To the participants in the 2008 Wharton Communications and Media Law Colloquium:

I have attached a draft of a manuscript describing the changes in the Internet and how they will affect policymaking.

Its primary audience is policymakers, although it should find an audience with academics as well.

I apologize both for its tardiness and how rough it is. Many of the arguments are simply sketched in. I rushed it out to get your reactions to the ideas contained in the manuscript while it is still in a very early stage when your help will be of the most benefit.

I look forward to seeing everyone on Friday.

Sincerely,
Christopher

The Transformation of the Internet

Christopher S. Yoo

INTRODUCTION

The emergence of the Internet as the world's dominant means of communication represents one of the most remarkable developments of the past decade. A medium that until the mid-1990s was primarily a way for academics to send e-mail and exchange files has become a juggernaut that has transformed almost every aspect of daily life, changing the way people work, play, and communicate and putting a previously unimaginable array of information at people's fingertips.

Its success has made people very used to the Internet as it is today. When it first emerged, the Internet was governed by a relatively simple set of principles. Business relationships were simple and uniform. End users could access any content without any fear of prioritization or discrimination.

What is often overlooked is how much the economic and technological environment has changed from when the Internet first arose. When it first arose, it was primarily the plaything of academics. The total universe of Internet users was relatively small and consisted of relatively sophisticated users with shared values and strong institutional support. These users employed a fairly narrow range of applications, such as e-mail and file transfer, which were predominantly text based (and thus used relatively little bandwidth) and which were relatively tolerant of delay. In addition, these end users employed a fairly uniform set of transmission technologies and connected to the network almost exclusively through personal computers. In addition, the pricing and other aspects of the business relationships between network providers was relatively simple.

In recent years, different network providers have begun deviating from this world of uniformity and universal interoperability. Critics have viewed these practices as attempts to exercise market power and as likely to harm innovation. These criticisms are rooted in a paradigmatic view that has been rendered obsolete by a number of developments that have reshaped the policy space in important, albeit largely unappreciated ways. Far from being an improper exercise of market power, these developments may represent nothing more than the normal effect of market forces and expected responses of an industry attempting to meet demand.

I. CHANGES IN THE INTERNET

A. Changes in the Population of End Users

1. Increases in the Number of Users and the Volume of Traffic

Perhaps the most dramatic change since the mid-1990s is explosive increase in the number of Internet users. What began as a means for academics to send e-mail and exchange files has exploded into a mass market phenomenon, skyrocketing from 16 million users in December 1995 to 1,463 million users in June 2008, an increase of more than ninety times.

Worldwide Internet Users (millions)

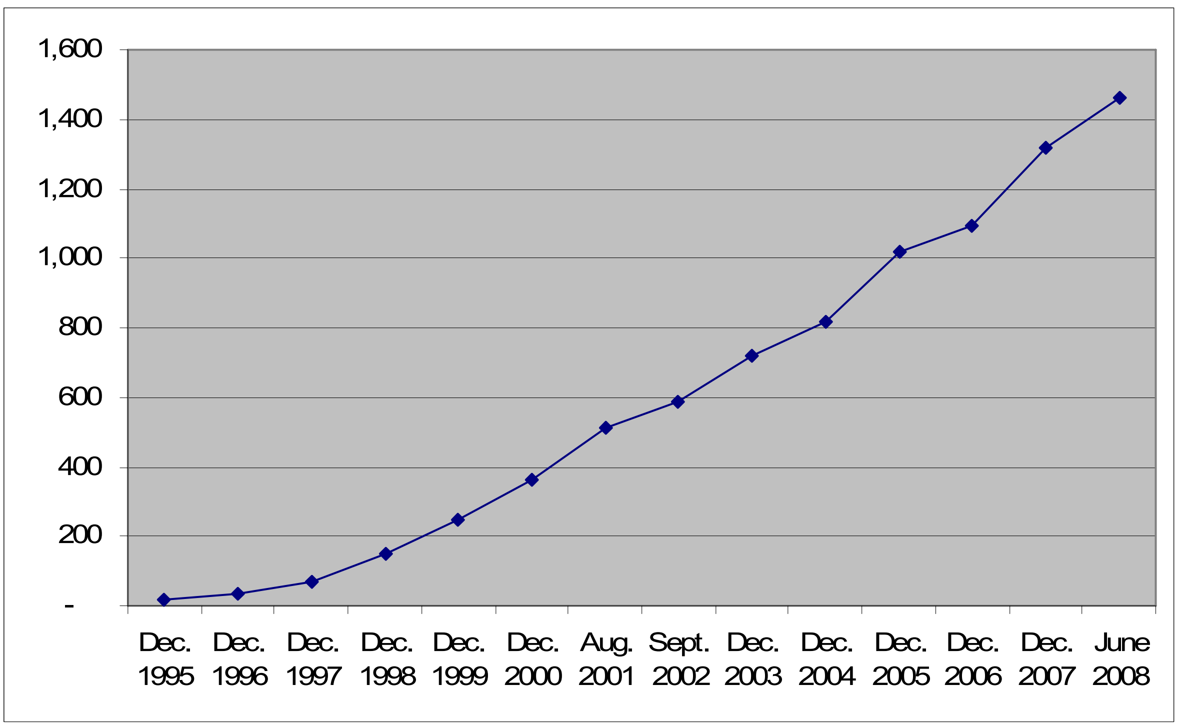

Source: Internet World Stats, Internet Growth Statistics, http://www.internetworldstats.com/emarketing.htm.

[Add in discussion about growth in the number of websites.]

The increase in the number of users has made the problems of managing the network much more complex. As an initial matter, the number of possible connections goes up quadratically with the number of endpoints, with a network consisting of *n* nodes consisting of $(1/2)n(n-1)$ potential links. Thus, even if we were just to take end users into account and ignore the growth in websites, the fact that the number of end users has increased more than ninety times since 1995 has caused the total number of connections to increase over 8900 times. In addition, just within the last year, the number of hops that the average Internet communication must travel has increased from _ to _.

2. Decreases in the Level of Institutional Support

The transformation of the Internet into a mass market phenomenon not only increased the number of end users. It also changed their character. As noted earlier, the universe of end users that initially populated the Internet consisted of academics, largely in the sciences. As a result, they were relatively technologically sophisticated. They also enjoyed a high level of institutional support from their universities.

Governance of the Internet was also greatly simplified by the fact that these researchers shared a common set of values. Interactions among them were reinforced by frequent contact in a wide variety of professional situations. As a result, the initial universe of Internet end users represented the type of close-knit community that can facilitate self governance without outside intervention. It also engendered a high degree of trust of other end users.

All of this has changed since the mid-1990s. The Internet has become dominated by individual end users operating without any institutional support aside from what their network provider can offer. This has created pressure to shift functions that used to be performed exclusively by end users, such as spam filtering and protection against viruses and other forms of malware, into the network itself.

In addition, the interests of Internet users has become increasingly heterogeneous and lacked the reinforcing mechanisms of belonging to the same community. As a result, Internet users have become increasingly distrustful of other users. To combat these threats, both end users and network providers are deploying firewalls, content filters, virus filters, and a wide range of other controls. While some of these continue to operate at the edge of the network, some of these functions are increasingly being transferred into the network's core.

B. Increases in the Variety of Ways the End Users Are Using the Network

As the number of Internet users has grown, the universe of end users has become increasingly heterogeneous. As the user base has expanded, people have wanted more things. The result is not only an increase in the number of users, but also an upsurge in the variety of demands that users are placing on the network.

The dramatic increase in the number and heterogeneity of Internet users has been paralleled by an increase in the volume of network traffic and the variety of ways that people are using the network.

1. The Shift to Bandwidth Intensive Applications

One of the most distinctive changes in the way people are using the network is an increase in bandwidth-intensive content and applications. The growth rates have varied over time. The most distinctive change was the sharp upsurge in 1995 and 1996 driven by the migration from websites that consisted almost exclusively of text, which required relatively little bandwidth to convey, to websites that incorporated large numbers of graphics, which required considerably more network capacity.

U.S. Internet Traffic
(Petabyes/month)

	Lower Estimate	Pct. Growth	Upper Estimate	Pct. Growth
1990	0.001		0.001	
1991	0.002	100%	0.002	100%
1992	0.004	100%	0.004	100%
1993	0.008	100%	0.008	100%
1994	0.016	100%	0.016	100%
1995	0.15	838%	0.15	838%
1996	1.5	900%	1.5	900%
1997	2.5	67%	4	167%
1998	5	100%	8	100%
1999	10	100%	16	100%
2000	20	100%	35	119%
2001	40	100%	70	100%
2002	80	100%	140	100%
2003	130	63%	210	50%
2004	200	54%	300	43%
2005	300	50%	500	67%
2006	450	50%	800	60%
2007	750	67%	1250	56%

Source: Minnesota Internet Traffic Studies, Internet Growth Trends & Moore's Law, http://www.dtc.umn.edu/mints/igrowth.html.

