WRITTEN TESTIMONY OF CHRISTOPHER S. YOO

Professor of Law and Communication
Co-Director, Center for Technology, Innovation, and Competition
University of Pennsylvania

Public *En Banc* Hearing on Broadband Network Management Practices

Before the Federal Communications Commission

Cambridge, Massachusetts

February 25, 2008

Chairman Martin and Members of the Commission, I am grateful for the opportunity to testify today on the subject of broadband network management practices.

Traffic on the Internet is undergoing explosive growth, but reports conflict as to just how explosive this growth will continue to be in the future. Some estimate that meeting this increase in demand will require more than $100 billion in new investments in network infrastructure.[1] Other observers project growth rates that are somewhat slower, but still nonetheless quite substantial.[2]

This uncertainty about the rate of traffic growth poses a dilemma for network providers that must undertake investments to meet this demand. If they follow the higher estimates, they may end up investing tens of billions of dollars in unnecessary network capacity in the areas that they serve. Following such an approach would slow national broadband deployment in higher-cost areas by taking up scarce capital and by increasing the breakeven number of customers

---

[1] *See, e.g.*, Nemertes Research, Press Release, User Demand for the Internet Could Outpace Network Capacity by 2010 (Nov. 19, 2007), *at* http://www.nemertes.com/press_releases/user_demand_internet_could_outpace_network_capacity_2010; Bret Swanson & George Gilder, *Estimating the Exaflood: The Impact of Video and Rich Media on the Internet* 12-14 (Jan. 2008), *available at* http://www.discovery.org/scripts/viewDB/index.php?command=view&id=4428.

[2] *See, e.g.*, Univ. of Minn. Digital Tech. Ctr., Minnesota Internet Traffic Studies (Aug. 30, 2007), *available at* http://www.dtc.umn.edu/mints/home.html [hereinafter Odlyzko Study].

needed for broadband service to be viable in any particular area. If they follow the lower estimates, they risk seeing congestion cause their networks slow to a crawl.

In addition, their investment decisions must be based not only on how large they estimate the total traffic will be, but also on the location of the traffic, its variability, and the upload-download pattern that the traffic will take. Any such estimate must take into account the technological differences in the various last-mile transmission technologies as well as projected developments in complementary technologies, such as RAM, hard disk, and chip technologies, which can have a dramatic effect on the demand for network services.

The difficulty in making these assessments is underscored by an article in yesterday's *New York Times*, which detailed how breakthroughs in nanotechnology may soon make it possible to manufacture semiconductor chips that can achieve unprecedented speeds.[3] Under these circumstances, even the best projections will occasionally miss the mark, and the fact that network capacity cannot be expanded instantaneously necessarily means that any underestimation would lead to congestion that degrades network performance.

That said, practices exist that can alleviate the harm caused by any underestimation that may occur. Specifically, network management can constitute an important safety valve that can alleviate network congestion when expanding capacity is not an option. In this sense, capacity expansion and network management are more properly regarded as alternative approaches to deal with the problem of congestion. Which will be preferable in any particular case will vary with the circumstances and with their relative costs. It is difficult, if not impossible, to determine *a priori* which will prove to be the better solution.

---

[3] *See* Michael Fitzgerald, *Trying to Put New Zip into Moore's Law*, N.Y. TIMES, Feb. 24, 2008.

It is for this reason that I believe that network management has an important role to play in building the Internet of the future. This claim is not merely theoretical. Indeed, the history of FCC regulation is replete with examples where network owners have used prioritization and network management to protect end users and to preserve network performance. Perhaps the most telling example is NSFNET's 1987 decision to respond to the surge in traffic caused by the arrival of the personal computer by prioritizing certain traffic over others.[4] The appropriateness of prioritization was further underscored by the FCC's decision last fall with respect to the public safety spectrum allocated by the 700 MHz auction, which recognized that a shared network in which higher value uses are given priority can reduce the costs of deployment and make more efficient use of the scarce bandwidth that is available.[5]

That said, the theoretical possibility exists that a network owner might degrade or block traffic to favor its own proprietary offerings. Should the risk be sufficiently great or the practice be sufficiently widespread, more intrusive regulatory intervention might be appropriate. Indeed, when the Commission identified such an abuse by Madison River Communications, it acted swiftly to address the problem.[6] Since that time, this Commission has concluded on five separate occasions over the past two and a half years (including the *Wireline Broadband Internet Access Services Order*,[7] the SBC-AT&T order,[8] the Verizon-MCI order,[9] the Adelphia order,[10] and the

---

4 *See* Christopher S. Yoo, *Beyond Network Neutrality*, 19 HARV. J.L. & TECH. 1, 22-23 (2005).

5 *See* Service Rules for the 698-746, 747-762 and 777-792 MHz Bands, Second Report and Order, 22 FCC Rcd 15289, 15431 ¶ 396 (2007) [hereinafter 700 MHz Second Report].

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

7 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).

8 SBC Communications, Inc. and AT&T Corp. Applications for Approval of Transfer of Control, Memorandum Opinion and Order, 20 FCC Rcd 18290, 18366-68 ¶¶ 140-144 (2005).

9 Verizon Communications, Inc. and MCI, Inc. Applications for Approval of Transfer of Control, Memorandum Opinion and Order, 20 FCC Rcd 18433, 18507-09 ¶¶ 139-143 (2005).

10 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 ¶¶ 217-223 (2006).

AT&T-BellSouth order[11]) that there was insufficient evidence of degradation and blocking to justify regulatory intervention, and the ongoing Notice of Inquiry[12] has identified only a handful of isolated instances. The case-by-case approach that this Commission adopted in *Madison River* would thus appear more than adequate to address. The potential benefits of the many possible alternatives suggest that a more categorical approach may prove costly in ways that are impossible to anticipate.

I. NETWORK CONGESTION IS A COMPLEX AND GROWING PROBLEM

As noted earlier, congestion has become a major problem on the Internet, and the pace of traffic growth has led some to suggest that it is likely to become even worse in the future. Although the increase in traffic is an important development, as I have detailed in my previous work, congestion is a complex phenomenon that on more than just total volume.[13] It also depends on the timing, location, and pattern of overall network traffic. In addition, networks ability to compensate for increases in demand by rerouting traffic can make network performance quite unpredictable. Thus, a disruption in one portion of the network can increase congestion in areas of the network located far from the point of disruption.

A The Emergence of New Bandwidth-Intensive Applications Is Causing a Rapid Increase in Network Traffic

Network usage is exploding, driven largely by the growth of new, bandwidth-intensive applications (most notably video). This has led some observers warn of an "exaflood," in which

---

[11] AT&T Inc. and BellSouth Corp. Application for Transfer of Control, Memorandum Opinion and Order, 22 FCC Rcd 5662, 5724-27 ¶¶ 116-120 & n.339, 5738-39 ¶¶ 151-153 (2007).

[12] *See* Broadband Industry Practices, Notice of Inquiry, 22 FCC Rcd 7894 (2007).

[13] *See* Daniel F. Spulber & Christopher S. Yoo, *On the Regulation of Networks as Complex Systems: A Graph Theory Approach*, 99 NW. U. L. REV. 1687 (2005); Christopher S. Yoo, *Network Neutrality and the Economics of Congestion*, 94 GEO. L.J. 1847 (2006).

