

BROADBAND

BRINGING HOME THE BITS

Committee on Broadband Last Mile Technology
Computer Science and Telecommunications Board
Division on Engineering and Physical Sciences
National Research Council

NATIONAL ACADEMY PRESS
Washington, D.C.

NATIONAL ACADEMY PRESS 2101 Constitution Avenue, NW Washington, DC 20418

NOTICE: The project that is the subject of this report was approved by the Governing Board of the National Research Council, whose members are drawn from the councils of the National Academy of Sciences, the National Academy of Engineering, and the Institute of Medicine. The members of the committee responsible for the report were chosen for their special competences and with regard for appropriate balance.

The majority of the support for this project was provided by the Defense Advanced Research Projects Agency under contract No. N00174-99-C-0052 and the National Science Foundation under grant No. ANI-9908155. Additional support was provided by the Association for Computing Machinery's Special Interest Group on Data Communication, Hewlett-Packard, Intel Corporation, Interval Research Corporation, WorldCom, Sun Microsystems, Texas Instruments, and Qwest. Any opinions, findings, conclusions, or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the sponsors.

International Standard Book Number 0-309-08273-0

Additional copies of this report are available from:

National Academy Press
2101 Constitution Ave., N.W.
Box 285
Washington, DC 20418
800-624-6242
202-334-3313 (in the Washington metropolitan area)
<http://www.nap.edu>

Copyright 2002 by the National Academy of Sciences. All rights reserved.

Printed in the United States of America

General Library System
University of Wisconsin - Madison
728 State Street
Madison, WI 53706-1494
U.S.A.

What Is Broadband?

WHY DEFINE "BROADBAND"?

The term "broadband" has become commonplace for describing the future of digital communications. It is widely used to refer to a range of technologies being offered or developed for the delivery of data communications services. Broadband refers most commonly to a new generation of high-speed transmission services aimed at residential and small business users. On its face, the term refers to the substantial bandwidth that a high-speed connection can provide to a user.¹ But is there a well-defined threshold that marks the boundary between broadband and narrowband services above which one can say a service is broadband? This chapter explores this and other dimensions of how one defines and characterizes broadband.

There is an obvious lower bound of broadband—that is, whether the service offers higher capacity than dial-up access (which is limited to 56 kilobits per second [kbps] per phone line) or even ISDN service (the basic service offers two 64-kbps channels; experienced service can be at 128 kbps or higher), which was introduced with the promise of providing higher-speed service but was never widely adopted in the U.S. residential market. The speeds offered by broadband local access services, such as cable modem and DSL, generally start well above this threshold (in the

¹Greater communications capacity translates into the ability to deliver a given amount of information faster, so "speed" is often used synonymously with "capacity" in this context.

hundreds of kilobits per second effective bandwidth or better), but they span a wide range of speeds, with consequences for the types of applications they are able to support. A service may, for example, be fast enough to support rapid Web browsing or a few channels of telephony, but too slow to support even a single TV-quality video stream.

Various groups have struggled to develop appropriate definitions of broadband, and these definitions have changed over time. In the 1980s and early 1990s, broadband referred to rates greater than 45 megabits per second (Mbps), and “wideband” referred to rates between 1.5 and 45 Mbps. Then, circa 1995, broadband commonly referred to anything 1.5 Mbps and higher in most circles; thus, it was an order of magnitude greater in capacity than ISDN service.² Mandated to report to Congress on deployment, the FCC, in its 2000 report on broadband deployment,³ defined an “advanced telecommunications service” as one that is at least 200 kbps in each direction. The speed and two-way requirement attempted to capture the intent expressed in the Telecommunications Act of 1996—that the speeds of what the act terms “advanced telecommunications services” should exceed the rates offered by the technologies available to residential customers at the time of the act’s passage.⁴ At the time of the act’s passage, residential customers were generally limited to dial-up service (then typically no more than 33.6 kbps). Hinting that a definition should also be, at least in part, driven by the requirements of user applications, the FCC report also observed that the 200-kbps threshold it selected is roughly the threshold above which the time it takes to load a Web page becomes comparable to the time it takes to turn the page of a book. But a 200-kbps service would be inadequate to support even a single TV-quality video stream to each house, let alone the multiple such streams that a family might reasonably use, which would require multiple megabits per second. Nor do TV-quality video streams represent the most bandwidth-demanding application one could imagine in the future.

These definitional questions also arise in an international context as other countries explore policy options concerning broadband. For ex-

²For a historical view of prospective services, see *IEEE Network* special issue on the North Carolina Information Highway, November/December 8(6), 1994.

³Federal Communications Commission. 2000. *Inquiry Concerning the Deployment of Advanced Telecommunications Capability to All Americans in a Reasonable and Timely Fashion, and Possible Steps to Accelerate Such Deployment Pursuant to Section 706 of the Telecommunications Act of 1996* (Second 706 report). CC Docket No. 98-146, Second Report, FCC 0-290 (rel. August 21). Available online at <http://www.fcc.gov/Bureaus/Common_Carrier/Orders/2000/fcc00290.pdf>.