Specifically, after doubling each year between 1990 and 1994, Internet traffic grew by eight to nine times each year in 1995 and 1996. In 1997, growth once again returned to the pattern of doubling roughly every year until 2003, after which point traffic growth slowed to a rate of roughly 50%-60% each year.

U.S. Internet Traffic (Petabytes/month)

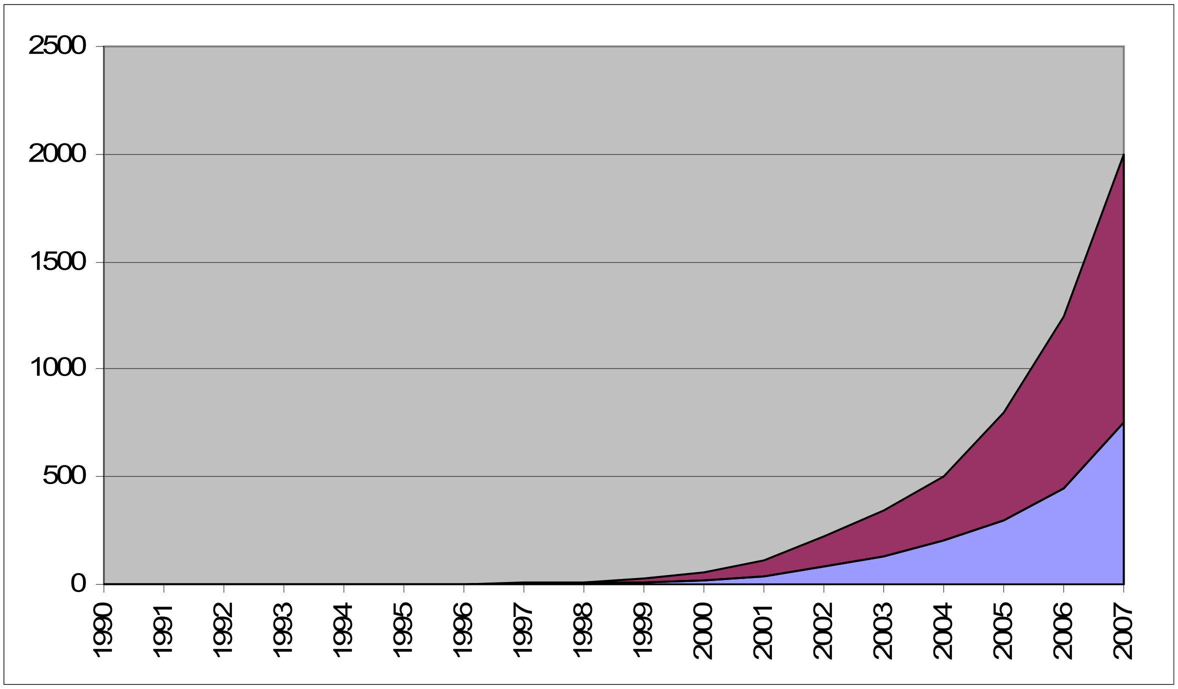

Source: Minnesota Internet Traffic Studies, Internet Growth Trends & Moore's Law, http://www.dtc.umn.edu/mints/igrowth.html.

Policy analysts are carefully studying whether the growth rates of Internet traffic may exceed the growth rates of network capacity, perhaps best captured by the popular concept known as "Moore's Law."[1] First put forth in 1965, Moore's Law is actually a prediction rather than a scientific principle. In its initial form, it predicted that the number of transistors that could be fit on a single chip would double each year.[2] Moore later revised his estimate to predict that the number of transistors that would fit on a chip would double every eighteen months.[3] The popular press has since generalized Moore's law from a prediction about the computing power of chips into a broader claim about the growth rate of network capacity. Because network capacity

1

2

3

depends on considerations aside from computing power, such as transmission and storage capacity, this reformulation is a significant extension from the original version of Moore's Law.

At times, the rate of bandwidth growth has struggled to keep up with the rate of traffic growth. Indeed, estimates indicate that from 2004 to 2006, the rate of growth of average and peak traffic on international backbones exceeded the rate of bandwidth growth by considerable margins. In 2007 and 2008, the rate of bandwidth growth accelerated and the rate of traffic growth slowed to the point where bandwidth growth exceeded the rate of traffic growth.

International Internet Traffic and Bandwidth Growth Rates

	Avg. Traffic	Peak Traffic	Bandwidth
2004	104%	105%	45%
2005	50%	57%	41%
2006	74%	57%	45%
2007	61%	63%	69%
2008	53%	58%	62%

The key question is whether the Internet will experience a repeat of what occurred in 1995 and 1996, in which the shift to more bandwidth-hungry form of Internet content will cause a sudden acceleration in traffic growth. In particular, the emergence of Internet video has caused some observers to warn of an "exaflood" that will cause the growth rates of Internet traffic to surge to 90% until 2015.[4] An oft-cited study by Nemertes research predicts that traffic will grow at roughly 100% per year, while capacity will grow at an annual rate of roughly 50%. This means that traffic growth will exhaust the usable network capacity by 2010 unless the world

---

[4] *See* Swanson & Gilder, *supra* note **Error! Bookmark not defined.**, at 8, 22; *see also* Bret Swanson & George Gilder, *Unleashing the "Exaflood"*, WALL ST. J., Feb. 22, 2008, at A15.

invests $137 billion in upgrading the Internet infrastructure.[5] Other observers, including Google and EDUCAUSE, have warned that the Internet will struggle to accommodate consumers' increasing demands for bandwidth.[6] Others have expressed skepticism about these claims. Cisco predicts that Internet traffic will grow at a 46% annual rate between 2007 and 2011,[7] which is a rate comparable to the growth of capacity. Similarly, the Minnesota Internet Traffic Studies (MINTS) conducted by the University of Minnesota similarly concludes that the evidence of the exaflood has yet to materialize, with the Internet continuing to grow at an annual rate of roughly 50% to 60%.[8] MINTS notes that the growth rates with wireless are considerably higher[9] and acknowledges that if Internet video deploys quickly enough, it could well overwhelm the system.[10]

2. The Emergence of Latency Sensitive Applications

Internet traffic is growing not only in terms of size, but also in sophistication. During the Internet's initial phase, the primary applications were e-mail and web browsing. For these applications, delays of a fraction of a second were virtually unnoticeable. The current Internet is increasingly dominated by more sophisticated applications such as streaming media, online

---

5 *See* Nemertes Research, The Internet Singularity, Delayed: Why Limits in Internet Capacity Will Stifle Innovation on the Web 31, 45 (Fall 2007), *at* http://www.nemertes.com/system/files/Internet+Singularity+Delayed+Fall+2007.pdf.

6 *See Internet Not Designed for TV, Google Warns*, PC MAG., Feb. 8, 2007 (quoting Google head of TV technology Vincent Dureau as stating at the Cable Europe Congress, "The web infrastructure and even Google's doesn't scale. It's not going to offer the quality of service that consumers expect."); John Windhausen Jr., A Blueprint for Big Broadband 7-11 (EDUCAUSE White Paper Jan. 2008) (also quoting studies by Jupiter Research and Technology Futures), *available at* http://www.educause.edu/ir/library/pdf/EPO0801.pdf.

7 *See* Cisco Systems, Global IP Traffic Forecast and Methodology, 2006-2011, at 1 (White Paper Jan. 14, 2008), *available at* http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/net_implementation_white_paper0900aecd806a81aa.pdf.

8 *See* Odlyzko Study, *supra* note **Error! Bookmark not defined.**.

9 http://www.dtc.umn.edu/mints/news/news_19.html.

10 http://www.dtc.umn.edu/mints/igrowth.html

gaming, telemedicine, and virtual worlds, which are often much more bandwidth intensive and much less tolerant of delay.

[Add more discussion here.]

3. The Partial Shift from Client-Server to Peer-to-Peer Architectures

Although the term, “peer-to-peer,” is often viewed as synonymous with file sharing or user-generated content, it actually embodies a more fundamental distinction. In the traditional Internet architecture, content and other files are stored in large computers at centralized locations (known as “servers”). End users (known as “clients”) use the network primarily to submit requests to those servers, usually by submitting a short bit of code such as a website address, also known as a uniform resource locator (URL). The server that hosts the requested files then downloads them to the.