Internet traffic will accelerate from its current growth rate of 50%-70% to 90%, which would cause Internet traffic to grow by 50 times between 2006 and 2015.[14] Another study indicates that average international backbone traffic grew at a 75% annual rate and peak traffic grew at 72% annual rate between 2004 and 2006, while bandwidth increased only at annual rate of 45%.[15] Even network neutrality proponents, such as Google and EDUCAUSE, have warned that the Internet will struggle to accommodate consumers' increasing demands for bandwidth.[16] 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 invests $137 billion in upgrading the Internet infrastructure.[17] Business executives have predicted annual growth rates ranging from 100% to 500%.[18]

Others have expressed skepticism about these claims. Cisco predicts that Internet traffic will grow at a 46% annual rate between 2007 and 2011,[19] which is a rate comparable to the growth of capacity. Similarly, Andrew Odlyzko of the University of Minnesota similarly concludes that the Internet is growing an annual rate of roughly 50%-60%.[20] Indeed, in 2007,

---

[14] *See* Swanson & Gilder, *supra* note 1, at 8, 22; *see also* Bret Swanson & George Gilder, *Unleashing the "Exaflood"*, WALL ST. J., Feb. 22, 2008, at A15.

[15] *See* TeleGeography Research, Global Internet Geography 3-4 (2007), *available at* http://www.telegeography.com/products/gig/samples07/GIG_Exec_Summary.pdf.

[16] *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.

[17] *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.

[18] *See* Odlyzko Study, *supra* note 2.

[19] *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.

[20] *See* Odlyzko Study, *supra* note 2.

bandwidth growth in the international backbone (68%) exceeded both average traffic growth (57%) and peak traffic growth (60%) for the first time in years.[21]

Sorting out the precise import of these studies is the fact that many (although not all) of them were produced by industry players or consultants who have a stake in particular outcomes. In addition, growth rates have not been constant over time. As Odlyzko notes, traffic did grow at an annual rates that exceeded 100% from the early 1990s until about 2002, and that it is quite possible that such rates could recur.[22]

Most tellingly, both sides acknowledge that demand is growing rapidly and that the network needs a massive investment in upgrading its infrastructure in order to prevent congestion from overwhelming the Internet. Uncertainty about the precise magnitude of that growth in demand, and thus the precise magnitude of the additional capacity needed to satisfy it, in essence forces network owners to confront a multibillion-dollar gamble about how much capacity to add.

B. Congestion Is Determined Not Just by Total Volume, But Also by the Timing and Location of Traffic and the Configuration of the Network

The problem facing network owners is compounded by the fact that congestion is not determined simply by the total volume being carried on the Internet. It is also determined by the timing with which that volume is introduced into the Internet. Users that consume a large amount of bandwidth may contribute little to network congestion if they do so at a time when overall network usage is low. Conversely, users that only consume a small amount of bandwidth

---

21 *See* TeleGeography Research, *supra* note 15, at 3-4.

22 *See* Univ. of Minn. Digital Tech. Ctr., Internet Growth Trends and Moore's Law (Aug. 30, 2007), *available at* http://www.dtc.umn.edu/mints/igrowth.html; Odlyzko Study, *supra* note 2.

may nonetheless cause a great deal of congestion if they do so at a time when overall network usage is high.

In addition, congestion is also function of location. In short, the level of congestion often depends as much as the number of close neighbors who also happen to be using the network as it does on total amount of traffic passing through the network at any time. The reality is that a large amount of traffic is generated by a small number of superheavy users whose usage is predominantly driven by filesharing. Although this is sometimes called the "80/20 rule," under the assumption that 20% of the users generate 80% of the traffic,[23] this appellation appears to be something of an understatement. The record indicates as few as 5% of end users may be responsible for generating from between 50%-80% of all Internet traffic.[24] Indeed, studies indicate that congestion becomes problematic when as few as 15 of the 500 or so end users sharing a fiber node attempt to run peer-to-peer filesharing programs through the same node.[25] Thus, any solution to congestion pricing must be able to account for differences in both time and space.

Furthermore, networks have the ability to respond to congestion by rerouting traffic along different paths. Although redirecting traffic in this manner can alleviate congestion in one part of the network, in so doing it can increase congestion in other parts of the network. This process of adjustment and readjustment can cause congestion to be transferred from one portion of the

---

[23] *See* SANDVINE, INC., NETWORK NEUTRALITY: A BROADBAND WILD WEST? 4 (Mar. 2005), *available at* http://www.sandvine.com/general/getfile.asp?FILEID=37; Martin & Westall, *supra* note 25, at 1.

[24] *See* David Vorhaus, Confronting the Albatross of P2P, at 1 (Yankee Group, May 31, 2007) (noting that 5% of users account for 50% of all traffic); Steven Levy, *Pay per Gig*, WASH. POST, Jan. 30, 2008, at D1 (quoting Time Warner Cable spokesman offering similar statistics); CTIA Comments at 4 (noting that one wireless provider reports that less than 4% of customers generate over 50% of traffic), *available at* http://gullfoss2.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6519841180.

[25] *See* James J. Martin & James M. Westall, *Assessing the Impact of BitTorrent on DOCSIS Networks*, *in* PROCEEDINGS OF THE 2007 IEEE BROADNETS (Sept. 2007), *available at* http://people.clemson.edu/~jmarty/papers/bittorrentBroadnets.pdf; *see also* Leslie Ellis, *BitTorrent's Swarms Have a Deadly Bite on Broadband Nets*, MULTICHANNEL NEWS, May 8, 2006, *available at* http://www.multichannel.com/article/CA6332098.html.

network in ways that vary with the particular configuration of the network and the amount of traffic being carried in each portion of the network. This can cause congestion to arise in areas located far away from the place where the new traffic is entering the network in ways that can be quite unpredictable.[26]

Network owners thus confront a problem. Network demand is growing at a rate that demands massive investments to upgrade the network. At the same time, there is a great deal of uncertainty over how large of an upgrade is really needed as well as how the network should be configured. Equally importantly, network owners face considerable uncertainty over the mix of traffic types that will comprise that volume. Unfortunately, all of this uncertainty affects decisions about network design as well as the total capacity to be deployed. In a world of limited capital, deploying extra capacity to reflect the higher estimates necessarily means reducing the area in which broadband service will be built out. Deploying too little capacity, in contrast runs the risk of severe degradation of service.

C. The Decision Whether to Design Networks for a Client-Server or a Peer-to-Peer Architecture

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.

---

[26] *See* Spulber & Yoo, *supra* note 13, at 1703-13.

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).[27]

---

[27] *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].

More recent versions of ADSL support higher bandwidth, but still allocates it asymmetrically.[28] The initial cable modem architecture, designed around DOCSIS 1.0, supported maximum theoretical speeds of 27 mbps downstream and 10 Mbps upstream[29] (with the actual downloads speeds reaching 6 Mbps[30]). 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.[31]

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 of "short-sighted[ness]" or "poor network design decisions."[32] 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

---

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

[29] *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).

[30] *See* Fourth Report, *supra* note 27, at 20554.

[31] *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.

[32] Comments of Free Press et al. at 21, 22.

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.[33] The ongoing transition of high definition television is likely to cause demand to increase still further.[34] 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.[35] Thus, in 2007, client-server traffic has retaken the lead from peer-to-peer, constituting 45% of all Internet traffic as compared with 37% of all traffic devoted to peer-to-peer.[36] Thus, some industry observers predict that video traffic will constitute over 80% of all Internet traffic by 2010.[37]

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.

---

33 *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.

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

35 *See* Swanson & Gilder, *supra* note 1, at 12-14.

36 *See* Ellacoya Networks, *supra* note 33.