⁴The FCC report defines a broader class of services—“high-speed services,” defined as services that exceed 200 kbps in at least one direction—of which advanced telecommunications services are a subset.

ample, two recent Swedish governmental commissions took a look at this same question (but without having to factor in the U.S. statutory language, as did the FCC) and adopted substantially higher-speed thresholds: symmetric 2 Mbps and 5 Mbps.⁵ In the end, neither the definitions of the FCC or those of the Swedish commissions are entirely satisfactory (indeed, FCC staff speaking at the committee's 2000 workshop acknowledged difficulties associated with the FCC definition).

Defining broadband is more than an academic exercise. Numerous groups would stand to benefit from workable definitions of what constitutes broadband. They include:

- *Consumers*, who would like to be able to evaluate service offerings to see if the offerings are likely to meet their needs;
- *Service providers*, who would like to develop, invest in, and deploy services that consumers will need and want;
- *Application and content developers*, who would like to understand and track the connectivity performance options available to consumers;
- *Policy makers or regulators*, who seek to monitor broadband service deployment and measure the impact of policy or regulatory decisions on deployment, define the characteristics of services eligible for tax credits or loans, or define the characteristics of services required in build-out commitments associated with regulatory relief; and
- *Public interest groups*, which seek to evaluate capabilities available to consumers and to understand the implications of alternative policy approaches that influence those capabilities.

Framed in this way, defining the term "broadband" in some sense also involves (1) identifying the kinds of applications that consumers are likely to find useful and desirable and (2) determining the benefits that different segments of the public anticipate from access to broadband services. The definition of broadband used by each of these groups will reflect that group's expectations and, consequently, can have a significant effect on decision making. Too limited a definition, such as establishing too low a data transmission rate as the broadband threshold, could result in a mismatch between expectations and capabilities, while a definition that is unrealistic in terms of technological capabilities, costs, or consumer demand could prompt inappropriate or poorly aimed policy interventions. The absence of a consensus on definitions will confuse political

⁵Swedish Special Infrastructure Commission (June 1999): Broadband should be defined as at least 2 Mbps (symmetrical) to the user. Swedish IT Commission (November 1999): Minimum 5 Mbps to the user.

debate on the subject and require ongoing debates about what definitions to use.

OVERVIEW OF THE TECHNICAL CHARACTERISTICS OF BROADBAND

Communications capacity, or speed, is only one of a set of performance characteristics of a service. That it is not the whole picture is easily seen in the contrast between dial-up access, where the modem must place a telephone call and negotiate a connection with the ISP's modem, and the services available today that are generally considered broadband—which frequently offer “always-on” connectivity as well as high speed. Along with speed and always-on are additional parameters such as bandwidth symmetry and addressability that are important components of a definition of broadband. Each of these is considered in the sections that follow.

Speed

The speed or bandwidth of a service—the rate at which one can transfer data to and/or from the home—is a function of multiple factors. Because the effective bandwidth reflects the capacity of the end-to-end connection between sender and receiver, the speed seen by a user can be constrained at any one of a number of points between the user's computer and the computer providing a particular service. However, speeds within the core network have been rising, at least in the United States and other developed nations, and the capacity of the network link between the user and the broadband provider's network is one of the crucial factors that determines how the broadband service can be used.

The better-than-dial-up criterion for broadband assures that a service is at least a little better than what was available before, but it does not address the question of whether the service is good enough. And while a 2- or 5-Mbps threshold would seem ample for most applications envisioned today, it might, on the one hand, prove inadequate in the future, or, on the other, raise questions about whether its costs today would exceed what customers are willing to pay for today. Later, this chapter explores several approaches to answering the fundamental question, How fast is fast enough?

As indicated earlier, the effective speed for interacting with an Internet host is not merely a function of the performance of the broadband local access link—it depends on the entire path between the host and the user, and also on the loading on the host computer. As a result, depending on the circumstances, improvements in the performance of one link does not necessarily improve overall performance—it may only shift the bottle-

neck. Network infrastructure such as caching and content hosting within the local ISP access networks also has a substantial effect on perceived performance to the end user and loading on the connections to the core Internet.

Where are the bottlenecks, and how might they shift as broadband access technologies are upgraded? Today, for dial-up users interacting with most commercial hosts, the bottleneck is the last mile dial-up connection. With the current generation of deployed broadband—cable modems, DSL, or wireless services—the location of the typical bottleneck, at least for routine Web access, is less clear. It may be in the last mile, within the local ISP network, at the upstream linkage between the cable-modem or DSL ISP and the Internet core, closer to the host, or even in the user's PC.

