The Information Interchange: Interconnection on the Internet

Michael Kende, *Senior Economist, Common Carrier Bureau*

Jason Oxman, *Counsel for Advanced Communications, Office of Plans and Policy*[*]

I. Introduction

The Information Highway, as it exists today, is not a monolithic, uniform network. Like the highway system itself, the Information Highway has smaller roads and on-ramps to allow users to access the system, but unlike the highway system, the Information Highway is in fact a network of networks, owned and operated by different companies, known as Internet backbones. This paper examines the Interchanges of the Information Highway - the policies and means by which backbones exchange traffic destined for each other's customers. Currently, such interconnection, often referred to as peering, is unregulated. This paper examines the history and future implications of unregulated Internet interconnection, and compares that to a scenario with regulation. The paper concludes that not only would interconnection regulation impose onerous and unwarranted obligations on backbone providers, but that the competitive environment that characterizes the Internet will continue to provide innovative solutions to any and all arguments that some would use to justify regulation.

In section two of this paper, we examine the history of Internet interconnection and describe current interconnection policies between Internet backbones. We then examine several pressures on the current system of interconnection in section three. These pressures may lead to calls for regulation of Internet backbones; in section four we describe the form that such regulation could take, and argue that such regulation would be harmful to the industry. In section five we conclude that, given the competitive marketplace in which Internet backbones operate, any regulation would be unnecessary in any event.

II. Background

A. Introduction

The Internet can be thought of as a hierarchy consisting of three levels of participants: end users, Internet service providers (ISPs), and Internet backbone providers (IBPs). At the

---

[*] The views expressed in this paper are those of the authors, and do not necessarily represent the views of the Federal Communications Commission, the Chairman, any Commissioners, or other staff.

bottom of the hierarchy are the end users who use the Internet to send and receive information. For simplicity, we include as end users content providers who create much of the information that is available on the Internet. End users access the Internet via Internet service providers such as AOL or Erols.[1] Small business and residential end users generally use modems to connect to their ISP over standard telephone lines, while larger businesses and content providers generally have dedicated access to their ISP over leased lines. At the top of the hierarchy are Internet backbone providers such as UUNET and PSINet. IBPs own or lease high speed fiber optic networks connected together by routers, which they use to deliver traffic to and from their customers. IBPs primarily sell wholesale Internet connectivity to ISPs that essentially resell this connectivity to their customers, but may also function as ISPs and have as customers end users who purchase retail Internet connectivity.[2]

Each IBP essentially forms its own network that enables all connected end users and content providers to communicate with one another. End users, however, are generally not interested in communicating just with end users connected to the same backbone. In order to provide end users universal connectivity, backbones must interconnect with one another and exchange traffic destined for each other's end users. It is this interconnection that makes the Internet the "network of networks" that it is today. As a result of widespread interconnection, end users have an implicit expectation of universal connectivity whenever they log on to the Internet, regardless of which ISP they choose. ISPs are therefore in the business of selling access to the entire Internet to their end-user customers. IBPs, selling wholesale services to ISPs, must in turn be able to provide access to the entire Internet to their customers. The driving force behind the need for these firms to deliver access to the whole Internet to customers, be they wholesale or retail, is what is known in the economics literature as network externalities.

B. Network Externalities

Network externalities arise when the value, or utility, that a consumer derives from a product or service increases as a function of the number of other consumers of the same or

[1] Of course, many ISPs such as AOL are also significant content providers. Here we focus on ISPs as providers of Internet access to end users.

[2] *See* Kevin Werbach, "Digital Tornado: the Internet and Telecommunications Policy" (OPP Working Paper Series No. 29, 1997)(Digital Tornado) at 10-12.

compatible products or services.[3] They are called network externalities because they generally arise for networks whose purpose it is to enable each user to communicate with every other user; as a result, by definition the more users there are, the more valuable the network. These benefits are externalities because a user, when deciding whether to join a network (or which network to join) only accounts for the private benefits that the network will bring her, and will not account for the fact that, by joining this network, the benefit of the network for other users increases. This latter effect is an externality.