In a peer-to-peer architecture, files are not stored in centralized locations. Instead, they are distributed across the network. Thus, unlike in a client-server architecture, in which computers that are connected to the edge of the network are divided into clients requesting files and servers hosting files, edge computers in a peer-to-peer architecture simultaneously request files and serve files. It is this less hierarchical structure that leads these types of edge computers to be called “peers” and this type of service to be called peer-to-peer.

Whether a network is comprised primarily of clients and servers or of peers has major architectural implications. If a network is organized around a client-server architecture, the traffic flowing from the server to the client tends to be larger than the traffic flowing in the other direction. As a result, it usually makes sense to divide the available bandwidth asymmetrically by devoting a greater proportion of the available bandwidth to downloads and a smaller proportion to uploads. Such asymmetry makes less sense if a network is organized around a peer-to-peer architecture, since each end user would represent an important source of upload traffic as well as download traffic.

At the time that network owners established the basic architectures for the major broadband technologies in the late 1990s, the Internet was dominated applications such as web browsing and e-mail that adhered to a client-server architecture. As a result, most network providers assigned bandwidth asymmetrically, devoting a greater proportion of the available bandwidth to downloading rather than uploading. For example, the dominant DSL technology is asymmetric DSL (ADSL), which initially supported theoretical speeds of up to 8 Mbps for downloading and 768 kbps for uploading (with actual download speeds reaching 3.0 Mbps).[11] More recent versions of ADSL support higher bandwidth, but still allocate it asymmetrically.[12] The initial cable modem architecture, designed around DOCSIS 1.0, supported maximum theoretical speeds of 27 mbps downstream and 10 Mbps upstream[13] (with the actual downloads speeds reaching 6 Mbps[14]). Finally, the service offered by wireless providers deploying EV-DO technologies are similarly asymmetrical, offering maximum download rates of 2 Mbps, with actual download rates ranging from 300-500 kbps and actual upload rates ranging from 40-50 kbps.[15]

These decisions were quite rational when they were made. I thus believe it is somewhat anachronistic and rather unfair to call the decision to adopt an asymmetric architecture the result

---

11 *See* 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 Telecommunications Act of 1996, Fourth Report, 19 FCC Rcd 20540, 20558 (2004) [hereinafter Fourth Report].

12 *See* DSL Forum, About ADSL, *available at* http://www.dslforum.org/learndsl/adslfaq.shtml (last visited Feb. 24, 2008).

13 *See* 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 Telecommunications Act of 1996, Third Report, 17 FCC Rcd 2844, 2917-18 ¶ 21 (2002).

14 *See* Fourth Report, *supra* note 11, at 20554.

15 *See id.* at 20560; CTIA—The Wireless Association, Wireless Broadband: High Speed Goes Mobile 3 (Apr. 2006), *available at* http://files.ctia.org/pdf/PositionPaper_CTIA_Broadband_04_06.pdf.

of “short-sighted[ness]” or “poor network design decisions.”[16] On the contrary, a network engineer in the mid-to-late 1990s, when cable modem and DSL systems first began to be widely deployed, would have had to have been exceptionally prescient to have foreseen the eventual emergence of peer-to-peer technologies. Since that time, network providers have begun developing new symmetric technologies, such as DOCSIS 2.0 for cable modem systems and symmetric DSL (SDSL) for wireline systems. DOCSIS 3.0 retains a degree of asymmetry, but to a lesser degree than DOCSIS 1.0. Very-High-Data-Rate DSL (VDSL) supports both symmetric and asymmetric services.

Indeed, even now it is far from clear whether a symmetric or an asymmetric architecture will eventually prove to be the better choice. For the four years preceding 2007, peer-to-peer traffic surpassed client-server traffic in terms of percentage of total bandwidth. A remarkable change occurred in 2007. Client-server traffic began to reassert itself, owing primarily to the expansion of streaming video services, such as YouTube. Some estimate that YouTube traffic constitutes 10% of all Internet traffic.[17] The ongoing transition of high definition television is likely to cause demand to increase still further.[18] Other video-based technologies, such as Internet distribution of movies (currently being deployed by Netflix), graphics-intensive online games (such as World of Warcraft) and virtual worlds (such as Second Life), and IPTV (currently being deployed by AT&T) are emerging as well.[19] Thus, in 2007, client-server traffic has retaken the lead from peer-to-peer, constituting 45% of all Internet traffic as compared with

---

16 Comments of Free Press et al. at 21, 22.

17 *See* Ellacoya Networks, Inc., Press Release, Ellacoya Data Shows Web Traffic Overtakes Peer-to-Peer (P2P) as Largest Percentage of Bandwidth on the Network (June 18, 2007), *at* http://www.ellacoya.com/news/pdf/2007/NXTcommEllacoyamediaalert.pdf.

18 *See* Bret Swanson, *The Coming Exaflood*, WALL ST. J., Feb. 20, 2007, at A11.

19 *See* Swanson & Gilder, *supra* note **Error! Bookmark not defined.**, at 12-14.

37% of all traffic devoted to peer-to-peer.[20] Thus, some industry observers predict that video traffic will constitute over 80% of all Internet traffic by 2010.[21]

Network providers thus confront a difficult decision. Not only must they determine the size and the location of the capacity to add. They must also determine the extent to which they should continue to embrace an asymmetric architecture based on their projections of the likely future success of applications such as BitTorrent and YouTube. Any imperfections in this projection is likely to have significant economic consequences.

C. Increases in the Variety of Communications Technologies

The Internet has also began to employ an ever wider variety of technologies. The variety increases both in terms of the technology used to provide the last-mile service as well as the devices employed by end users to interface with the network.

1. Greater Variety of Network Technologies

One of the biggest changes is the increase in the variety of last-mile network technologies. The different technologies vary widely in terms of availability of bandwidth and their susceptibility to local congestion.

When the Internet first emerged, end users connected to the network almost exclusively through high-speed connections provided by telecommunications companies. When the Internet began to be deployed to mass consumers, they relied on a narrowband architecture that used dial-up modems to make calls through the local telephone network to offices maintained by Internet

---

20 *See* Ellacoya Networks, *supra* note 17.

21 *See* William B. Norton, *Video Internet: The Next Wave of Massive Disruption to the U.S. Peering Ecosystem* (v0.91) at 2 (Sept. 29, 2006), *available at* http://www-tc.pbs.org/cringely/pulpit/media/InternetVideo0.91.pdf

service providers. Because local telephone lines are dedicated to individual customers, this technology was not subject to local connections. The bandwidth, however, was quite constrained, reaching a maximum theoretical speed of 56.8 kilobits per second (kbps) and typically offering delivered speeds of no more than 30 kbps.

During the _, Internet users began to migrate to broadband architectures. The two initial principal broadband technologies used relatively minor changes to the existing infrastructure to provide broadband service. Digital subscriber line (DSL) service takes advantage of the fact that that conventional voice communications only occupy the lower transmission frequencies (typically those ranging from 300 to 3400 Hz). DSL uses the higher frequencies (i.e., those above 20 kHz) to convey data communications through the same telephone line without interfering with voice communications. Reconditioning telephone lines for DSL service required minimal capital investments of roughly $400 to $800 per customer.[22] Because DSL uses telephone lines dedicated to each consumer, it is not subject to local congestion. The primary variant of DSL, known as asymmetric DSL, supports download speeds of up to 3 megabits per second (Mbps) and upload speeds of 768 kbps.

Cable modem service uses the coaxial cables used to provide multichannel television services and was initially deployed around a standard known as DOCSIS 1.0. Cable modem service has the advantage of providing greater bandwidth, supporting maximum theoretical speeds of 27 Mbps downstream and 10 Mbps upstream, and with actual speeds being closer to 6 Mbps. A key difference is that cable modem customers share bandwidth with the 300-500 households located in the same neighborhood, which makes it subject to local congestion.

22

Deploying cable modem service also required capital investments of between $800 and $1000 per subscriber.

Cable modem took the early lead, prompting some observers to raise concerns that no other technologies would catch up. As the FCC predicted, ADSL began to catch up over time.