37 *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

D. Differences Among Transmission Technologies in the Ways Congestion Can Arise

Not only must network owners must base their investment plans on their projections of the magnitude and the shape of the traffic that will have to support. Although all of the technologies are subject to congestion in content/application servers, backbone, and regional ISPs, they differ in the extent to which they are subject to congestion more locally. Indeed, the failure to consider these key differences reflects the extent to which thinking about broadband networks reflects the Internet's origin as a wireline technology connecting desktop personal computers.

Figure 1
Architectures of the Major Broadband Transmission Technologies

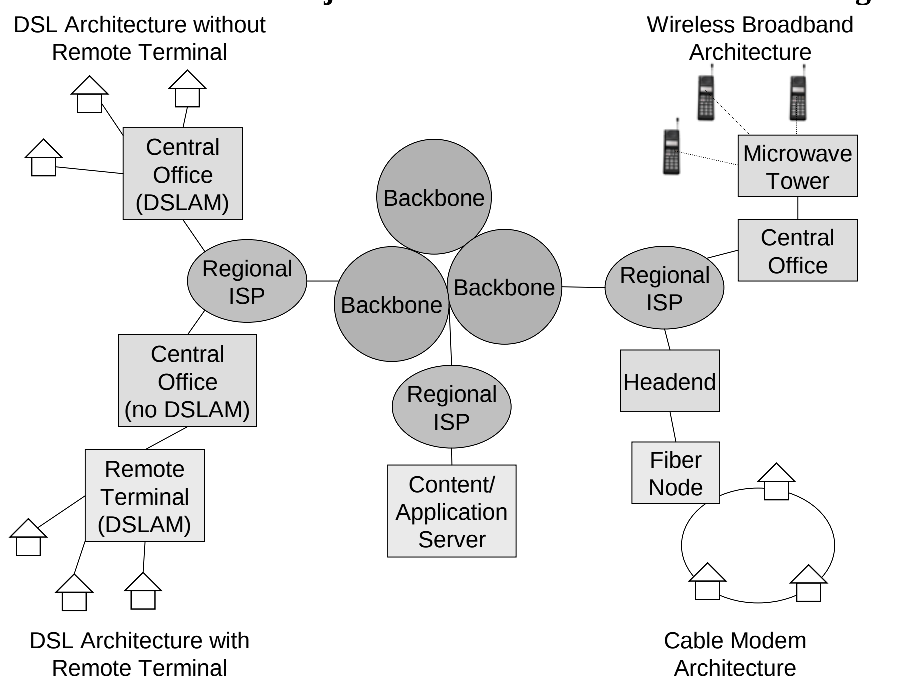

Consider first the architecture of DSL. DSL customers typically connect through a pair of copper wires to a major facility maintained by the local telephone company known as the central office, in which the telephone company maintains a piece of equipment known as a DSL access multiplexer (DSLAM) to separate the voice traffic from the data traffic. Because DSL customers connect to the DSLAM through a dedicated connection, their traffic is not typically aggregated with other traffic until it reaches the central office. As a result, the local connection between DSL customers' premises and the central office is not subject to congestion. The primary constraint is that modern ADSL can only serve customers located within 18,000 feet of a DSLAM. To serve these customers, local telephone companies sometimes deploy DSLAMs in satellite facilities known as remote terminals, which are in turn connected to the central office through optical fiber. Because DSL customers have dedicated connections to the DSLAM, their traffic is not aggregated with other traffic until it reaches the remote terminal. As a result, DSL customers do not share bandwidth with other customers in the link between their premises and the remote terminal, and thus that portion of the network is not subject to congestion.

The situation is quite different with respect to cable modem systems, which are based on a hybrid fiber coaxial (HFC) architecture. Under an HFC architecture, the copper coaxial cables connecting individual customers' premises are reconfigured into a ring configuration and connected to a satellite facility known as a neighborhood node. The node is in turn connected by optical fiber to the major cable facility known as the head end. Unlike under DSL, traffic generated by individual cable modem customers shares bandwidth with the traffic generated by their neighbors from the moment it leaves their house. As a result, the quality of service that any particular cable modem customer receives is considerably more sensitive to the bandwidth consumption of their immediate neighbors than is DSL.

The congestion problems confronted by wireless broadband providers are even more severe. Wireless broadband providers connect to the Internet through transponders located on microwave towers and other high-altitude locations. Because the capacity of any one transponder is limited, customers attempting to connect to the same tower compete for bandwidth with their neighbors. Thus, like cable modem service, wireline broadband service is sensitive to local congestion in addition to the other forms of congestion.

This problem is exacerbated in the case of wireless broadband by two other considerations. First, wireless broadband operates under bandwidth constraints that are much more restrictive than those faced by DSL or cable modem systems. Second, in DSL and cable modem systems, broadband traffic is carried in a different channel than traffic associated with the other services provided by the company. In the case of DSL, conventional voice traffic is transmitted through a different channel than data traffic. As a result, broadband traffic cannot degrade conventional voice traffic in a wireline network no matter how much it increases. Similarly, in a cable network, conventional video traffic is transmitted through a different channel than data traffic. Again, broadband traffic cannot degrade conventional video traffic in a cable network no matter how large it grows.

This is not true in the case of wireless. Wireless broadband shares bandwidth with the voice services offered by wireless companies. Consequently, any congestion that may arise in a wireless network degrades not only the quality of Internet broadband services provided; it also degrades the conventional voice services that represent the wireless providers' the core business.

It should thus come as no surprise that different types of providers vary in their tolerance for local congestion, with some taking more aggressive efforts to manage it and some taking less. It should also come as no surprise that different types of providers would manage congestion on a different geographic scale, depending on the nature of their technology. These technological realities caution strongly against adopting a one-size-fits-all approach to network management. Indeed, any regulatory solution that might be imposed must be carefully tailored to take these important variations into account.

II. DIFFERENT APPROACHES CAN PLAY A POSITIVE ROLE IN NETWORK MANAGEMENT

The rapid growth in the volume of network traffic is forcing network providers to undertake significant investments to expand network capacity.[38] As noted earlier, such investment decisions depend on forecasts of the expected volume growth, the precise timing and variability of network flows, the geographic locations of network growth, the extent to which it will consist of conventional downloads or peer-to-peer, and the technological limitations associated with different transmission technologies.

Faced with this complex reality, the Commission should expect network providers to adopt a variety of approaches to managing congestion. In addition, as I shall discuss in greater detail below, each of the available tools for network management are subject to its own strengths and weaknesses. Which mechanism will prove the best way to manage congestion will vary from case to case. This suggests that no particular mechanism for managing congestion should be discarded. Instead, this suggests that there is considerable benefit in preserving flexibility. Indeed, the presence of alternative institutional mechanisms for managing congestion serves as a valuable source of insurance against any misestimation by network providers of the nature of future network demand.

A. Building More Bandwidth

Network owners can respond to the increasing demand for network services and the emergence of applications that are increasingly bandwidth intensive and increasingly intolerant of delay in one of several ways.[39] Some have suggested that the preferred approach would be for

---

[38] The analysis in this Part is adapted from Christopher S. Yoo, *Network Neutrality, Consumers, and Innovation*, 2008 U. CHI. LEGAL F. (forthcoming 2008).