Network externalities can be direct or indirect. Network externalities are direct for networks that consumers use to communicate with one another; the more consumers there are on the network, the more valuable the network is for each consumer.[4] The phone system is a classic example of a system providing direct network externalities. The only benefit of such a system comes from access to the network of users. Network externalities are indirect for systems that require both hardware and software in order to provide benefits.[5] As more consumers buy hardware, more software compatible with this hardware will be produced, making the hardware more valuable to users. A classic example of this is the compact disk system; as more consumers purchased compact disk players, companies increased the variety of compact disks available, making the players more valuable to their owners.[6] These network externalities are indirect because consumers do not purchase the systems to communicate with others, yet they indirectly benefit from the adoption decision of these other consumers.

One of the aspects of the Internet that makes it unique is that it offers a combination of both direct and indirect network externalities. Users of applications such as email and IP telephony derive direct network externalities from the system; the more Internet users there are, the more valuable the Internet is for such communications. Users of applications such as the World Wide Web derive indirect network externalities from the system; the more Internet users

---

3 *See* Michael L. Katz and Carl Shapiro, "Systems Competition and Network Effects," *Journal of Economic Perspectives*, Vol. 8, No. 2, Spring 1994, at 93-115; Nicholas Economides, "The Economics of Networks," *International Journal of Industrial Organization*, Vol. 14, No. 2, March 1996.

4 *See* Michael L. Katz and Carl Shapiro, "Network Externalities, Competition, and Compatibility," *American Economic Review*, Vol. 75, June 1985, at 424-440.

5 *See* Jeffrey Church and Neil Gandal, "Network Effects, Software Provision, and Standardization," *Journal of Industrial Economics*, Vol. 40, March 1992, at 85-104.

6 For an empirical description of the interplay between compact disc hardware sales and the availability of compact discs, *see* Neil Gandal, Michael Kende, and Rafael Rob, "The Dynamics of Technological Adoption in

there are, the more Web content will be developed, making the Internet more valuable for its users. The ability to provide direct and indirect network externalities to customers provides an almost overpowering incentive for Internet backbones to cooperate with one another by interconnecting their networks.

C. Peering and Transit

During the development of the Internet, there was only one backbone, and therefore interconnection between backbones was not an issue.[7] In 1986, the National Science Foundation (NSF), funded the NSFNET, a 56 Kbps network created to enable long-distance access to five supercomputer centers across the country. In 1987, a partnership of Merit Network, Inc., IBM, and MCI began to manage the NSFNET, which became a T-1 (1.544 Mbps) network connecting thirteen sites in 1988. The issue of interconnection later arose as a number of commercial backbones came into being, and eventually supplanted the NSFNET.

In 1991, a number of commercial backbone operators including PSINet, UUNET, and CerfNET established the Commercial Internet Exchange (CIX). The CIX consisted of a router, housed in Santa Clara, California, that was set up for the purpose of interconnecting these commercial backbones and enabling them to exchange their end users' traffic. Then, in 1993, the NSF decided to leave the management of the backbone to competing, commercial backbones. In order to facilitate the growth of overlapping competing backbones, the NSF designed a system of geographically dispersed Network Access Points (NAPs), each consisting of a shared switch or LAN used to exchange traffic. The four original NAPs were in San Francisco (operated by PacBell), Chicago (BellCore and Ameritech), New York (SprintLink) and Washington, DC (MFS). Backbones could choose to interconnect with one another at any or all of these NAPs. In 1995, this network of commercial backbones and NAPs became the Internet as we now know it.

The interconnection of commercial backbones is unregulated. In the beginning, it was governed mainly by the informal interactions that characterized the Internet at the time. The NSF did not establish any interconnection rules at the NAPs, and Internet backbones are not

---

Hardware/Software Systems: The Case of Compact Disc Players, Centre for Economic Policy Research, Discussion Paper No. 2078, February 1999.