Advanced Service Lines
(Over 200 kbps in Both Directions)

Technology	June ’00	June ’01	June ’02	June ’03	June ’04	June ’05	June ’06	June ’07
ADSL	326,816	998,883	1,852,879	2,536,368	3,768,019	13,176,095	18,310,957	23,381,289
SDSL and Traditional Wireline	758,594	1,088,066	1,186,680	1,215,713	1,407,121	869,772	946,874	1,027,937
Cable	1,469,130	3,329,976	6,819,395	11,935,866	17,567,468	22,745,012	28,878,587	33,939,919
Fiber	40,627	81,204	104,015	110,829	129,636	314,229	684,729	1,400,565
Satellite and Wireless	3,649	73,476	66,073	64,393	93,805			
Satellite						10,966	27,489	57,202
Fixed Wireless						191,229	333,209	553,919
Mobile Wireless						21,079	1,914,456	9,189,830
Power Line and Other						4,174	5,209	5,420
Total	2,598,816	5,571,605	10,002,942	15,863,169	22,966,048	37,332,557	51,101,510	69,556,081

High-Speed Lines
(Over 200 kbps in at Least One Direction)

Technology	June ’00	June ’01	June ’02	June ’03	June ’04	June ’05	June ’06	June ’07
ADSL	951,583	2,693,834	5,101,493	7,675,114	11,398,199	16,316,309	22,584,255	27,516,171
SDSL and Traditional Wireline	758,594	1,088,066	1,186,680	1,215,713	1,407,121	898,468	948,134	1,028,654
Cable Modem	2,284,491	5,184,141	9,172,895	13,684,225	18,592,636	24,017,442	29,174,494	34,408,553
Fiber	46,635	81,248	105,991	111,386	130,928	315,651	685,823	1,402,652
Satellite and Wireless	65,615	194,707	220,588	309,006	421,690			
Satellite						376,837	495,365	668,803
Fixed Wireless						208,695	361,113	586,141
Mobile Wireless						379,536	11,016,520	35,305,253
Power Line and Other						4,872	5,208	5,420
Total	4,106,918	9,241,996	15,787,647	22,995,444	31,950,574	42,517,810	65,270,912	100,921,647

In more recent years, the increasing demands for bandwidth is causing broadband providers to undertake more extensive investments to upgrade their networks. Unlike the initial transition to broadband, which only required reconditioning existing cable and telephone technologies, the new strategies require significantly greater capital investments. Verizon has committed to invest $23 billion between 2004 and 2010 to deploy its fiber-optic based FiOS network, at which point FiOS should be available to deliver upload and download speeds of up to 20 Mbps in half of Verizon's service area.[23] Wall Street's initial reaction was rather skeptical, but lately has become more receptive.[24]

AT&T is pursuing a different strategy, committing $6.5 billion to deploy its new uVerse system based on a telephone-based technology known as VDSL to 60% of its service area. [EXPLAIN BUILDOUT OF REMOTE TERMINALS.] Because all DSL technologies use telephone lines, uVerse is not subject to local congestion. Its limited bandwidth available through VDSL has required AT&T to take certain steps to make its service viable within its more limited bandwidth. For example, instead of following cable television's approach of transmitting all of the video channels regardless of whether anyone is watching them, uVerse uses a technology known as "switched digital video" so that it only transmits the channel being watched, an approach that is much more economical in terms of bandwidth usage. It also compensates for its lack of bandwidth is to reserve bandwidth for video.

These developments have forced cable companies to respond, with Comcast investing additional billions to upgrade portions of its network to DOCSIS 3.0. Thus, technologies now vary widely across providers and across any particular provider's service area.

---

23 *See* Verizon | All About FiOS, http://newscenter.verizon.com/kit/fios-symmetrical-internet-service/all-about-fios.html (last visited Dec. 19, 2008).

24

In addition, mobile wireless providers are in the process of upgrading their networks to support Internet applications. They are subject to severe bandwidth constraint. Mobile wireless providers are also more vulnerable to local congestion than any other technology, since its core voice product shares bandwidth with all other broadband services. Despite these limitations, these services have grown sharply, soaring from having no subscribers at the beginning of 2005 to capturing 13% of the market for advanced services lines by the middle of 2007.[25] The impact of mobile wireless broadband becomes all the more striking if one considers the other category of services tracked by the FCC, high-speed lines, defined to be services that provide at least 200 kbps in at least one direction. Measured against this standard, mobile wireless providers have actually become the industry leader, capturing nearly 35% of the market for high-speed lines, as compared with 34% and 27% for cable modem and DSL service respectively.[26]

Competition is also beginning to emerge from unlicensed wireless technologies, such as Wi-Fi and WiMax. The deployment of new wireless broadband services based on the recently auctioned spectrum in the 700 MHz range[27] promises to further diversify the number of last-mile technologies in the future.

2. Greater Variety of Devices Connected to the Network

Another major technological change is with respect to the devices that end users employ to connect to the network. When the Internet first arose, almost all users connected to the Internet through desktop personal computers (PCs). By their nature, PCs are rich in memory and

---

[25] *See* HIGH-SPEED SERVICES JUNE 2007 REPORT, *supra* note **Error! Bookmark not defined.**, at tbl.2 (reporting that as of June 2007, mobile wireless controlled 9,189,830 out of a total of 69,556,081 advanced services lines).

[26] *See id.* at tbl.1 (reporting that as of June 2007, mobile wireless controlled 35,305,253 out of a total of 100,921,647 high-speed lines, compared with 34,408,553 for cable modem and 27,516,171 for ADSL).

[27] *See* Public Notice, FCC, Auction of 700 MHz Band Closes (Mar. 20, 2008), *available at* http://hraunfoss.fcc.gov/edocs_public/attachmatch/DA-08-595A1.pdf.

computing power. In addition, they are particularly sensitive to conflicts between applications Because they are plugged into the wall, PC-based connections did not need to economize on power consumption.

End users now use a much wider variety of devices to obtain Internet service. The most striking is with respect to mobile wireless devices. Unlike PCs, these devices are very limited in terms of memory and computing power. They are extremely sensitive to conflicts between applications. And their performance is very sensitive to power consumption.

D. Increases in the Diversity of Business Relationships

Last-mile providers serve only one link in the chain, providing the link between end users and a distribution facility located in the their metropolitan area. The other key providers in the system are backbone providers and regional ISPs.

Backbone providers provide high-bandwidth, long-haul service between a limited number of recognized locations. The original backbone supported by the National Science Foundation provided connections among a small number of academic institutions. The restrictions the NSF placed on commercial uses of the backbone led a group of private companies to create an additional interconnection point known as the “commercial internet exchange” (CIX) located in Santa Clara, California.

The NSFNET Backbone, circa 1988

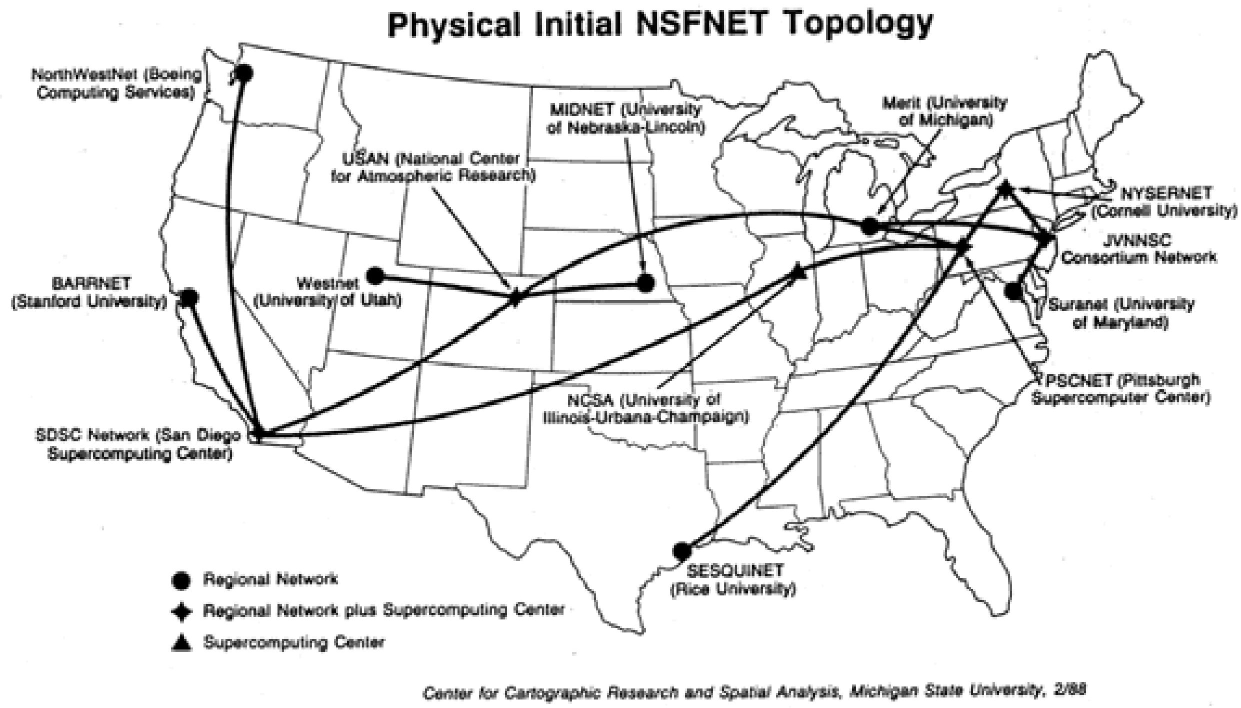

Eventually, the NSF backbone was privatized and is now run by companies such as AT&T, Cable & Wireless, Level 3, Qwest, and Verizon. These companies expanded the number of network access points and made them more rational given the pattern of traffic.