[39] *See* Yoo, *supra* note 4, at 22-23, 70-71.

network owners to meet these demands by expanding network capacity.[40] Perhaps most famously, Lawrence Lessig suggests that although congestion is a real problem, it can be solved by increasing bandwidth rather than by giving network owners more control over network flows. Although Lessig recognizes that this vision of a world with "infinite" bandwidth contradicts the basic economic notion that all commodities are inherently scarce, he nonetheless states, "I'm willing to believe in the potential of essentially infinite bandwidth. And I am happy to imagine the scarcity-centric economist proven wrong."[41]

As appealing as Lessig's vision may be to some, I find it to be problematic. The reason is that in the absence of some way to shape traffic on the network, the only way to preserve the quality of service to network customers is by maintaining excess capacity to serve as insurance against surges in demand. Such excess capacity can be quite costly. Studies that have attempted to quantify the size of the excess capacity needed to preserve quality of service estimate that a network without prioritization might have to maintain up to have to maintain up to 60% more capacity than a network offering prioritized service when traffic is moderate and 100% more capacity when traffic is heavy.[42] Other studies estimate that maintaining such excess capacity would cost consumers somewhere between $140 and $400 each month.[43] Increasing the capital costs of building out any particular geographic area also has the effect of increasing the number

---

40 *See, e.g.*, Petition of Free Press for Declaratory Ruling that Degrading an Internet Application Violates the FCC's Internet Policy Statement and Does Not Meet an Exception for "Reasonable Network Management," Petition for Declaratory Ruling at 26 (filed Nov. 1, 2007), *available at* http://fjallfoss.fcc.gov/prod/ecfs/retrieve.cgi?native_or_pdf=pdf&id_document=6519825121 [hereinafter Free Press Petition].

41 *See* LAWRENCE LESSIG, THE FUTURE OF IDEAS 47 (2001).

42 *See* Murat Yuksel et al, *The Value of Supporting Class-of-Service in IP Backbones*, PROCEEDINGS OF IEEE INTERNATIONAL WORKSHOP ON QUALITY OF SERVICE (June 2007), *available at* http://www.ecse.rpi.edu/Hompages/shivkuma/research/projects/cos-support.htm.

43 *See* Richard N. Clarke, Costs of Neutral/Unmanaged IP Networks (Aug. 2007) (placing the cost of excess capacity at between $143 and $416 per month), *available at* http://ssrn.com/abstract=903433; George Ford et al., The Efficiency Risk of Network Neutrality Rules (Phoenix Center Policy Bulletin NO. 16, May 2006) (placing the cost of excess capacity at between $300 and $400 per month), *available at* http://ssrn.com/abstract=925347.

of subscribers that the network provider must attract to break even, which in turn has the effect of limiting the ability to build out rural and other high-cost areas.[44]

Put a different way, adding capacity and using network management techniques represent alternative approaches to solving the problems of congestion, with the social desirability of one approach or the other depending on their costs. There would thus seem to be no reason to prefer one solution to the other *a priori*. Indeed, a recent study by two noted economists demonstrates how preventing network owners from employing multiple classes of service may be inefficient and reduce overall welfare.[45]

Moreover, it is inevitable that network demand will sometimes diverge from network providers' predictions. Although network owners are currently planning to expand capacity, the specific investments depend on their best judgment of how much additional bandwidth should be added and where, as well as the precise type of architecture needed to meet that demand. When the network owner overestimates future demand, quality of service does not suffer, although the additional investment would tie capital that would be better used building out less populous areas. Any underestimation of future demand, in contrast, causes network performance to degrade, and the fact that capacity cannot be added instantaneously necessarily dictates that network owners turn to some alternative approach to reducing congestion.

Network management thus has the potential to play two distinct roles. First, it can represent a way to respond to the increase in demand that is less expensive than adding capacity. Second, it can serve as an important backup plan that insures against the possibility that network traffic might grow more quickly than anticipated.

---

[44] *See* Embarq Comments at 7; LARIAT Comments at 2-4; National Telecommunications Cooperative Association Comments at 5-6.

[45] *See* Benjamin E. Hermalin & Michael L. Katz, *The Economics of Product-Line Restrictions With an Application to the Network Neutrality Debate*, 19 INFO. ECON. & POL'Y 215 (2007).

B. Metered Pricing

Some network neutrality proponents also advocate some form of metered pricing,[46] although other network neutrality proponents disagree.[47] Under this approach, network owners monitor every end users' and content/application providers' bandwidth usage and charge them prices that reflects their total bandwidth consumption. Indeed, this is precisely the pricing regime reflected in the contracts between network providers and content/application providers, which typically bases access charges on the total bandwidth consumed over a thirty-day period. Although contracts between network providers and end users have tended to employ unmetered, all-you-can-eat pricing, Time Warner has recently begun to experiment with usage-sensitive pricing.[48]

The classic solution to the problems of congestion is to charge network users a price that reflects the exact cost of congestion that they impose on other network users. Doing so would give end users the incentive to not to consume more network resources when doing so would be socially harmful.[49]

The problem with metered pricing based simply on the number of bits transmitted is that it ignores the fact that the amount congestion generated by any particular end user depends as much on the timing of that usage in relation to the usage patterns of all other end users as it does

---

46 *See* Free Press Petition, *supra* note 40, at 29; Brett Frischmann & Barbara van Schewick, *Net Neutrality and the Economics of the Information Superhighway: A Reply to Professor Yoo*, 47 JURIMETRICS J. 383, 395-97 (2007); Tim Wu, *Network Neutrality, Broadband Discrimination*, 2 J. ON TELECOMM. & HIGH TECH. L. 141, 154 (2003).

47 *See* Free Press, Press Release, Time Warner Metering Exposes America's Bigger Broadband Problems (Jan. 17, 2008), *available at* http://www.freepress.net/press/release.php?id=328; Time Warner Metered Pricing: Not the Solution (Jan. 17, 2008), *available at* http://www.savetheinternet.com/blog/2008/01/17/time-warner%e2%80%99s-metered-pricing-not-the-solution; Marvin Ammori, Time Warner Goes Back to the Future (Jan. 25, 2008), *available at* http://www.savetheinternet.com/blog/2008/01/25/back-to-the-future-time-warner-broadband-plan-recalls-aols-walled-garden; Fred von Lohmann, Time Warner Puts a Meter on the Internet (Jan. 22, 2008), *available at* http://www.eff.org/deeplinks/2008/01/time-warners-puts-meter-internet.

48 *See* Time Warner Comments at 24.

49 *See* Yoo, *supra* note 13, at 1864-65.

on the total amount of bandwidth consumed. For example, a heavy bandwidth user might impose minimal congestion if it confines its usage to times when few other users are using the network. Conversely, a light bandwidth user might nonetheless become a significant source of congestion should it choose to use the network at a time of heavy network usage. Thus, merely counting bits can represent a poor measure of congestion costs and thus may not provide sufficient incentive for individual end users to behave in a way that maximizes consumer welfare.

C. Time-of-Day Pricing

Some of the shortcomings of metered pricing can be addressed through the use of time-of-day or peak-load pricing.[50] Under this approach, individual end users face higher usage charges during those times of day when the overall network usage is likely to be highest. Indeed, peak-load pricing schemes should be quite familiar to those with long distance plans that lower rates on evenings, nights, and weekends and with wireless plans that offer free night and weekend minutes.

Regulators' experience with a form of pricing for telephone service known as "local measured service" serves as a cautionary tale about the potential shortcomings of time-of-day pricing. Like Internet access, local telephone service is typically priced on an all-you-can-eat basis. Some network elements, most notably the copper loop connecting individual customers to the telephone company's central office, are not shared with other end users and are thus not subject to congestion. Other network elements, such as switching, are shared with other customers and thus are subject to congestion.

---

50 *See* Frischmann & van Schewick, *supra* note 46, at 396, 405.