[7] *See* Digital Tornado at 13-16 for a brief history of the Internet. *See also* Robert H'obbes' Zakon,"Hobbes' Internet Timeline v4.1," at http://www.isoc.org/guest/zakon/Internet/History/HIT.html.

regulated by the Federal Communications Commission or any other government agency.[8] The backbones developed a system of interconnection known as peering. As a system of interconnection, peering has a number of distinctive characteristics. First, peering partners only exchange traffic that originates with the customer of one backbone and terminates with the customer of another backbone. As part of a peering arrangement, a backbone would not act as an intermediary and accept the traffic of one peering partner and transit this traffic to another peering partner.[9] Second, peering partners exchange traffic on a settlements-free basis, also known as bill-and-keep or sender-keeps-all.[10] The only cost involved in peering is that each partner pays for the equipment and transmission capacity needed for the two peers to meet at each peering point. Third, peering partners have adopted what is known as "hot potato routing;" they each exchange their customers' traffic destined for the other backbone at the nearest possible exchange point. Finally, the recipients of traffic only promise to undertake "best efforts" to transmit data to its destination rather than guarantee any level of performance in delivering packets received from peering partners.

Because each bilateral peering arrangement only allows backbones to exchange traffic destined for each other's customers, backbones need a significant number of peering arrangements in order to have access to the full Internet. UUNET, for instance, claims to have over fifty private DS-3 peering connections with other ISPs internationally.[11] As a result of this and other considerations, there are few backbones that rely solely on peering to meet their interconnection needs.[12] The alternative to peering is a transit arrangement between backbones. There are two main differences between transit and peering. The first is that, in a transit arrangement, one backbone pays another backbone for interconnection, and therefore becomes a wholesale customer of the other backbone. In return, unlike peering, the backbone selling the

---

8 *See* Jason Oxman, "The FCC and the Unregulation of the Internet" (OPP Working Paper Series No. 31, 1999)(Unregulation of the Internet).

9 *See* Intermedia Communications, *Peering White Paper*, 1998, http://www.intermedia.com (Intermedia White Paper) at n.1 for a definition of peering.

10 In a bill-and-keep or sender-keep-all arrangement, each partner bills its own customers only for the origination of traffic and does not pay the other partner for terminating this traffic. In a settlement arrangement, on the other hand, the network on which the traffic originates pays the other network for terminating the traffic. If traffic flows between the two networks are balanced, the net settlement that each pays is zero, and therefore a bill and keep arrangement may be preferable because the networks do not have to spend money on billing systems.

11 UUNET, Network FAQ, http://www.uu.net/land.en/network/faq.html. Private peering is discussed below.

12 It is not known exactly how many backbones rely only on peering, because peering information is not publicly available; all peering agreements are subject to non-disclosure agreements.

transit services will route traffic from the transit customer to its peering partners as well as its other customers. Many backbones have adopted a hybrid approach to interconnection, peering with a number of backbones and paying for transit from one or more backbones in order to have access to backbones with whom they do not peer.

There is no accepted convention that governs when two backbones will or should decide to peer with one another, nor is it any easy matter to devise one. The term peer suggests equality, and one convention may be that backbones of equal size will peer. However, there are many measures of backbone size, *e.g.* geographic spread, capacity; traffic volume, or number of customers. It is unlikely that two backbones will be similar along many dimensions. One may have fewer, larger customers than the other, another may reach into Europe or Asia, and so forth. The question then becomes, how the backbones weigh one variable against another. Given the complexity of such judgements, it may be best to use a definition of equality proposed byone industry participant, that companies will peer when they perceive equal benefit from peering based on their own subjective terms, rather than any objective terms.[13]

The system of peering as it developed originally has evolved over time. Originally, most peering took place at the NAPs, as it was efficient for each backbone to interconnect with as many backbones as possible at the same location. Given the rapid growth in Internet traffic, the NAPs soon became congested, leading to delayed and dropped packets. Intermedia Business Solutions argues that at one point packet loss at the MAE-East NAP reached up to 20 percent.[14] As a result of this congestion, many backbones began to interconnect directly with one another. This system has come to be known as private peering, as opposed to the public peering that takes place at the NAPs.[15] Although this system developed partly in response to congestion at the NAPs, in some cases it may be more efficient for the backbones. For instance, traffic that may have been routed to a NAP, just to switch backbones and terminate in the same city can be exchanged within the same city if the two backbones privately peer in that city. This alleviates

---

13 Geoff Huston, “Interconnection, Peering and Settlements,” January 1999, <http://www.telstra.net/gih/peerdocs/peer.html>

