with digital rights management: content generators have used the threat of withholding contracts and other legal maneuvers to induce consumer electronics manufacturers to produce devices that assist in the protection of content, although such devices also limit access and use of content that has fewer or no protections. See Julie E. Cohen, "Normal Discipline in the Age of Crisis," Georgetown University Law Center, Public Law and Legal Theory Research Paper No. 572486, August 4, 2004.
[62] Growth in the Internet marketplace has attracted government attention to competitive conduct there, as evidenced by attention to antitrust issues on the occasion of certain mergers or acquisitions.
[63] The International Telecommunication Union's promotion of "Next Generation Network" standards suggests an

effort that combines traditional players and governments in promoting seemingly benign customary features such as priority access in emergencies, as well as law enforcement access (wiretapping) that may have broader implications. See <www.itu.int/ITU-T/ngn/index.phtml>.

[64] See Federal Communications Commission, The Telecommunications Act of 1996, available at <www.fcc.gov/telecom.html>.

[65] See <www.fcc.gov/calea/> and <www.cdt.org/digi_tele/>.

[66] Similar kinds of concerns have been raised even for seemingly benign applications: the development of the Platform for Internet Content Selection (PICS) Web standard was more or less derailed when critics expressed concern that PICS could be used by governments to achieve censorship.

[67] See, in this volume, chapter 5, "Cyberspace and Infrastructure," and chapter 23, "Cyber- power and Critical Information Protection: A Critical Assessment of Federal Efforts."

[68] The effects of the 9/11 attacks on conventional telecommunications and the Internet reinforced the value of diversity, among other things, and should have boosted Internet service provider interest in this subject. See Computer Science and Telecommunications Board, *The Internet Under Crisis Conditions: Learning from September 11* (Washington, DC: National Academies Press, 2002).

[69] Such decisions as have been made are limited to subjects such as articulation of principles and definitions (for example, of Internet governance).

[70] See Computer Science and Telecommunications Board, *Broadband: Bringing Home the Bits* (Washington, DC: National Academies Press, 2002).

[71] See Computer Science and Telecommunications Board, *Renewing U.S. Telecommunications Research* (Washington, DC: National Academies Press, 2006).

[72] See, for example, <http://europa.eu.int/information_societv/research/index_en.htm>.

[73] See <www.nets-find.net/>.

[74] See Global Environment for Network Initiatives Web site at <www.geni.net/>.