The presence of these congestible elements led many analysts to suggest that metering local telephone usage based on the time of day would yield substantial consumer benefits and convinced some local telephone companies to experiment with local measured service. Empirical studies indicate that local telephone companies' experiments with local measured service either harmed consumers or yielded benefits that were too small to cover the metering costs.[51]

The reason is that peak-load pricing schemes cause inefficiencies of their own. These problems are illustrated in Figure 2, in which the time of day is represented on the horizontal access and the total congestion (measured in congestion cost) is represented on the vertical access. Assume that the goal is to set a peak-load price during the busiest time of the day, represented in Figure 2 as the interval between t_1 and t_2. Some degree of inefficiency will result regardless of whether the network sets price at the lowest congestion cost during this period (represented by p_1), the highest congestion cost during this period (represented by p_2), or a price set somewhere in between (represented by p_3).

Consider first price p_1. Because p_1 falls below the congestion cost created by incremental usage at every point during the peak-load period, setting price at p_1 would encourage end users to increase their consumption of network resources even when the congestion cost of doing so would exceed the benefits. On the other hand, because p_2 exceeds the congestion cost created by incremental usage at every point during the peak load period, pricing at p_2 would deter usage even though increasing usage would cause consumer welfare. Setting price in

---

[51] *See* Rolla Edward Park &Bridger M. Mitchell, Optimal Peak-Load Pricing for Local Telephone Call 6, 32 (Rand Paper No. R-3404-1-RC March 1987); Lewis Perl, Impacts of Local Measured Service in South Central Bell's Service Area in Kentucky (May 21, 1985); Bridger Mitchell, *Optimal Pricing of Local Telephone Service*, 68 AM. ECON. REV. 517 (1978). For an overview, see Alfred E. Kahn & William B. Shew, *Current Issues in Telecommunications Regulation: Pricing*, 4 YALE J. ON REG. 191, 237-38 (1987).

Figure 2
Inefficiencies of Peak Load Pricing

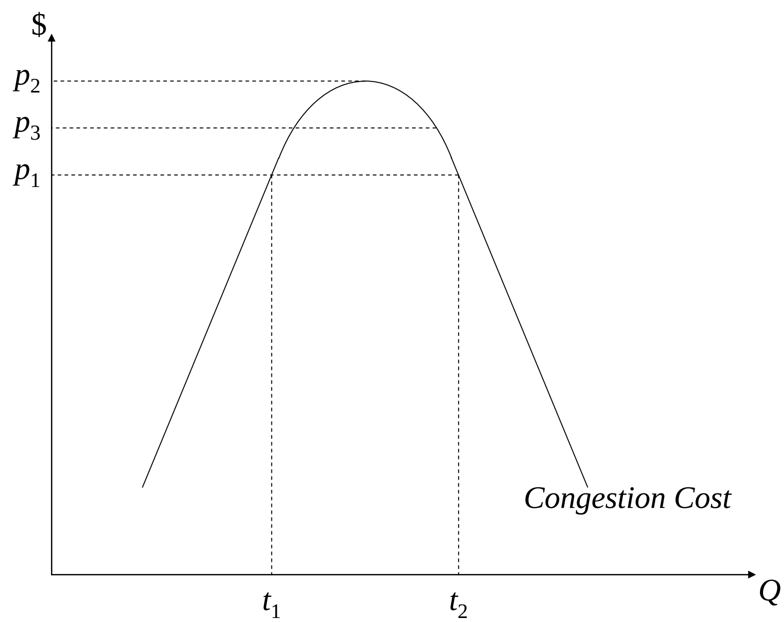

between at p_3 gives rise to both of these problems during different portions of the peak load period. During the middle of the peak-load period, p_3 would fall below the congestion costs associated with incremental usage and thus would provide end users with the incentive to increase their consumption even when the congestion costs imposed on others would exceed the benefits that that end user would derive from doing so. At the beginning and ending portions of the peak-load period, p_3 would exceed the congestion cost, in which case pricing at p_3 would deter additional usage even when increasing consumption would cause consumer welfare to increase.

An additional problem is that end users inevitably respond to the imposition of peak-load pricing by shifting some of their usage to the periods immediately preceding and following the peak-load period. The result is to create "shoulders" in the distribution of traffic on either side of the peak-load period. If this reallocation is sufficiently large, it can cause congestion costs outside the peak-load period to rise to welfare-reducing levels. As a result, networks that utilize peak-load pricing typically find it necessary also to impose near-peak rates (sometimes also called “shoulder rates”) during the period immediately preceding and following the peak-load period. Near-peak rates suffer from the same consumer welfare problems discussed above associated with peak-load rates, albeit to a smaller degree. The resulting pricing scheme also increases the complexity of the pricing scheme confronting consumers, by requiring them to incur the costs of keeping track of the price at any particular time of day and adjusting their behavior accordingly day. In addition, if end users are allowed to choose between metered pricing plan and an all-you-can-eat pricing plan, high-volume users have the strategic incentive to opt for the latter.

In the case of local measured service, these problems combined to dissipate the consumer benefits associated with peak-load pricing. There are aspects to Internet traffic likely to make peak-load pricing of broadband service even less likely to benefit consumers. As an initial matter, Internet traffic is much more variable than telephone traffic. For example, web browsing tends to generate sharp peaks of bandwidth usage followed by long periods of inactivity while the end user reads the webpage that has just been loaded. The result is that congestion on the Internet is likely to arise much more abruptly and be much more transient than on telephone networks, which makes it much more difficult whether and to what degree additional usage by one consumer will adversely affect other consumers.[52]

Congestion on the Internet may also often be quite localized. For example, it is quite possible that congestion will be quite high in one neighborhood while simultaneously being quite low in another portion of the same metropolitan area, depending on the size and bandwidth

---

[52] *See* Spulber & Yoo, *supra* note 13, at 1700.

intensiveness of the traffic being generated by end users in either part of the city at any given time. This means that a properly functioning congestion-pricing scheme would have to do more than impose different prices during different times of the day. At any particular time, it would also have to charge different prices in different portions of the network, depending on the local congestion conditions in any particular geographic area.[53]

Lastly, any congestion-based pricing system would have to take into account packet switched networks' ability to compensate for surges in demand by routing around areas of congestion in ways that circuit switched traffic associated with conventional telephone service cannot. While the ability to reroute traffic may mean that increases in congestion need not necessarily degrade network performance, the ability to route around trouble spots can also have the effect of transferring congestion to areas of the network that are located quite distant from the location where network flows are increasing. This can make congestion quite unpredictable.[54] Any congestion-based pricing would thus have to be able to incorporate information about the precise level of network flows and capacity in all portions of the network at any particular time in order to determine the magnitude of the congestion cost caused by any particular increase in network traffic. Such information was relatively easy to collect in local telephone systems, which have historically been dominated by incumbent local exchange carriers well positioned to collect such information. The Internet, however, operates on very different principles. Indeed, the decentralized nature of the Internet necessarily implies that no player has access to all of the information needed to assess all of the systemic effects.[55] Thus, although imposing bandwidth

---

[53] *See id.* at 1709-11.

[54] *See id.* at 1703-07, 1711.

[55] *See* Daniel F. Spulber & Christopher S. Yoo, *Rethinking Access to Last-Mile Broadband Networks* (forthcoming 2008).

tiering or peak-load pricing would capture some of the aspects of congestion pricing, institutional considerations force the outcomes under both regimes to fall short of the ideal.

D. Other Institutional Solutions

There thus may be good reason that network owners have not simply deployed metered pricing or peak-load pricing schemes and have instead begun to experiment with different institutional solutions. One particularly interesting solution to the problems of congestion is content delivery networks like Akamai, which reportedly handles more than fifteen percent of the world's web traffic.[56] 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 the cache maintained by Akamai.