From an applications perspective, DSL and cable modem broadband offerings today remove barriers to many applications. Advanced fiber-to-the-home (FTTH) networks with very high capacities (such as gigabit ethernet) enable additional applications, but they also illuminate a new set of barriers—such as the cost of core Internet connectivity at extremely high speeds—which present an obstacle to the widespread deployment of these applications. An FTTH network offers enormous amounts of bandwidth (e.g., gigabit Ethernet speeds) *within* the service area, but the fiber network's connection to the core Internet service providers may in fact be concentrated into a much slower link (say T1, 1.544 Mbps) that is shared by all users of the network. In this regard, residential fiber broadband networks will come to resemble the networking situation on university or corporate campuses, where local bandwidth is plentiful, but connectivity beyond the local campus is a comparatively constrained shared resource.⁶ The link to core Internet service providers is typically going to be paid for on a leased basis, and its costs rise as the bandwidth of the link increases. In contrast, high local bandwidth within the community incurs mainly the fixed capital costs of installing and lighting the fiber network (which can be financed for a long period of time). Thus, where very fast FTTH networks are deployed, they can have the property that access to hosts within the community served is very fast, but that more general access to Internet sites is much slower; it may thus be possible to exchange high-definition video with a neighbor or a local community center, but difficult in the short term to extend this level of performance beyond a modest geographical region.

⁶This phenomenon has been experienced by universities that saw new bandwidth-intensive applications such as Napster clog their backbone Internet connections.

Latency and Jitter

While the net throughput is the most significant enabler of many applications, two additional parameters are crucial for applications that depend on real-time delivery of information or interaction, such as telephone or interactive game playing. "Latency," or delay, is a measure of how long it takes to deliver a packet across the network to its destination. Latency is a function of the distance the packet travels (speed of light, which is of particular significance for traffic carried over geosynchronous satellites), the length of time the packet waits in queues within the network, and the delay that results from retransmission when a packet is dropped due to congestion within the network. Latency especially affects applications that depend on interaction, such as human-to-human conversations, games, and the like. "Jitter" measures the variation in latency, resulting from such factors as variations in the path taken by each packet, variable queue lengths, or variations in the level of congestion within the network. Even if the average latency is acceptable, high jitter may make the application unusable nonetheless.

Symmetry Between Upstream and Downstream Capacity

Today, telecommunications services, including broadband, do not necessarily provide the same capacity up- and downstream. At one extreme, digital cable television service and direct broadcast satellite service provide a very high data rate digital connection into the home. These services may also provide a low data rate return path—over the same link or over an alternative return link using a phone line—to enable enhanced services such as pay-per-view. However, most users probably would not think of these services as broadband—they expect broadband to include high-speed Internet access (perhaps along with these predominantly one-way services). On the other hand, broadband does not necessarily imply that one must have anything close to symmetric bandwidth to and from the premises—though some would argue that it will, over time, as a consequence of the minimum bandwidth particular applications require.

The asymmetric services typically found in today's residential broadband services were designed with one of two asymmetrical application classes in mind. One class is Web browsing, where a low-bandwidth upstream connection serves to carry a user's requests for Web pages, and the higher downstream connection returns the content the user has requested; e-commerce or other applications in which users interact via entering information in Web forms involve a similarly asymmetric communications model. The other class, audio or video delivery, in which a small amount of data is sent upstream to select and direct delivery of a

particular stream (delivery of packets for playback in near-real time), is even more asymmetric.

While Web browsing has been a dominant application of residential broadband, accompanied by more limited audio/video streaming, peer-to-peer applications have surged recently. These applications, which use many individual computers instead of a central server to distribute content, require significant upstream capacity for each computer. They have, as a result, presented ISPs with traffic loads that are at odds with the ISPs' assumptions about asymmetric traffic⁷ and have raised questions about what shape user demand will take in the long term. Similar pressures result from other applications in which users host content on their local machines, creating upstream demand whenever this content is requested. These pressures, at odds with the capabilities of today's networks, have also led some broadband ISPs to prohibit customers who subscribe to consumer/residential services from running servers on their computers.

It is not clear at this point how traffic patterns will evolve as applications mature and as the population of broadband network users moves beyond early adopters. There is at least some reason to believe that the traffic patterns will in fact be asymmetric, though perhaps not as strongly as some of the broadband ISPs initially assumed in their network design and pricing models. But it is also important to recognize that some of the demand for symmetric bandwidth is a political rather than a business or engineering proposition. If the networks are designed to make it impossible, or very expensive, for individuals to originate the kind of traffic associated with the provision of services or content to significant audiences, it would foreclose the possibility that high-traffic upstream services will emerge on a highly distributed, grass-roots basis. This prospect points toward a model of the future broadband-enabled Internet as an environment dominated by commercially provided services connecting to customers—an outcome that in the view of some would fall short of broadband's full potential.

More generally, while much of the focus on broadband has been on its potential as a channel for delivering information, broadband also provides a more general communications channel (into and out of the premises). On the one hand, e-mail and instant messaging are prominent examples of communications applications that do not depend on large amounts of upstream bandwidth (or downstream, for that matter), but

⁷For example, according to Jim Hannan of Sprint Wireless at the committee's June 2000 workshop, "As a point of data, in our experience, we would love it if [the ratio of downstream to upstream] were 10 to 1. You know, our network model said worst case: 8 to 1. Unfortunately, our experience is 2 1/2 or 3 to 1, downstream to upstream."

that provide evidence of demand for convenient, Internet-based communication. On the other hand, as consumers start transmitting video clips (produced using increasingly inexpensive digital video cameras), bandwidth requirements could increase significantly. On the horizon are a number of communications applications—telephony being the most obvious—that place increasing demand on the upstream channel.