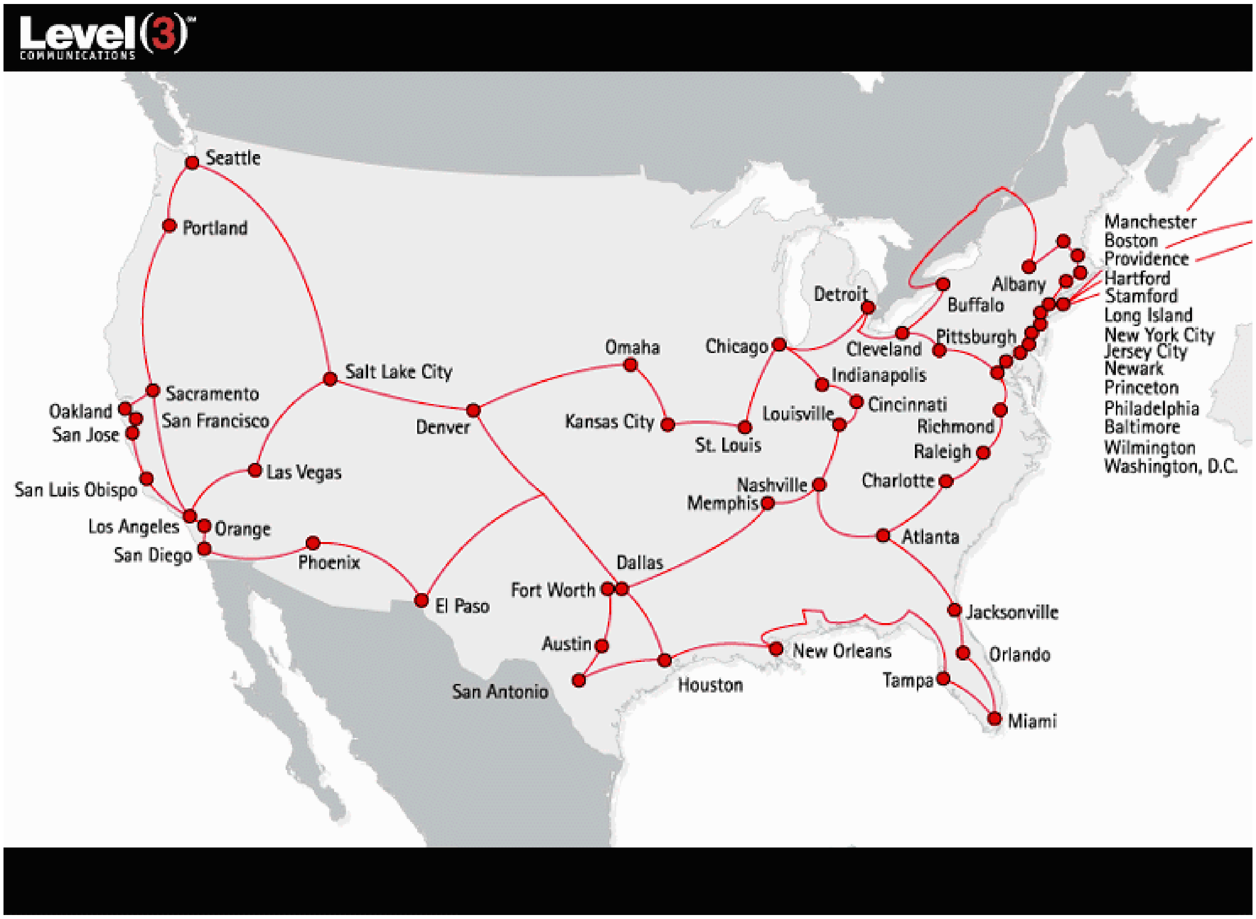

Now consider what would happen to traffic originating in a city that is not a network access point, say in Minneapolis. The service between the closest network access point (located in Chicago) and the local distribution facilities located in Minneapolis are provided by a third category of network provider known as regional ISPs.

Basic Architecture of the Internet

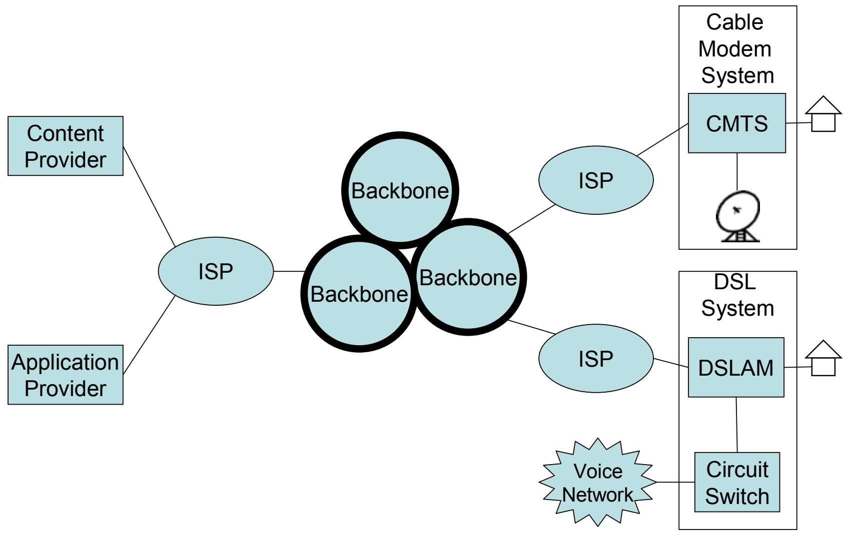

Interestingly, although the Internet is often portrayed as a cloud, in its original form, it was quite hierarchical. Each content provider maintained a relationship with precisely one regional ISP. Each regional ISP maintained a relationship with precisely one backbone provider. The result was a network architecture in which the path connecting any two points was unique. The architecture was also quite hierarchical and concentrated the economic power of the backbone providers.

In addition, the business relationships between different types of providers was relatively uniform. End users paying last-mile providers a flat-rate price for unlimited usage. Content and application providers pay last-mile providers based on the bandwidth that they use.[28] Last-mile providers pay transit fees to backbone providers. Backbone providers typically exchange traffic

---

[28] The norm is a system known as 95th percentile billing. Drawing on the insight that what matters is not the total bandwidth used, but rather peak usage, last-mile providers sample the bandwidth usage of content and application providers every six minutes and charge them based on peak usage. In addition, they forgive them the absolute peak usage and simply base charges on the volume in the 95th percentile time period over a thirty-day period.

under "peering" arrangements that are similar to telecommunications arrangements known as "bill and keep." Under peering arrangements, the originating backbone collects and retains all of the compensation for the transaction notwithstanding the fact that other backbones also incur costs to terminate the transaction. So long as the traffic initiated and terminated by each backbone is roughly equal in value, peering allows backbones to forego the costs of metering and billing these termination costs without suffering any adverse economic impact. Peering is not economical, however, in cases where the value of the traffic being terminated is not reciprocal. Backbones typically require that other backbones meet certain minimum levels of traffic before they will peer with them. Smaller backbones unable to meet these requirements will have to pay transit fees. In addition

These business relationships and pricing structures suffered from a number of drawbacks. All-you-can-eat pricing charged to end users encouraged congestion. [Provide further explanation.] Many regional ISPs paid significant amounts in transit fees. The uniqueness of each connection between any two points made guaranteeing network reliability more difficult, as the failure or congestion of any particular network component would degrade the network's quality of service. Moreover, the fact that no money changes hands when backbones interconnect with one another makes it much more difficult for network participants to compensate each other value that is provided by another network participant on the other side of a peering point. Instead of having that value internalized into the network transactions, peering forces the network participants to use some other alternative institutional arrangement aside from price to replace those price signals.