This process often allows the resulting traffic to bypass the public backbone altogether, which in turn protects the query from any backbone congestion that may exist. 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.

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. Akamai thus represents a creative way to use server hardware as a substitute for network hardware.

---

56 *See* Yoo, *supra* note 13, at 1882-83.

The problem from the standpoint of network neutrality is that it is commercial service that is only available to content and applications providers willing to pay a premium above and beyond the basic Internet access fees that all content and applications providers pay. It thus violates the basic network neutrality principles that all like traffic travel at the same speed and that network owners be prohibited from charging content and applications providers more for higher-speed service.

On some occasions, network owners have taken to blocking access to websites when proven to be harmful. The best known of these examples is the practice of denying computers sending suspiciously large volumes access to port 25, which is the port that plays a key role in e-mail. Some networks estimate that this practice reduces the total amount of spam by as much as twenty percent.[57] Again, blocking port 25 violates the principle of treating all like content alike and may well have the effect of blocking legitimate e-mails. Yet, the practice of blocking port 25 is relatively uncontroversial.

In addition, ISPs that detect end users using applications that consume large amounts of bandwidth (such as leaving their browser open to a website like ESPN.com that streams video continuously or engaging in large amounts of peer-to-peer file sharing), will suggest to the end users that they change their practices or purchase a higher-bandwidth service that more accurately reflects the amount of congestion they are imposing on other end users. If the end user is unwilling to change, the ISP may choose to cease doing business with the customer.

I recount other examples of alternative institutional solutions short of imposing full-fledged congestion-based pricing in my other work.[58] All of these practices are to some degree

---

57 *See* Jim Hu, *Comcast Takes Hard Line Against Spam*, CNET News.com, June 10, 2004, *at* http://news.com.com/2100-1038_3-5230615.html.

58 *See* Yoo, *supra* note 13, at 1874-85.

inconsistent with the principles advocated by network neutrality proponents. In pointing out these practices, I make not attempt to show that any particular practice is always beneficial or always harmful or to make any assessment of which is likely to prove best. Indeed, rapid pace of change in terms of cost and functionality would make any such assessment too ephemeral a basis for policymaking.

My point is that policymakers will find it difficult, if not impossible, to determine the relative merits of any of these alternative institutional solutions at any particular time, let alone keep up with the rapid pace of technological change. So long as some plausible argument exists that a practice might be socially beneficial, the better course is to establish rules that give network operators the flexibility to experiment with that practice until its precise impact on consumers can be determined.

E. Prioritization

An alternative way to manage congestion is to prioritize time-sensitive and higher value traffic over other traffic.[59] Contrary to the suggestions of some network neutrality advocates, and as the following examples will show, prioritization is an often-used solution to the problems of congestion that enhances both network performance and consumer welfare.

1. NSFNET

An example from the days when the National Science Foundation ran the backbone illustrates the point nicely.[60] Before 1987, end users tended to connect to the backbone through

---

[59] Participants in the debate tend to characterize these practices either as "prioritization," "traffic shaping," or "network management" on the one hand or as "degradation" or "throttling" on the other hand. The difference is largely a matter of semantics. Both aspects are opposite sides of the same coin.

[60] *See* Yoo, *supra* note 4, at 22-23

dumb terminals. The total traffic they were able to introduce into the system was generally limited by the speed with which they were able to type keystrokes. All of this was changed by the emergence of the personal computer and modem technologies that made it possible to attach those PCs directly to the network. The increased functionality provided by the shift to personal computers increased the intensity of the demands that end users were placing on the network. The resulting congestion caused terminal sessions to run unacceptably slowly, and the fact that fixed cost investments cannot be made instantaneously created an inevitable delay in adding network capacity. This is precisely the type of technology- and demand-driven exogenous shock that makes network management so difficult.

NSFNET's solution was to reprogram its routers to give terminal sessions higher priority than file transfer sessions until additional bandwidth could be added. Indeed, the NSF did so in a way that is quite logical. It distinguished between sessions in which a human was sitting on the other end of the line (when delays of a few seconds can become intolerable) and gave them priority over machine-to-machine transfers that were already likely to take several minutes. The National Science Foundation's actions were not designed to harm consumers or innovation. The NSF prioritized traffic simply to minimize the practical impact of any disruption.

Indeed, such solutions need not be temporary: in a technologically dynamic world, one would expect that the relative cost 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 shortsighted to tie network managers' hands by limiting their flexibility in their choice of network management solutions.

2. Wireless

In addition, wireless broadband providers have begun to use prioritization techniques to compensate for the inherent bandwidth limits of spectrum-based communications. For example, some networks are attempting to leverage the fact that the physics of wave propagation dictates that the available bandwidth can vary as a person walks across a room. Some networks have begun to experiment with protocols that give priority to latency-sensitive applications, such as voice, and hold delay-tolerant applications until the end user reaches a location where the available bandwidth is relatively large, at which point the network will download all of the data associated with these delay-tolerant applications.[61]

This is an innovative solution to a real problem. Because those protocols discriminate based on application, it is the precisely the type of solution that network neutrality would prohibit.

3. Public Safety Spectrum in the 700 MHz Auction

Another prominent, recent example is the public/private partnership created by the FCC to govern the public safety spectrum allocated through the 700 MHz auction. The FCC's solution is to give allow public safety and commercial traffic to share bandwidth, but to allow the former to preempt 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."[62] This decision effectively concedes that prioritization of higher value traffic

---

[61] *See* Eun Ho Choi et al., *Throughput of the 1x EV-DO System with Various Scheduling Algorithms*, *in* IEEE EIGHTH INTERNATIONAL SYMPOSIUM ON SPREAD SPECTRUM TECHNIQUES AND APPLICATIONS 359 (Sept. 4, 2004); Charles L. Jackson, Wireless Handsets Are Part of the Network (Apr. 27, 2007), *available at* http://files.ctia.org/pdf/Comments_CTIA_SkypeOpposition_AppendixC_43007.pdf.

[62] 700 MHz Second Report, *supra* note 5, at 15431 ¶ 396.

represents an effective way to lower the cost of providing service, while at the same time representing a creative solution to extant bandwidth limitations.

4. Comcast's Network Management of BitTorrent

These previous examples shed new light Comcast's recent efforts to manage its network. As noted earlier, the fact that cable modem subscribers share bandwidth from the moment their traffic leaves their homes renders cable modem systems more vulnerable to local congestion than wireline technologies, such as DSL and FiOS. As the above-quoted studies show, cable modem systems are vulnerable to substantial degradation of service if as few as 15 of the 500 end users sharing a fiber node engage in file sharing.[63]

Comcast is not alone in singling out peer-to-peer filesharing as a primary source of congestion. As the record in this proceeding has shown, colleges, universities, and public school systems across the country have chosen to block student use of peer-to-peer filesharing systems. Although these schools' actions are driven in part by concern over potential copyright liability, they are also motivated in no small part by the need to manage congestion. Indeed, many report substantial reductions in congestion and network cost.[64] These schools clearly did not do so over any desire to harm competition or end users. The similarity between Comcast's targeting of peer-to-peer file sharing and the decisions by colleges and universities to target peer-to-peer file sharing both reflect the reality that an empirical matter such programs represent a primary source of congestion. The fact that these schools are in effect end users rather than network providers

[63] *See supra* note 25 and accompanying text.