The detailed discussion of application classes below suggests that the jury is still out on the long-term implications of such applications for symmetry demands in broadband services. Nevertheless, a number of pressures for increased upstream capacity are evident.

Always-On

In addition to higher bandwidth, a broadband connection also generally provides an always-available connection to the Internet. One principal implication of always-on broadband service is that, for the first time, residential users have nearly instant access to Web or other Internet services on demand. Before the advent of broadband services, residential and many small business Internet users were confined to using a dial-up line to access the Internet. With dial-up, the user faces a noticeable delay—the sum of the time it takes to place a call between the user and ISP modems, the time it takes for the two modems to negotiate a connection, and the time it takes to log in (generally by authenticating the user via a password) to the ISP. The delay is increased if the user makes a habit of turning off the PC between sessions, since the time it takes the computer to boot up must also be added to the time it takes before a user can access the Internet. By eliminating the need to place a telephone call, broadband services greatly reduce the time required. While there is some delay associated with negotiating communications parameters when the customer's modem is powered up, these devices are designed to be left on all of the time, meaning that there is continuous connectivity between the modem and the network to which it is attached. Laptop computers have had power management features for some years; more recently this capability has been added to desktop computers. Power management capabilities make it possible to have computers "sleep" (quickly switched to a low-power state) and then be reawakened whenever the user wishes to access Internet resources.

The term "always-on" might conjure up visions of some sort of compelled use in which computers or applications must be left running all of the time. Always-on does not imply this; it refers merely to a characteristic of broadband networks that enables network communications to be initiated at any time. Users remain free to close software programs or shut down computers as they wish. Of course, some applications and com-

puter devices will be designed to work best when they are always connected, and many users may choose to keep some computers or applications in an always-connected state.

Research has shown that removing the start-up delay changes the way that users perceive and use the Internet. Because the overhead associated with accessing the Internet becomes very small, there is more casual use of the network for very short tasks—sending a short message or looking up a piece of information. This change also has the effect of significantly reducing the length of a typical “session,” as users begin to regard the network as an always-available utility, even though total use may stay the same or increase. Users also may change their behavior to leave their PCs on more of the time, either fully powered up or in sleep mode.⁸

The popularity of instant messaging (and chat rooms) despite the fact that the majority of household Internet users connect via dial-up service demonstrates the high level of user interest in using the Internet for communications applications with a real-time dimension (in contrast to the delays associated with e-mail). However, because dial-up users are unlikely to be connected to the Internet at any given time, it is generally not assumed in today’s applications that they are always connected. In an always-connected broadband environment, these applications become much more powerful. For example, the value of an Internet telephony application is limited if calls can only be placed if the person being called happens to have an active dial-up Internet connection at the time the call is placed. Many other applications are most useful when the equivalents of telephony ringing and signaling capabilities are available.

Bandwidth aside, the combination of quick access and new applications is compelling. Indeed, one sees this value reflected in value-stratification practices of several DSL service providers. They provide two tiers of residential service—a cheaper one that requires users to go through a log-in screen when they wish to connect and a more expensive one that provides continuous connectivity. Results from the INDEX project at the University of California, Berkeley, which was designed to learn how much people are willing to pay for various levels of bandwidth on demand, provide further support for the proposition that a considerable fraction of the value that users attach to broadband may stem from its always-on quality. In that experiment, users tended to keep a basic 8-kbps service, which was provided free and was on all the time, but on average they placed a surprisingly low value on their waiting time, which in the ex-

⁸Ken Anderson and Anne Page McClard. 1998. *Always On: Broadband Living Enabled*. Technical report, Broadband Innovation Group, MediaOneLabs. October.

periment could be avoided by purchasing a higher-bandwidth capability on demand.⁹

There will be important new applications—health monitoring, security, and the like—which will be possible only with the always-on characteristic, but users will choose whether they want to use those applications. The notion is familiar—telephones are generally left connected, ready to respond to a ring signal. What always-on should conjure up, however, are concerns about security. Today's end-user computing devices are vulnerable to a variety of network attacks.¹⁰ Always-on connectivity increases their exposure to these threats (see below).

Connectivity Sharing and Home Networks

Another attribute that users sometimes associate with broadband access is that of a premises network. Dial-up access is generally done from a single machine. The speed of the dial-up connection is slow even for a single machine, so trying to share that bandwidth among multiple machines is not generally very desirable. Moreover, it is common for each PC to have an analog modem. Thus, users have generally arranged to time-share the household phone lines sequentially among a number of machines in a home (though quite possibly not without disharmony resulting from contention over access to the phone lines). A broadband connection, however, by virtue of its always-on nature and greater capacity, makes it reasonable to support multiple machines concurrently. Thus, broadband Internet access and use of home networks will increasingly be interrelated.¹¹

Spurred in large part by the initial deployments of broadband services to the home, a variety of home networking technologies are available in the consumer market (Box 2.1). The year 2000 represented something of a turning point for the mass-marketing of these devices, seen in the increasing number of vendors offering products and in falling prices.