14 Intermedia White Paper at 2.

15 Intermedia states that it adopted a “dual peering policy aimed at promotion open public peering with all ISPs, while selectively seeking out strategic, private peering partners … [in order to offer customers] the highest standard of Internet connectivity available. *Id*. at 3. Intermedia states that this “will create a win-win solution for everyone and a better management approach to the Internet.

the strain on the NAPs. According to one estimate, 80% of Internet traffic is now exchanged via private peering.[16]

There have been a number of allegations in the past, notably during the MCI WorldCom merger proceeding, that the whole system of interconnection between backbones is at risk. At least one industry observer argues that the emerging system of private peering will enable the top backbones to act in an anticompetitive manner by excluding smaller backbones from private peering and then raising prices.[17] However, there are fears that even private peering is at risk. Partly, this is due to variations in the strategies and growth of individual backbones that have led to some breakdowns in the system. Partly, this is also a result of the enormous pressure that the exponential growth of the Internet has put on the backbones and the interconnection points connecting the backbones. The final, and perhaps terminal, source of stress may arise from the increasing reliance the end users place on the Internet for critical functions whose service quality backbones may not be able to guarantee under the current system.

III. Interconnection Issues

A. End of Peering

Internet backbones face conflicting incentives. On one hand, they have an incentive to cooperate with one another in order to provide their customers with access to the full range of Internet users and content. On the other hand, these same backbones have an incentive to compete with one another for both retail and wholesale customers. The need for backbone *X* to interconnect with backbone *Y* in order to provide access to backbone *Y*'s customers creates what might be termed a *competitive network externality*; this interconnection also enables backbone *Y* to provide its customers access to backbone *X*'s customers. As long as *X* and *Y* are relatively equally sized, there is a strong incentive for them to continue to cooperate with one another in spite of competitive network externalities; if either unilaterally stops interconnecting it has no guarantees that it will benefit more than the other from such an action. This situation seems to characterize the early days of the commercial Internet, when a number of backbones were relatively similar in size. Recently, there have been allegations that as certain backbones grew

[16] This estimate was given by Michael Gaddis, chief technical officer of SAVVIS Communications. Randy Barrett, "ISP Survival Guide," inter@ctive week online, December 7, 1998.

[17] Jack Rickard, "Yet another unique moment in time. Peering redux – back to the future and the essentials of a competitive Internet," Editor's Notes, May 1998, *Boardwatch Magazine*.

they began to engage in uncooperative, if not anti-competitive, practices. We examine the validity of these claims here.

In early 1997, UUNET, followed by several other large backbones, attempted to end peering with a number of smaller backbones and instead charge them for transit.[18] A number of commenters argued that this move was anti-competitive, while the president and CEO of UUNET, John Sidgmore, argued that "a few years ago all ISPs were generally the same size and used each other's infrastructures to a more or less equal extent… that situation no longer exists and consequently there are many cases where peering is not appropriate."[19] A year later, there was much concern that when WorldCom, which had purchased UUNET and several other backbones, merged with MCI, the combined backbone would be dominant and could exercise market power against its smaller competitors in a variety of ways, including refusing to peer with them.[20] For instance, in their argument opposing the MCI WorldCom merger, Level 3 argued that MCI and WorldCom were refusing to peer with Level 3, and that the merger would increase their incentive to discriminate against rivals seeking to interconnect.[21]

We ask here whether it is anti-competitive for a backbone to deny peering to another backbone or whether there might be another justification for such an action. We define anti-competitive to mean the ability to raise prices while keeping out potential entrants that may otherwise constrain prices. We thus seek to identify harms that effect consumers, not competitors. Actions that effect competitors may not harm consumers – for instance, a merger may increase the efficiency of a firm and result in lower prices which may harm competitors, but is not anti-competitive as it benefits customers. First, we examine whether there may be

---

18 It is not clear whether this attempt was successful. *See Id.*

19 Steve Lohr, "Internet Growth Brings Up Questions of Governance," *New York Times*, May 12, 1997, http://www.nytimes.com/library/cyber/week/051297regulate.html, and Jon Swartz, "UUNET to Begin Charging Some ISPs for Net Traffic," *San Francisco Chronicle*, May 13, 1997, http://www.sfgate.com/cgi-bin/article/archive/1997/05/13/BU27478.DTL. UUNET press release, "UUNET Details Peering Strategy," May 12, 1997. http://www.us.uu.net/press/1997/peering.html.