[64] *See* National Cable & Telecommunications Association Comments at 6-7; Information Technology and Innovation Foundation Comments at 6; Comcast Comments at 20; Time Warner Cable Comments at 17.

does not alter the fact that blocking particular application may constitute a quick-and-dirty solution that can be implemented cost-effectively.[65]

Most importantly, the NSFNET experience shows the extent to which the way Comcast is managing its network in a way that minimizes the disruption to end users. Limiting its policy to one-way, upload-only sessions all but ensures that, as was the case with the NSFNET's policies, network management is invoked only with respect to computer-to-computer communication and not when a live person is waiting on the other end of the line. In addition, these policies are limited in both time and space, arising only in limited geographic areas and only in those areas when congestion rises to a level that it will degrade the service provided to others. As a result, the vast majority of BitTorrent sessions proceed without any interference. In addition, the resulting delays are typically limited to no more than a few seconds.

F. The Dangers of in Locking into or Prohibiting Any Particular Approach Ex Ante

Congestion represents a dynamic and constantly changing problem for networks. Changes in traffic size, traffic patterns, network configuration, and the relative cost of each solution causes the optimal solution to shift constantly.

The dynamic nature of the problem cautions against categorical solutions that either prohibit certain network management practices or lock the network in to any particular approach to dealing with the problem of congestion. As I have advocated at length elsewhere,[66] the insights of modern competition policy, illustrated by the Supreme Court's antitrust jurisprudence governing when to use case-by-case analysis and when to use categorical per se rules, militates

---

65 *See* Yoo, *supra* note 13, at 1879-80.

66 *See* Yoo, *supra* note 4, at 6-7; Christopher S. Yoo, *What Can Antitrust Contribute to the Network Neutrality Debate?*, 1 INT'L J. COMM. 493, 508-10 (2007).

strongly in favor of adhering to the case-by-case approach embodied in the FCC's *Madison River* decision.[67]

III. THE INCREASING NUMBER AND DIVERSITY OF INTERNET PLAYERS INCREASES THE DANGER OF OPPORTUNISTIC BEHAVIOR

Pressure to manage networks is not only the result of the increase in network traffic. It is also coming from changes in the nature of the people connected to the Internet.[68] It is often noted that the Internet is a network of networks. In the absence of some brooding omnipresence overseeing the entire network, the Internet depends on the cooperation of everyone connected to it. In particular, under the current configuration of the Internet, the management of congestion depends as much on the devices attached to the network owned by end users as it does on the network itself.

For example, when packets arrive at a router in the network faster than the router can process them, the router will queue those packets in a buffer. If the buffer becomes full and the arrival speed exceeds the router's ability to clear the queue, rather than allow the queue to become infinite, the router begins to drop packets. As the sending computer receives notice that some of the packets it is sending are being lost, TCP instructs the computer to slow down the pace with which it is transmitting packets.[69] To use Yochai Benkler's phrase, the Internet thus depends on the willingness of every network participant to "shar[e] nicely."[70]

---

[67] *See* Madison River Commc'ns, LLC, Order, 20 FCC Rcd 4295 (2005).

[68] *See* Christopher S. Yoo, *Would Mandating Network Neutrality Help or Hurt Competition? A Comment on the End-to-End Debate*, 3 J. ON TELECOMM. & HIGH TECH. L. 23, 34-37 (2004).

[69] *See* Edward W. Felten, Nuts and Bolts of Network Neutrality 2-3 (July 6, 2006), *available at* http:// itpolicy.princeton.edu/pub/neutrality.pdf.

[70] *See* Yochai Benkler, *Sharing Nicely: On Shareable Goods and the Emergence of Sharing as a Modality of Economic Production*, 114 YALE L.J. 273 (2004).

Such cooperation was quite likely when the Internet used primarily by academics, which would represent precisely the type of close-knit community in which such cooperative solutions work well.[71] The commercialization of the Internet beginning in the mid-1990s has led to a dramatic change. Internet users have become much more heterogeneous. The result is that we cannot rely on common values to hold things together.[72] The increasing risk that end users will act in ways that are not consistent with good behavior places ever-growing emphasis on trust.[73]

Computer scientists have long realized that it may be in the selfish best interest of a device attached to the network if all other devices adhered to these cooperative norms, while it continued to send packets into the network indiscriminately.[74] In essence, a device can opportunistically free ride on the fact that other devices play nicely, much like a sole driver illegally speeding through the carpool lanes. It may well make sense for a network provider to deny access to a device or an application that refused to cooperate in this cooperative approach to managing congestion

A weak version of this problem may be at work with respect to BitTorrent. To be clear, there is no evidence that BitTorrent is not honoring the traffic management protocols of TCP. The problem is not that BitTorrent is breaking the rules, but rather that the current rules do not require that it take the impact it is having on other end users into account. This sheds new light

---

71 *See* ROBERT C. ELLICKSON, ORDER WITHOUT LAW: HOW NEIGHBORS SETTLE DISPUTES (1991).

72 *See* Steven A. Hetcher, *The Emergence of Website Privacy Norms*, 7 MICH. TELECOMM. & TECH. L. REV. 97, 101 (2000-2001) (noting that "the Internet may present an especially difficult context for the emergence of efficient norms, however, as online participants would appear to be anything but close-knit").

73 *See* David Clark & Marjory Blumenthal, The End-to-End Argument and Application Design: The Role of Trust (Aug. 2007) (paper presented at the 34th annual Telecommunications Policy Research Conference), *available at* http://web.si.umich.edu/tprc/papers/2007/748/End%202%20end%20and%20trust%2010%20final%20TPRC.pdf.

74 *See* Stefan Savage et al., *TCP Congestion Control with a Misbehaving Receiver*, 29 ACM COMPUTER COMM. REV. 71 (1999), *available at* http://www.cs.ucsd.edu/~savage/papers/CCR99.pdf; Rob Sherwood et al., *Misbehaving TCP Receivers Can Cause Internet-Wide Congestion Collapse, in* PROCEEDINGS OF THE 12TH ANNUAL ACM CONFERENCE ON COMPUTER AND COMMUNICATIONS SECURITY 383 (Vijay Alturi et al. eds., 2005), *available at* http://www.cs.umd.edu/~capveg/optack/optack-ccs05.pdf; Jon M. Peha, *The Benefits and Risks of Mandating Network Neutrality, and the Quest for a Balanced Policy*, 1 INT'L J. COMM. 644, 651 (2007), *at* http://ijoc.org/ojs/index.php/ijoc/article/view/154/90.

on the significance of BitTorrent founder Bram Cohen's statement that, "My whole idea was, 'Let's use up a lot of bandwidth.'" When asked about the problems this would cause for the network, he replied, "Why should I care?"[75] In this case, the problem is not that BitTorrent is deliberately attempting to harm the service provided to other end users, but rather that the current pricing regime does not force BitTorrent users to take the impact that they are having on other end users into account.

These concerns may also provide some insight into the reluctance of some wireless providers to provide short codes. (In the case of Verizon's refusal to provide a short code to NARAL, Verizon has already acknowledged that this was a mistake, and I have suggested earlier that given the beneficial role that network management can play, an occasional obsolete policy does not justify wide scale regulation.) Network providers have some reason to be cautious in issuing short codes to potentially noncooperative parties. Verizon has already acknowledged that it does not issue short codes to those who would use them simply to send advertisements, distribute pornographic wallpaper or ringtones containing profanity or racial slurs, or charge excessive amounts. In addition, once short codes are established, it is possible that they might be resold, much as phone lists, address lists, and e-mail lists have been resold to telemarketers, junk mailers, and spammers. If so, those who participate in short codes may find themselves the recipients of unwanted communications.