⁹Hal R. Varian. 2000. *Estimating the Demand for Bandwidth*. Technical report, University of California at Berkeley. August 1999, revised August 29, 2000. Available online at <<http://www.index.berkeley.edu>>.

¹⁰See CSTB, NRC. 1998. *Trust in Cyberspace*. National Academy Press, Washington, D.C.

¹¹Another possibility, most plausible with wireless access in which each device has its own antenna, is that the devices within the home would each have their own connection to the network infrastructure. While this configuration, in which there is no home network, may have advantages for the consumer in terms of ease of management, it also has the limitation that interdevice communications would have to be routed through the local access link to the core network and then return through another local access link. Such an architecture would likely preclude such applications as having a DVD player transmit a program to a remotely located video display.

BOX 2.1 In-Home Networking Technologies

Wired Ethernet

This is the industry-standard networking technology commonly used in office and business settings. Ethernet today routinely runs at 100 Mbps, which provides adequate performance for many applications, even video applications. (Quality-of-service management, which is not provided by standard Ethernet, may, however, be required for very-high-end applications.) The primary difficulty with using wired Ethernet is that most homes are not wired with the requisite category 5 wiring. There are long-term efforts aimed at promoting the installation of networking-compatible wiring in new construction;¹ rewiring of existing dwellings is generally restricted to early technology adopters who are willing to deal with the cost and disruption associated with installing new wires within a house. With a wired infrastructure in place, Ethernet is a very inexpensive solution, with interface hardware available from many vendors and frequently incorporated into computers. It also offers ample headroom for future speed upgrades (1 gigabit per second [Gbps] has been demonstrated over four category 5 twisted pairs of 150 feet). For the slowly growing set of users that have the requisite wiring, Ethernet is a likely technology of choice.

Phone Line Networking

Phone line networking operates by using a high-frequency carrier superimposed over a home's existing analog voice telephone wiring, allowing any standard phone jack within the house to become a network port. For homes that have phone jacks already installed in locations where networked devices are likely to be located, as is the case in many U.S. homes, this is an attractive solution. An industry standards group, the Home Phoneline Networking Alliance (HPNA), has defined two standards: HPNA 1, which provides a 1-Mbps data rate, and HPNA 2, which increases the data rate to 10 Mbps while maintaining backward compatibility with the first version. The HPNA specification does not interfere with the normal voice operation of the phone line, and it operates in a different band from a G.992.2 asymmetric digital subscriber line (ADSL) (G.lite) so that it can coexist with that standard as well. HPNA products are currently offered by a number of vendors for attachment to existing computers via parallel port, universal serial bus (USB), or via a peripheral component interconnect (PCI) add-in card. Additionally, a number of PC vendors are shipping home PCs with HPNA integrated into the system directly. The technology is relatively inexpensive and HPNA is being bundled with home computers today, much as most home PCs include analog modems. HPNA has the drawback of interfering with broadcast radio and some, but not all, types of DSL.

¹Wiring America's Homes, sponsored by the Home Automation Association. See <<http://www.homeautomation.org/wah.html>>.

Wireless Networking

Wireless solutions are very attractive in that they require no wires to be installed. They allow networked computers to be located anywhere within the home, and they support mobility within the house. In the past, wireless technologies suffered from being too expensive for broad consumer acceptance, had lackluster performance, and were too power-hungry for use in battery-powered devices. More recently, two standards using the **unlicensed 2.4-gigahertz (GHz) band that overcome these shortcomings** have emerged: 802.11b and HomeRF. 802.11b has been defined by the Institute of Electrical and Electronics Engineers (IEEE) and supported by the Wireless Ethernet Compatibility Alliance industry consortium. It uses direct sequence spread spectrum rather than frequency hopping, and operates at up to 11 Mbps for data. A number of vendors have announced 802.11b products. 802.11b is also being widely supported in business environments and public access points and is also being **used in a number of community access** initiatives. With both standards, there are issues of **interference** with other uses and **other** technologies that make use of the same radio spectrum.

Products using a new standard from IEEE (802.11a), which use the 5-GHz band and promise data rates of 54 Mbps, are under development. Another wireless LAN standard operating in the 2.4-GHz band, 802.11g, would boost speeds to as much as 54M bit/sec while proving backward compatibility with the 802.11b. Another relevant wireless standard is Bluetooth. Currently Bluetooth targets data rates of less than 1 Mbps, although work is ongoing to define higher-speed versions. However, it focuses now and probably in the future on short-range "personal area networking," at less than 10 meters (m). While Bluetooth will likely have applications in the home, the range may restrict it from being a general-purpose solution for home networking. Also, because both Bluetooth and the current generation of wireless LANs use the same spectrum, there are still unresolved issues of how to efficiently share the spectrum and avoid interference.

Powerline Networking

Because power outlets are ubiquitous in homes, there has been long-standing interest in technology that would provide high-speed networking over household AC power wiring. However, there is a great deal of interference present on power wiring, because these wires are used to supply power to motors and other very electrically noisy devices. In addition, many homes utilize both phases of the supply power in different circuits within the home; communicating across outlets connected to different phases requires installation of bridges between the phases, which may be difficult for consumers themselves to install.