In recent years, network providers have been experimenting with a wide range of alternative business relationships to mitigate these problems. For example, regional ISPs have begun to affiliate with more than one backbone through a practice called "multihoming." Regional ISPs that were too small to peer with the backbones avoided paying transit charges by entering into agreements to exchange traffic directly with one another on a settlement free basis, a practice which has become to be known as "secondary peering." These changes yielded multiple benefits. The avoidance of transit charges reduced the costs borne by end users. The creation of alternatives to backbone services has made the Internet less hierarchical and weakened the position of top-level backbones. Secondary peering and multihoming also made the network more robust by creating additional paths connecting particular points. Indeed, studies estimate that as much as seventy percent of the endpoints in the internet can interconnect without passing through the public backbone.[29]

Another particularly innovative is content delivery networks like Akamai, which reportedly handles more than fifteen percent of the world's web traffic.[30] Akamai caches web content at over fourteen thousand locations throughout the Internet. When an end user sends a request for a webpage, the last-mile broadband provider checks to see whether that webpage is hosted by Akamai. If so, the last-mile provider redirects the query to a local cache maintained by Akamai.

The sheer number of caches all but guarantees that the closest Akamai cache will be located closer to the end user than the server hosting the primary webpage. As a result, content served by Akamai is less likely to be plagued by problems of latency. It also reduces network costs. For example, suppose that multiple end users located in Philadelphia wanted to access content hosted by CNN.com. Rather than burdening the network with multiple requests to CNN.com's servers, Akamai can satisfy all of these requests with making a single query to

---

[29] See *id*.

[30] *See* Yoo, *supra* note **Error! Bookmark not defined.**, at 1882-83.

CNN.com's server and then satisfying all of the subsequent requests solely with local resources. In addition, the redundancy in Akamai's server network not only insulates the content Akamai hosts from denial of service attacks. It also allows the system to redirect requires to other caches when particular caches are overly congested. To the extent that Akamai uses connects to its caches through private connections, content hosted by Akamai can bypass the public backbone altogether. Akamai thus represents a creative way to use storage as a substitute for networking resources. In commercializing its service, Akamai also offers a way to transfer money across a peering point within the context of the business relationships that comprise the Internet itself.

In addition to establishing new types of business relationships, network providers are experimenting with different types of pricing regimes within existing business relationships. For example, last-mile providers are moving away from all-you-can-eat pricing for end users and are experimenting with different forms of usage sensitive pricing.[31] In addition, the pricing relationships between network providers have begun to move away from the polar extremes of peering and transit and begun to engage a much wider array of interconnection arrangements.

Consider the institution of peering. The decision to exchange symmetric traffic on a settlement free basis presumes that both parties on either side of the peering point incur similar costs to originate and terminate traffic. The problem is that this is often not the case. In fact, last-mile providers can be divided into what some commentators have called *content* networks and *eyeball* networks.[32] Content networks are last-mile providers like Abovenet and Cogent that primarily serve large content providers, such as Google and Yahoo. Eyeball networks are last-mile providers like Comcast and Verizon that primarily serve end users.

31

32 Faratin.

There are a number of differences between content and eyeball networks that lead to significant asymmetries in cost. For example, web browsing typically requires only a small amount of traffic flowing toward the content network, in this case the uniform record locator (URL) of the website being accessed. Since the traffic being returned contains the content requested, it is likely to be much larger. In addition, establishing a content network is relatively cheap, requiring the establishment of high-speed connections to a relatively small number of locations. The costs of establishing an eyeball network are much higher, requiring the wiring of entire neighborhoods and requiring certain economies of density in order to be profitable.

In addition, the business models that have come to dominate the Internet have increasingly emphasized advertising revenue. Advertising revenue tends to flow in to content and application providers. The value of the network to advertisers, however, depends largely on the number of end users they can reach. Thus, the Internet represents the classic case of a two-sided market. [Explain better.]

With the costs being higher for eyeball networks than content networks, with advertising revenue flowing primarily to content rather than eyeball networks, and with the magnitude of that advertising revenue depending on the ability to reach a large number of end users, many regard it natural for cash to flow from the content networks to the eyeball networks. Indeed, this is precisely the situation that has obtained for years in the broadcast television industry. [Explain.] This accords with the insights of two-sided markets, which suggests that under these circumstances, it is often optimal for content networks to subsidize eyeball networks.

As a result, network providers who otherwise have the symmetry of traffic volume to peer with one another are increasingly entering into "paid peering" relationships. Other providers are entering into "partial transit" arrangements, in which a network provider would pay for the ability to hand off traffic to another network provider without undertaking the obligation to receive traffic in return.

The net result is to increase the complexity of both the network architecture and the economic relationships governing them. The particular cost will depend in no small part on the precise location of the relevant endpoints in the overall topology of the Internet. For example, an end users accessing content hosted by a regional ISP with which its ISP has a secondary peering relationship will save the transit costs associated with accessing that content through the public backbone. Indeed, accessing that content will be cheaper than accessing content provided by a similarly situated content provider residing on a regional ISP not covered by a secondary peering relationship. Similarly, content hosted on Akamai will generally be delivered more quickly and more reliably than content hosted elsewhere. These deviations from the seamless web of complete equality and interoperability are the product of the natural effort to economize on costs and to find creative ways to organize business relationships.

E. The Shift from Person-to-Person to Mass Communications

Although the early Internet did some contain some community-oriented applications such as bulletin boards, the dominant applications (e-mail and file transfer) focused on person-to-person communications. The emergence of webpages have increasingly turned the Internet into a medium for mass communications. The emergence of the Internet as a platform for disseminating video content will make this increasingly true in the future.

At first, many scholars predicted that the Internet would provide the means for individuals to speak directly to mass audiences without having to deal with the media companies who served as intermediaries.[33] What we have found is that we depend on search engines, bloggers, e-mail newsletters, and a host of other intermediaries to sort through the avalanche of content that ever-growing avalanche of content available on the web. Indeed, the Supreme Court's decisions on broadcast and cable television have repeatedly emphasized how this type of exercise of editorial discretion serves important free speech values.

It is thus not a question of whether, but which, entity will serve as the intermediary. Indeed, many current policy disputes can be understood as competition among search engines, software designers, and network providers in an attempt to serve that role.

F. The Maturation of the Industry

Some of the changes of the Internet involve the maturation of the industry. Some observers bemoan the fact that garage innovators have found it increasingly difficult to break through.

Much of this should be unsurprising. The early days of any industry are typically wide open and often witness the emergence of new startups into major players. As industries mature, it is not unusual for the hurdle for startups to become much higher.

In addition, as the industry matures, the stakes have become higher. The strategy of simply leveraging the existing infrastructure is rapidly nearing the end of its usefulness. Soon further increases in bandwidth will require more substantial capital investments.

In addition, the nature of competition has changed. It began as a race to serve new customers who are currently not being served. As broadband adoption approaches saturation, firms increasingly have to compete not by getting new customers, but rather by providing greater

---

33 *See, e.g.*, THOMAS L. FRIEDMAN, THE WORLD IS FLAT: A BRIEF HISTORY OF THE TWENTY-FIRST CENTURY 102 (2005); Eugene Volokh, *Cheap Speech and What It Will Do*, 104 YALE L.J. 1805 (1995).

value to the customers they already have. This shift from extensive to intensive competition marks a sea change in the history of the Internet. It bears noting that intensive competition is much riskier and that firms all too often turn to the government to insulate them from those risks.

II. POLICY IMPLICATIONS

The failure to appreciate these insights has caused the policy debate to overlook some important implications.

A. Changes in the Optimal Level of Standardization

The increasing heterogeneity of consumer demand has potentially significant implications for Internet policy. Network economic effects are often oversimplified to focus solely on the fact that the value of the network may increase with the number of users reachable through the network. A closer reading of the economic literature reveals that network economic effects only represent one of many competing considerations that affect optimality. Indeed, standardizing on any particular protocol inevitably necessitates a tradeoff. It creates value for consumers by increasing the size of the network to which they belong. At the same time, it reduces value for some consumers by requiring them to forego using a different protocol that would yield additional consumer benefits. Whether standardization is optimal depends on which of these two effects dominates. If everyone wants the same thing from the network, one would expect the former to dominate the latter. As consumer preferences become increasingly heterogeneous, at some point one would expect an equilibrium with different sets of users using different noninteroperable protocols to become both stable and optimal. Although many commentators refuse to consider the possibility, under certain conditions the optimal number of networks may be greater than one.