Indeed, end users are expecting network owners to play larger roles in screening out unwanted communications. Network owners are expected to screen out spam and protect end users from viruses and other forms of malware. In other words, end users already expect network owners to exercise some degree of editorial discretion on their behalf. In addition,

---

[75] *See* David Downs, *BitTorrent, Comcast, EFF Antipathetic to FCC Regulation of P2P Traffic*, S.F. WEEKLY, Jan. 23, 2008,

every spam filter imposes some limits on end user control and occasionally screens out some communications that end users would like to receive.

Network owners trying to satisfy their customers must thus strike a careful balance. The question is whether FCC policy should prevent network owners from exercising some degree of discretion on behalf of their subscribers and, if so, how such a policy could be reconciled with allowing network owners to continue to screen out spam.

IV. POLICYMAKERS SHOULD BE CAREFUL NOT TO CONFUSE TECHNOLOGICAL CHANGE OR NORMAL ECONOMIC BARGAINING WITH MARKET FAILURE

The normal functioning of a network necessarily involves a great deal of change.[76] The distributed nature of the Internet and the rapid pace of technological change necessarily means that the Internet may not always run smoothly. Network providers will struggle to decide how to accommodate new technological developments. Market participants who bargain with each other in the normal course of doing business may occasionally deadlock and be unable to reach agreement. Such problems should not necessarily be regarded as problems in need of regulatory redress. On the contrary, a certain degree of toleration is required if the market for network services is to go through the normal process of reequilibration.

A. New Technological Developments

At times, some new development may arise to which the market may need some time to adjust. For example, when they first arose, network owners prohibited the use of VPNs and home networking devices.[77] This would ultimately prove short lived. Consumer pressure soon induced the network owners to change course. Although some have pointed to this development

[76] The analysis in this Part is adapted from Yoo, *supra* note 38.
[77] For an analysis, Yoo, *supra* note 13, at 1877, 1879-80.

as demonstrating the need to impose network neutrality regulation,[78] I think it demonstrates the opposite. It shows how consumers preferences exercised through the competitive process can force openness in ways that render regulation unnecessary.

B. Bargaining

At other times, the network owners and content and applications providers may disagree regarding the value of the services they provide. Rather than jump in, regulators should allow the give and take of the bargaining process to work its way through. The point is illustrated by two recent examples involving negotiations between the cable and satellite television industry on the one hand and network-owned broadcast stations on the other. In both the case of cable and satellite television, the broadcast stations have two options. They can either invoke their right free carriage on a cable or satellite system, or, if they think they can obtain some form of economic compensation from the cable or satellite system, they can forego those rights and instead negotiate their own carriage arrangements through arms-length negotiations. Most such negotiations are resolved amicably enough. But sometimes differences of opinion about the relative value the other party is providing and the strength of one's bargaining position can lead to temporary deadlock.

Indeed, this appears to be exactly what happened between Disney/ABC and Time Warner Cable in late 1999 and early 2000. The television stations owned by Disney/ABC decided not to exercise their rights to free carriage on local cable systems and instead opted to negotiate their own retransmission consent agreements. The previous agreement between Disney and Time Warner Cable expired in December 1999, and the parties negotiated a $1 billion agreement that

---

[78] *See* Wu, *supra* note 46, at 143, 152-53, 157-58, 160, 165.

would have lasted ten years. Just as that deal was about to be consummated, America Online announced its agreement to acquire Time Warner. ABC immediately asked for an additional fee increase of $300 million, and Time Warner refused. After five months of short-term extensions of the previous agreement and additional negotiations failed to yield an agreement, Time Warner dropped all of the ABC-owned stations on May 1, 2000.[79] ABC filed a complaint with FCC the same day, with FCC Chairman William Kennard warning that "[t]he television sets of average consumers should never be held hostage in these disputes" and criticizing "[t]he game of brink[s]manship" being played by the parties.[80] In the shadow of impending FCC action, Time Warner capitulated and put the stations back on after only one day. The FCC ruled Time Warner's actions illegal,[81] with Chairman Kennard again warning, "No company should use consumers as pawns in a private contract dispute."[82] Time Warner reached an agreement with ABC later that month, with ABC receiving the $300 million increase it sought.[83]

A similar dispute arose in March 2004, when DBS provider EchoStar was unable to reach a carriage agreement with Viacom/CBS. EchoStar cut Viacom programming on March 9, which left 9 million subscribers without MTV programming and 2 million subscribers without CBS. This time the FCC followed very different course of action. Rather than criticizing the parties, FCC Chairman Michael Powell simply acknowledged, "That's what sometimes happens in the market. Consumers usually lose and so do both parties. It usually doesn't happen very long."[84]

---

79 *See* Marc Gunther, *Dumb and Dumber*, FORTUNE, May 29, 2000, at 140.

80 Statement of FCC Chairman William E. Kennard Regarding Disney/ABC and Time Warner Dispute (May 2, 2000), *available at* http://www.fcc.gov/Speeches/Kennard/Statements/2000/stwek035.html.

81 Time Warner Cable, Emergency Petition of ABC, Inc. for Declaratory Ruling and Enforcement Order, 15 FCC Rcd 7882 (2000).

82 Statement of FCC Chairman William E. Kennard on Ruling in Time Warner-Disney Dispute (May 3, 2000), *available at* http://www.fcc.gov/Speeches/Kennard/Statements/2000/stwek036.html.

83 *Time Warner, Disney Sign Long-Term Deal*, TELEVISION DIGEST, May 29, 2000 (available at 2000 WLNR 48749-8).

84 Jonathan D. Salant, *EchoStar, Viacom Expect Quick Resolution*, KAN. CITY STAR, Mar. 11, 2004, at C3.

The parties settled the dispute two days later. Although both sides claimed victory,[85] the ultimate terms were closer to what Viacom initially sought.[86]

Taken together, these episodes demonstrate the extent to differences of opinion about the value of the services being provided are a natural (indeed, an essential) part of a market-based economy. People of disagree, and the bargaining process sorts out which side is mistaken fairly quickly. In order for this mechanism to work properly, the parties must enough breathing room in which to engage in arms-length negotiating. Interfering with these negotiations threatens to eliminate the give-and-take upon which normal economic processes depend. Policymakers must thus be careful not to regard the inability to reach agreement as a definitive sign of market failure or the necessity of government intervention. On the contrary, a certain amount of deadlock is the sign of a properly functioning economic market.

CONCLUSION

The Internet is constantly changing, and there are almost as many predictions of its ultimate shape as there are network analysts. An ever-growing number of users is using the network in increasingly diverse ways to accomplish increasingly diverse goals. New technological developments place new pressures on the network and change the relative cost of the various ways to address the growing problem of congestion. In addition, broadband service is being provided by an increasingly diverse array of transmission technologies that differ widely in their susceptibility to congestion as well as with respect to the available solutions. In addition,

---

[85] *See* Kris Hudson, *EchoStar, Viacom Claim Win: Angry Viewers Cited as Motivating Factor*, DENVER POST, Mar. 12, 2004, at C1.

[86] *See* Phyllis Furman, *Viacom Seen Getting Best of EchoStar*, N.Y. DAILY NEWS, Mar. 12, 2004, at 76.

ISPs are entering increasingly diverse relationships that are altering the basic economics of the Internet.

Every tool available for network management has its strengths and its weaknesses. It is thus unlikely that any one-size-fits-all solution will be the right choice in all circumstances and at all times. In such a dynamic environment, the Commission should be careful to preserve the flexibility of the network's ability to respond to congestion. Taking tools out of the network management toolkit may well have the unfortunate effect of limiting our ability to harness the benefits of the growing phenomenon that is the Internet for all consumers.