In multidwelling units there are also potential interference problems, analogous to those faced in wireless networking, associated with multiple users using common lines. To date, the most common use of networking over the powerlines has been for X-10 home control devices that permit remote operation of lights and the like. In the year 2000, several companies began shipping, or announced, powerline-based data networking products, with claims of data rates up to 10 Mbps. The HomePlug Powerline Alliance available online at <www.homeplug.org> has been formed and is promoting its 10-Mbps baseline standard. Field trials and certification lab are reportedly approaching.

In 2001, a number of products that integrate home networking technology have been announced. Gateways connect to DSL or cable modems and provide home networking via a variety of technologies—the range alone indicates progress in standards-setting and growing technology maturity in this arena. Vendors also are integrating these functions into the modems themselves, aided by the minimal cost of adding home networking functionality to the silicon that implements the modem. Such integration extends to computers and Internet appliances as well, with these devices incorporating one or more home networking technologies. These trends work toward making broadband installation a simpler process for the consumer, eliminating the need for additional wiring, and lessening the need for visits by installers.

Finally, even for a single computer on a broadband connection, users often have an expectation of being able to multiplex among a number of applications. Dial-up access, in contrast, often constrains a user to be doing one thing at a time with the system. This may be because it is simply too slow to have multiple activities sharing the connection, or it may even be because the resulting effective slowdown of the connection simply renders one of the applications unusable. For example, listening to Internet audio while also downloading files is likely to make the audio drop out over a dial-up connection, whereas simultaneously listening to Internet radio, downloading files, and surfing the Web is quite feasible over a high-speed connection. Although rooted in what is enabled by the speed of the connection, this change is as much a change in user behavior as it is about a new technological capability.

Addressability

A critical requirement of many applications is that a user's computer be addressable in some fashion by software running on computers elsewhere on the Internet. This means that someone on the Internet can initiate communications with the user, much as a telephone caller can place a call to a subscriber by dialing the subscriber's telephone number. Addressability also enables such functionality as a user being able to run a server that other Internet users can access (a capability demonstrated with Napster and Gnutella).

While addressability is commonplace in the business use of the Internet, it is the exception for residential users today. One reason is provider policy—not all service providers allow inbound connections in their policies. Another reason is the technical means of connecting the user to the Internet: Addressability is most easily provided when each computer device within the home has its own globally addressable Internet Protocol (IP) address. Many residential customers are provided dynamically as-

signed addresses, and even if they have a static address, they may only have one or a few—meaning that each computer within the home may not have a globally routable address as a result of the use of network address translation by either the ISP or within the home gateway used to provide connectivity to multiple computers within the home.

Addressability is a double-edged sword. Being able to address home computers from other computers attached to the Internet enables powerful new applications, but it carries with it issues of security and privacy that will need to be solved. For example, exposing computers to the Internet on a continuous basis makes them more attractive and potentially vulnerable targets for attacks aimed at destroying data stored on them or otherwise disrupting their operation, viewing information stored on them, or manipulating them for use in launching attacks on other Internet services. ISPs may take some steps to help protect users, such as filtering out possibly hostile traffic (e.g., blocking transmission of NetBIOS packets, which could be used to alter or delete files and are generally only intended for use within local area networks), but the security of home computers depends in large part on the security of the computers or gateways installed within the home. That security depends, in turn, on the adequacy of add-on devices such as firewalls, which users have to configure to filter traffic and deliver warnings appropriately, and on the quality of typical home-system software, which tends to be low from the perspective of security.¹² Another risk is posed by always-connected, addressable sensor devices such as cameras. Breaches of security could enable outsiders to read or take control of these devices, allowing them to view or otherwise monitor what happens within people's homes. And tampering with control devices could cause direct physical harm—for instance, someone with access to a household's networked climate controls might maliciously turn off the heat during a cold spell. In a development that signals increased awareness of security issues and suggests a willingness to trade off flexibility for additional security (despite arguments in favor of the end-to-end principle in technical communities), consumers have been purchasing gateways that incorporate firewalls as well as installing software-based firewalls on individual computers.

Controls on Applications and Content

Questions about potential limitations on the content and applications that will be available to the typical consumer of a broadband service have figured in the cable open access debate, which features a conflating of

¹²See CSTB, *Trust in Cyberspace*, 1998.

practice, motivation, and perceived impacts that varies with point of view. The issues are outlined here because they contribute to both perceptions and realities of commercial broadband offerings, which will be implemented based on choices made by providers that must decide among technological and strategy options. Given the newness of the marketplace, it is easy for critics to assume intent based on experience with traditional media, especially cable, and it is hard to predict what practices will succeed in the broadband marketplace, regardless of their fate with traditional media. About the only certainties at this point are that the service providers are trying to make money, that content providers seek access to users, and that provider policies and practices are evolving at the same time as the technologies and businesses.