The changes in the network architecture that many commentators decry as an attempt by network providers to exercise monopoly power may thus be nothing more than network providers' attempt to satisfy consumer demand. Indeed, one would expect different network providers begin to tailor their offerings to different subsets of consumers. These changes may thus reflect nothing more than the natural economic response to increasing heterogeneity of consumer preferences. Indeed, targeting specific subsets of end users can increase the competitiveness of the Internet by making it easier for more network providers to survive in an industry characterized by both demand-side and supply-side economies of scale. Conversely, standardization limits the dimensions along which network providers can compete to price and networks size, which are considerations that inherently favor the largest players. Thus, mandating adherence to a standardize protocol in the name of protecting consumers may actually have the effect of harming them instead.

B. Experiments with Different Technological Solutions

Different network providers are experimenting with new solutions to the technological challenges that confront them. Network providers are pursuing a number of strategies to meet this rapidly increasing demand. Some are responding simply by building bigger pipes. The best example of this strategy is Verizon, which is investing $23 billion to make its fiber-based FiOS system.

Another alternative to building bigger pipes is to engage in some form of network management that protects quality of service by giving higher priority to traffic associated with applications that are particularly sensitive to delay. This would allow network owners to forego investments in excess capacity to protect against the degradation of quality of service during short-term traffic surges. In a world in which the relative cost of each solution is constantly changing, there is a good argument against regarding either alternative as being off the table.[34] Preventing networks from prioritizing will make buildouts more costly, which in turn will exacerbate the digital divide and adversely affect rural buildouts. It may also represent the only solution when additional bandwidth is simply not available.

Indeed, the FCC explicitly recognized the benefits of prioritization when establishing a public/private partnership to govern the public safety spectrum allocated through the 700 MHz auction. The FCC's solution was to allow public safety and commercial traffic to share bandwidth, but to give the former priority over the latter when the network becomes congested. As the FCC noted, such prioritized sharing should "both help to defray the costs of build-out and ensure that the spectrum is used efficiently."[35] This decision acknowledged that prioritization of higher value traffic represents an effective way to lower the cost of providing service, while at the same time representing a creative solution to extant bandwidth limitations. Furthermore, a recent study estimated that a network without prioritization might have to maintain up to 60 percent more capacity than a network offering prioritized service, assuming that links operate on average at 60 percent of capacity and that delay-sensitive traffic represents 20 percent of all traffic.[36]

A good example of this strategy is AT&T, which has committed $6.5 billion to deploy its new uVerse network based on a telephone-based technology known as VDSL to 60% of its

---

34 See Yoo, 19 Harv J L & Tech at 22–23, 70–71 (cited in note **Error! Bookmark not defined.**) ("[O]ne would expect that the relative costs of different types of solutions to change over time. Sometimes increases in bandwidth would be cheaper than reliance on network management techniques, and vice versa. It would thus be short-sighted to tie network managers' hands by limiting their flexibility in their choice of network management solutions.").

35 Service Rules for the 698–746, 747–762 and 777–792 MHz Bands, Second Report and Order, 22 FCC Rec 15289, 15431 ¶ 396 (2007).

36 See Joseph D. Houle et al, *The Evolving Internet—Traffic, Engineering, and Roles* 5 (paper presented at the 35th Annual Telecommunications Policy Research Conference), available at <http://www.cse.unr.edu/~yuksem/my-papers/2007-tprc.pdf> (last visited Apr 10, 2008).

service area. Rather than undertaking investments of the scale of Verizon's FiOS network, AT&T is attempting to leverage the existing infrastructure. The problem is that the bandwidth provided by VDSL is quite limited. The only way they can provide adequate voice, video, and data services is to reserve bandwidth for video.

The wireless industry provides an excellent example. As noted earlier, the limited amount of available bandwidth and wireless technologies' susceptibility to congestion represent particularly difficult challenges. Different networks are experimenting with different solutions. Some networks are experimenting with smarter and more sophisticated devices that give higher priority to applications such as voice, which are very sensitive to delays of as little as a fraction of a second, over applications such as e-mail, for which delays of a fraction of a second are essentially unnoticeable. This requires a very tight integration of network, devices, and applications in ways that may limit the network's interoperability. Other providers are going to the opposite extreme by making the devices attached to the network less sophisticated rather than more sophisticated and instead offloading a greater number of functions in to the network itself.

Perhaps most interesting are the providers using unlicensed spectrum to provide wireless broadband in rural areas. These providers are under such extreme bandwidth constraints that they openly throttle video for the simple reason that the presence of even a handful of video downloaders would render the service unusable for every customer.

These examples underscore that the ability to prioritize traffic and the viability of new last-mile providers are deeply linked. It is hard to know which strategy will prevail and how things might change in the future. The uncertainty the eventual proportion of client-server and peer-to-peer traffic complicates decisions about bandwidth deployment still further.

Network providers must thus make decisions that involve difficult tradeoffs based on their best guess of what the future will bring. These considerations underscore the problems associated with any one-size-fits-all solution to the Internet. The network now consists of very different transmission technologies, each of which is susceptible to different problems and different solutions. In addition, the number of potential solutions is vast, including building additional bandwidth, storing content locally, and network management.

The difficulty of anticipating which of these solutions will prove best in each context is underscored dramatically by the AOL-Time Warner merger. When it was announced in 2001, many regarded the "walled garden" approach in which AOL gave preferential treatment to its own propriety content as a profound threat to the Internet. These threats never materialized, demonstrated most eloquently by Time Warner's recent announcement that it was selling off AOL at a loss of $200 billion.

My second point is to draw on the lessons of past efforts to implement access mandates similar to network neutrality. Past regulatory efforts have found that such interconnection and nondiscrimination mandates only work when the interface and the product being regulated is relatively simple. As the Supreme Court recognized in its *Trinko* decision, the situation is quite different when the interface is complex. When that is the case, disputes over access are likely to be "highly technical" and "extremely numerous, given the incessant, complex, and constantly changing interaction" between providers."[37] Thus, in order to protect against "death by a thousand cuts," any regulator would have to undertake comprehensive oversight of essentially all facets of the business relationship between the parties. The challenge of doing so would be particularly demanding in industries like broadband, which are undergoing rapid technological

---

[37] Verizon Commc'ns Inc. v. Law Offices of Curtis V. Trinko, LLP, 540 U.S. 398, 414 (2004).

change.[38] This has led many commentators to conclude that any attempts to mandate access to such complex technologies are likely to prove futile.[39] Indeed, past efforts to impose similar access regimes, such the controversy over protocol conversion and vertical switching services under the Computer Inquiries, leased access to cable television networks, and unbundled access to network elements under the 1996 Act, have become bogged down in incessant controversies and litigation.

These problems demonstrate the potential dangers of regulatory intervention and underscore the importance of making sure that the scope of intervention is commensurate with the scope of the problem. It bears noting that the OECD,[40] the FCC (on multiple occasions over the past two and one half years),[41] the Justice Department,[42] the FTC,[43] and leading Internet gurus David Farber and Bob Kahn[44] have concluded that the factual record does not justify the type of regulatory intervention that network neutrality proponents seek. The FCC's current

---

38 Christopher S. Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1, 39-45 (2005), *available at* http://ssrn.com/abstract=742404; Christopher S. Yoo, *Network Neutrality and the Economics of Congestion*, 94 GEO. L.J. 1847, 1896-97 (2006), *available at* http://ssrn.com/abstract=825669.

39 *See, e.g.*, Paul L. Joskow & Roger G. Noll, *The Bell Doctrine: Applications in Telecommunications, Electricity, and Other Network Industries*, 51 STAN. L. REV. 1249 (1999); Gerald R. Faulhaber, *Policy-Induced Competition: The Telecommunications Experiments*, 15 INFO. ECON. & POL'Y 73 (2003).

40 OECD Report, Internet Traffic Prioritisation: An Overview 5 (Apr. 6, 2007), *available at* http://www.oecd.org/dataoecd//43/63/38405781.pdf.

41 AT&T Inc and BellSouth Corp Application for Transfer of Control, Memorandum Opinion and Order, 22 FCC Rcd 5662, 5724-27 ¶¶ 116-20 & n 339, 5738-39 ¶¶ 151-53 (2007); Applications for Consent to the Assignment and/or Transfer of Control of Licenses, Adelphia Communications Corporation, Assignors, to Time Warner Cable Inc, Assignees, et al, Memorandum Opinion and Order, 21 FCC Rcd 8203, 8296-99 ¶¶ 217-23 (2006); Verizon Communications, Inc and MCI, Inc Applications for Approval of Transfer of Control, Memorandum Opinion and Order, 20 FCC Rcd 18433, 18507-09 ¶¶ 139-43 (2005); SBC Communications, Inc and AT&T Corp Applications for Approval of Transfer of Control, Memorandum Opinion and Order, 20 FCC Rcd 18290, 18366-68 ¶¶ 140-44 (2005); Appropriate Framework for Broadband Access to the Internet over Wireline Facilities, Report and Order and Notice of Proposed Rulemaking, 20 FCC Rcd 14853, 14904 ¶ 96 (2005).

42 Ex parte Filing of the Department of Justice, Broadband Industry Practices Before the FCC, WC Docket No. 07-52 (filed Sept. 6, 2007), *available at* http://www.usdoj.gov/atr/public/comments/225767.pdf.

43 Federal Trade Commission, Staff Report on Broadband Connectivity Competition Policy 10, 11 (June 2007), *available at* http://www.ftc.gov/reports/broadband/v070000report.pdf.

44 David Farber & Michael Katz, *Hold Off on Net Neutrality*, WASHINGTON POST, January 19, 2007, at A19; Andrew Orlowski, *Father of Internet Warns Against Net Neutrality*, THE REGISTER, Jan. 18, 2007, *available at* http://www.theregister.co.uk/2007/01/18/kahn_net_neutrality_warning/ (quoting co-developer of TCP/IP Robert Kahn).

Notice of Inquiry was haled as an opportunity for network neutrality proponents to demonstrate the types of harms wrought by the absence of mandated network neutrality.[45] The proceeding only turned up a few isolated instances that do not appear to support broadscale regulatory intervention.[46]

On the other hand, the Internet has a long history of adjusting to these types of problems by itself. Indeed, many examples to which network neutrality proponents point, such as network providers' initial resistance to virtual private networks (VPNs) and home networking equipment such as WiFi routers, are better regarded examples of how the private decisions of consumers and network providers can solve such problems without regulatory intervention. Comcast's recent accommodation of BitTorrent and Pando and Verizon's recent commitment to open networks represent more recent examples of the same phenomenon.

The better solution is to pursue what I have called "network diversity," in which different providers are permitted to experiment with different approaches and to let the choices of consumers control the ultimate outcome.[47] A case-by-case, after-the-fact approach would appear to strike a better would balance that preserves room for experimentation while simultaneously ensuring that any problems that may emerge will be addressed. The FCC's enforcement action

---

45 Broadband Industry Practices, Notice of Inquiry, 22 FCC Rcd 7894 (2007).

46 *See* Kara Rowland, *FCC Set for Airwaves Auction*, WASH. TIMES, Jan. 16, 2008, at C8 (quoting FCC Chairman Kevin Martin as calling network neutrality regulation unnecessary).

47 Yoo, *Beyond Network Neutrality*, *supra* note 38; *Network Neutrality and the Economics of Congestion*, *supra* note 38; *Christopher S. Yoo, Would Mandating Broadband Network Neutrality Help or Hurt Competition? A Comment on the End-to-End Debate*, 3 J. ON TELECOMM. & HIGH TECH. L. 23 (2004), *available at* http://ssrn.com/abstract=495502.

against Madison River[48] and Chairman Kevin Martin's recent testimony before the Senate Commerce Committee[49] attest to the agency's readiness to play this role.

D. Increasing Diversity in Business Relationships

The changes discussed above are also likely to make Internet pricing to become more varied and complex. Such a development should be embraced, rather than resisted. Price mechanisms are essential for signaling the extent to which markets are in short-run disequilibrium and for providing the incentive to reallocating resources in a more efficient manner. Prices also provide the incentive for upgrading existing networks and for investing in new network technologies.

Interestingly, even though many new business arrangements are criticized as attempts for existing players to exercise monopoly power, the increasing diversity of business relationships and pricing regimes often serves to weaken rather than strengthen market power. In addition, moving away from the hierarchical nature of the previous regime makes the network more robust to system failure.

Technological solutions are complicated by the fact that the Internet is comprised of numerous autonomous actors whose actions can be hard to coordinate. The lack of coordination raises the real danger of strategic behavior. Indeed, the recent dispute between Comcast and BitTorrent can be understood as a struggle over whether optimization of traffic patterns in light of congestion will occur at the network level or the application level.

---

48 Madison River Commc'ns, LLC, Order, 20 FCC Rcd 4295 (2005).

49 Written Statement of Kevin Martin, Chairman, Federal Communications Commission, Before the United States Senate Committee on Commerce, Science and Transportation 4-5 (Apr. 22, 2008), *available at* http://hraunfoss.fcc.gov/edocs_public/attachmatch/DOC-281690A1.pdf.

C. The Possibility of Incomplete Convergence and the Myth of the One Screen

Much of the desire to preserve complete interoperability is the implicit belief in total convergence, i.e., that every end user will subscribe to a single network provider that will support all available communications. There are many reasons to expect that most consumers will maintain relationships with more than one network provider. Part of the motivation is technological. For example, wireless providers offer mobility features that wireline providers cannot. Part of the motivation is reliability. Indeed, an increasing number of end users are maintaining more than one connection to protect against network failure.

The likelihood of multiple providers broadens the policy space in important ways. Making sure that every possible source of content and applications is available through every possible network provider becomes less important. The possibility that consumers may subscribe to more than one network provider makes competition between large, vertically integrated providers more benign from the standpoint of competition policy, which in turn broadens the range of business strategies that industry players can pursue without raising policy concerns.

These insights support a range of broader conclusions. Adherence to the status quo is no longer tenable. At the same time, the changes discussed above indicate that one-size-fits-all solutions are probably no longer viable. Indeed, many possible outcomes may result from this competition, including ones in which the network remains unified on a different protocol as well as ones in which the network we now know as the Internet becomes balkanized into separate, noninteroperable networks. It is impossible to know at this point which will be the best solution. Any policy solution must be flexible enough to allow network providers to experiment with different approaches. Any constriction of this freedom threatens to choke off the process by which technological and economic progress occurs.

E. The Inevitability of Intermediation

As noted earlier, the proper question is not whether some actor will serve as an intermediary to help sift through the ever-growing avalanche of content confronting consumers every day. A long line of Supreme Court precedents in the context of broadcasting and cable television has long noted the critical role that editorial discretion plays when mass communications are involved. At the same time, these precedents underscore the danger of allowing the government to substitute its editorial judgment for that of private actors. Although the Supreme Court has recognized three justifications, specifically the scarcity doctrine, pervasiveness, and natural monopoly, that might justify some forms of regulation, the Court has already held the first inapplicable to the Internet and is currently reevaluating the second, while technology is effectively undermining the third. In an environment as dynamic as the Internet, it is difficult to predict the roles that search engines, browser manufacturers, content aggregators, and network operators might play in exercising that discretion. Moreover, technological change will no doubt require a constant readjustment of the roles played by these actors. Indeed, modern theories of administrative action suggest that government intervention would systematically overrepresent established players that already have revenue streams sufficient to support lobbying efforts and that have customer bases that can serve as constituencies for preserving the status quo. An alternative is to allow the various actors to offer different approaches to exercising the editorial discretion needed to manage the avalanche of content. Conversely, allowing regulation to predetermine or freeze into place the roles played by each actor carries with it considerable risk.

F. Pressures Toward Nonanonymity

Describe how the pressures of virus protection and spam are forcing the network toward nonanonymity.

G. The Maturation of the Industry

Many observers mourn the fact that it is more difficult for garage innovators to succeed on the Internet. As the recent emergence of MySpace reveals, it is far from clear that such success stories have disappeared entirely. That said, it should come as no surprise that entry may become more difficult as the wide-open nature of competition in a nascent industry is replaced by an environment requiring greater technical and business requirements as the industry matures. In this respect, the lessons of the first competitive era of telephone service that emerged after the initial Bell patents lapsed in the mid 1890s provide a useful example. In the early days of the competitive era, the presence of a large number of uncommitted customers made entry quite easy, as new entrants could survive simply by serving new customers. After the network had largely been built out, the nature of competition changed from extensive competition, in which competitors raced to extend the reach of the network, to intensive competition, in which networks competed directly with one another to provide greater value to consumers. Intensive competition is much riskier and much more difficult. Such phases of intensive competition also raise serious dangers that industry players will look to regulation as a form of flight from competition.

H. The Return of Conventional Economic Analysis

[Describe how the emphasis on the end-to-end argument has obscured the issues and how a return to conventional industrial organization principles would yield better results.]

CONCLUSION

These change make clear that policymakers cannot blindly adhere to the status quo.