Already, there are various models for Internet service, ranging from the ISP that provides only basic IP connectivity to the ISP that provides a modest bundle of services and content, to ISPs such as America Online (AOL) and Microsoft Network (MSN) that aim to provide a wide range of products and services. Providers can seek various degrees of control over particular applications that run over their service or restrict access to particular content, perhaps simply by making it much easier to access preferred content. This may happen for various reasons, and the effects may be either primary (to promote use of certain content) or secondary (to make use of certain content less convenient).

If some content (e.g., from sources with business relationships with providers) is cached and easily accessed, other content may appear to be harder to get to. Uncached Web content will, for example, be slower to load—especially if the source is far from the user and/or on a network with poor connectivity to the user's. In the extreme case, access to non-cached content might be poor enough to make it seem effectively filtered; consumer advocates express this concern about the fate of content from nonprofit sources, but the concern remains hypothetical.

Service providers provision bandwidth, especially upstream, based on a particular business model—which makes assumptions about who is sending how much of what to whom—or on the assumption that a certain fraction of traffic can be cached. Or, providers might use restrictions—such as restrictions on virtual private networks or on operating household-based servers—as a means of value stratification, charging more to those who value, use, and will pay for more flexibility or capacity.

Actions that restrict upstream communication raise concerns about innovation enabled by end users' being able to originate content or applications. A targeted approach by ISPs might alleviate some concerns. For example, a provider that is concerned with upstream bandwidth scarcity might more effectively deal with excessive bandwidth use (relative to provisioning assumptions) by applying measures that monitor or restrict

bandwidth consumption rather than by prohibiting all users from, say, running servers. While some legacy equipment might not support bandwidth monitoring or control, most ISP equipment today would permit this. Opinions will differ as to whether restrictions reflect legitimate operational considerations; valid business decisions to differentiate customers; or unreasonable attempts to limit customer access to applications, content, and services.

Assessment of provider conduct should distinguish between the use of caching and similar techniques, which are aimed at improving access to some content by moving it (via distributed copies) to locations closer to users, and the use of filtering, which limits access to some content entirely. Steering or restricting customers to certain content runs counter to the traditional Internet model, in which Internet service is deemed synonymous with access to all Internet content. The success of these “walled-garden” models depends on whether consumers want the preferred content (or, in the view of critics, have no alternative). Actual consumer preferences and reactions to their experiences as users confronting differential ease of access to content are an important but unknown factor.

Implications of Network Design/Architecture

Another parameter that deserves consideration is whether “broadband” refers exclusively to Internet service or is a more inclusive term that refers to a set of data communications services. Is the point of broadband to bring the Internet to the home or small business at much higher speed and with characteristics such as always-on, or is broadband really about delivering to the home a bundle of digital services, which include IP service, that are demultiplexed at a gateway? Cable television systems today deliver both IP data and digital television signals. Higher-quality pictures and greater system capacity than analog cable systems could deliver were the original motivation for deploying hybrid fiber coax (HFC) in cable systems—IP data capabilities were added later. This video via MPEG (a standard for video compression from the Motion Picture Experts Group) streams is largely one-way (possibly with a low-capacity return channel to allow the selection of content or other interactive features), and the content is offered through various service bundles and pay-per-view options defined by the service provider. Looking forward, to what extent will services be delivered using plain-vanilla IP versus more specialized protocols and architectures? Running video and audio over plain IP, for example, is not without problems today, and these, together with business considerations, may well lead providers to devise other network protocols and systems to deliver audio and video alongside IP (for other applications), perhaps coming out of the set-top box into

proprietary or consumer-electronics-oriented interfaces that feed TVs and other appliances—possibly with additional features such as intellectual property protection engineered in at a very low level.

Another characteristic of broadband networks is the increasing sophistication of network monitoring capabilities. Usually discussed in the context of quality of service, commercial broadband providers are motivated to track individual (per-home) usage patterns to accurately assess and supply different levels of broadband capabilities. For example, users could pay for larger upstream bandwidth to support a home business.

APPROACHES TO DEFINING BROADBAND

The discussion above suggests the pitfalls of picking a specific bandwidth threshold as defining broadband, and Chapter 3 displays the wide variety of requirements posed by different classes of applications. A better approach may be to define a broadband service as one that meets a certain set of objectives—objectives that will change over time.

One important task in engineering any system is to examine the trade-offs between the performance of different components. The first definition offered here takes into account the multiple performance parameters that affect the overall performance of a broadband service. For example, one can compensate for limited bandwidth by compressing the data being transmitted, a technique that is widely relied on today. Or, one can take advantage of ample local storage capacity as might be found on today's computer hard drives to store content so that it does not have to be transmitted over the connection at the time it is needed.¹³ This perspective motivates one definition of broadband:

Broadband Definition 1. Local access link performance should not be the limiting factor in a user's capability for running today's applications.

To illustrate how this definition could be used in practice, consider how this applies to Web browsing. Speed here is inherently limited by factors independent from those imposed by the bandwidth of the local access connection. The speed-of-light transit time poses a fundamental limit to the rate at which data can be sent across the network using the data transmission protocols on which the Web is based. For today's typical Web page, a user "cruising the Web" will not see any material performance improvement once his or her access link has a capacity of about

¹³One local storage strategy, caching, keeps local copies of frequently used items. Another strategy, replication, preloads information so that it is available for later use.

1 Mbps.¹⁴ In other words, upgrading an individual user's 1-Mbps link with one 10 times faster would not speed up the transfer of a typical Web page. So for this application and at this point in time, an access link provides broadband service when it provides a capacity approaching 1 Mbps. It is important to realize that this is a statement about *today's* Web content and protocols and *today's* network. The availability of higher-performance links would likely give rise to richer content that would take advantage of that availability and would also provide incentives to colocate caches or streaming servers close to the broadband access point, leading to an improved user experience as a result of the faster local access link.

This viewpoint can also be useful in identifying where to make investments to alleviate potential bandwidth bottlenecks within the network. With a major source of congestion residing in the transit circuits that connect broadband providers to the Internet, improving local access performance above that bound will not improve the net performance seen by the user accessing the Internet. Everywhere in the Internet (except possibly on the access link), the average traffic of a sufficiently large number of users aggregated together on a given link will be constant on average, even if the traffic demands of individual users varies considerably. However, the guarantee is statistical in nature, meaning that there is always the potential for fluctuations that result in congestion somewhere in the network. In contrast, an individual access link is subject to much more predictability. An unshared link (e.g., a DSL connection to the central office) will be a bottleneck only if it is simply too slow (has too little bandwidth). Even for local access technologies with a shared local medium (e.g., an HFC system or the feeder to a DSL remote terminal), the implications of sharing can be understood fairly easily, because one knows (based on provisioning and traffic engineering) how many potential users there are on any given segment. Thus, where there is a compelling application for which the access link bandwidth is identified as the limiting factor, a well-defined investment in upgrading the capacity of the link can solve the problem. But for most other links within the network, where loading can only be addressed on a statistical basis and where there are many options for making investments that might improve matters, it may be less clear where to invest—which in turn may mean that making an

¹⁴The transfer of a "typical" Web page (one used by the World Wide Web Consortium as a benchmark for performance) from a server for which the network latency is a typical 100 milliseconds. If one varies the bottleneck bandwidth and plots the total page transfer time, the curve approaches an asymptote of about 6 seconds when the transfer speed reaches about 1 Mbps.

investment in the local access link will not improve application performance. (Other strategies may be used in such circumstances, such as caching or replication within the network of the access provider.)

Broadband Definition 1, however, gives an answer that is only correct for a given set of applications at one point in time. What happens when new applications come along? In fact, the performance of broadband access will be a key factor influencing the emergence of new applications, since new applications that demand higher transfer speeds cannot take off until there is a critical mass of users with the access capacity to use them. This motivates an alternative definition of broadband:

Broadband Definition 2. Broadband services should provide sufficient performance—and wide enough penetration of services reaching that performance level—to encourage the development of new applications.

This definition implies that a broadband access system is defined both by a technical evolution path and an economic evolution path that will allow it to play its part in the chicken-and-egg application cycle. The subscriber link is viewed as a potential bottleneck that inhibits innovation and constrains the development of new services elsewhere in the network. Those providing services over the Internet who feel constrained by the premises-link bottleneck may not be able to fully incorporate the benefits of relaxing this bottleneck in their own investment decisions (i.e., their incentives to “subsidize” broader deployment fall short of the true collective benefits)—because they are unable to internalize the benefits realized by other service providers.

One example of how this view comes into play is the asymmetry of broadband services. Anticipating a number of new applications that require greater upstream capacity, one can project increasing demand for upstream bandwidth arising from new applications. Yet, if connections remain highly asymmetric over the long run, then applications that need significant upstream capacity will be slower to appear. Because it takes into account the dynamic interplay between deployed technology and applications as well as the interplay between technology and economic developments, this second definition is likely to be the more useful one in planning and policy making.

Whichever of these definitions one adopts, it is quite apparent that a single number—be it 200 kbps or 2 Mbps—is not a useful definition of broadband (even if one focuses only on the bandwidth issue). However, not all values (from zero to infinity) will be equally meaningful. Applications such as those discussed in Chapter 3 tend to cluster into classes characterized by bandwidth and other performance requirements. This suggests that there will be a series of milestones along the way, with

multiple peaks that may well correspond to, or catalyze, application and infrastructure deployment milestones.

Today's residential broadband capabilities, which are typified by several hundred kilobits per second to several megabits per second downstream and several hundreds of kilobits per second upstream, support such applications as Web browsing, e-mail, messaging, games, and audio download and streaming. These are possible with dial-up, but their performance and convenience are significantly improved with broadband. At downstream speeds of several tens of megabits per second, new applications are enabled, including streaming of high-quality video, such as MPEG-2 (a standard defined by the Moving Picture Experts Group) or high-definition television (HDTV), download of full-length (70- to 90-minute) audiovisual files in tens of minutes rather than hours, and rapid download of other large data files. Reaching this plateau would enable true television-personal computing convergence. With comparable upstream speeds, computer-mediated multimedia communications become possible, including distance education, telecommuting, and so forth. With FTTH, a new performance plateau with gigabit speeds both up- and downstream would be reached. The applications that would take full advantage of this capacity remain to be seen.