20 *See Application of WorldCom, Inc. and MCI Communications Corporation for Transfer of Control of MCI Communications Corporation to WorldCom, Inc.*, CC Docket No. 97-211, Memorandum Opinion and Order, 13 FCC Rcd 18025, 18103-18115, paras. 142-156 (1998)(*MCI/WorldCom Order*). The issue was resolved when MCI sold its Internet assets to Cable and Wireless in order to satisfy regulators' concerns. *See* European Commission Press Release, "Commission Clears WorldCom and MCI Merger Subject to Conditions," July 8, 1998; DOJ Press Release, "Justice Department Clears WorldCom/MCI Merger After MCI Agrees to Sell its Internet Business," July 15, 1998; *MCI/WorldCom Order*, 13 FCC Rcd at 18109-18115, paras. 151-156.

21 *See* Letter from Terrence J. Ferguson, Senior Vice President and General Counsel, Level 3 Communications, to Magalie Roman Salas, May 29, 1998, Attach.

legitimate reasons for a backbone to discontinue peering with a smaller rival. Then we examine whether the decision to deny peering may be anticompetitive.[22]

There may be a number of legitimate reasons for one backbone not to agree to peer with another backbone. One reason may be that one of the backbones perceives that peering would enable the other backbone to free-ride. This may arise if one national backbone has a presence on both coasts, and another only has a regional presence on one coast. If the two backbones peered on the regional backbone's coast, when a customer of the regional backbone requests a web page from a customer of the national backbone whose server is on the other coast, then the national carrier would carry both the request from one coast to the other and the response back. The national backbone may thus refuse to peer on the grounds that it is bearing the expense for a national infrastructure that the regional carrier can then benefit from at no cost. As a result, a number of backbones include in their publicly stated peering policies that the companies be willing and able to peer at a number of geographically diverse locations.[23]

Another example of perceived or actual free-riding that may arise in a peering relationship comes from the "hot-potato routing" that characterizes peering arrangements. One backbone may choose for a variety of reasons to focus on providing service to web sites. Given hot-potato routing, this web-centric backbone will receive and carry any request for a web page from a customer of a peered backbone at the first available connection and it will send the requested page back over to requesting customer's backbone at the nearest connection point. Since, as a rule, web pages are much larger in volume than the corresponding request, the result is that backbones that peer with web-centric backbones will carry much more traffic volume than the web-centric backbones. They may then refuse to peer on the grounds that they are bearing

---

[22] Level 3's Chairman, James Crowe, claimed that MCI and WorldCom's refusal to peer with Level 3 constituted "monopolistic behavior." Joan Engebretson, "Level 3: whiner or visionary," *Telephony Magazine*, May 25, 1998 at 7. *See also* John J. Keller, "Level 3 Assails the WorldCom-MCI Deal," *The Wall Street Journal*, May 20, 1998, at B10.

[23] For instance, UUNET's North American Peering Policy stated, among other things, that "a peering candidate needs to meet UUNET at minimally four geographically diverse locations across the US," with a minimum requirement of an East Coast location and a West Coast location, with, "ideally" two Midwest locations. Second Joint Reply of WorldCom, Inc. and MCI Communications Corporation, in the WorldCom MCI Merger Proceeding, filed March 20, 1998, at App. D (UUNET Peering Policy). Before divesting its Internet backbone, MCI's Public Peering Policy contained a similar provision, that a peering partner have connections to "at least four geographically dispersed public interconnection points where MCI is also connected." *Id*. at App. E (MCI Peering Policy).

the expense for more capacity than the backbone that is actually hosting the content that utilizes this capacity.[24]

Finally, a backbone may refuse to peer with a smaller backbone if the smaller backbone will not be able to exchange a certain minimum level of traffic. In both MCI and UUNET's peering policies, on file with the Commission as a result of the MCI/WorldCom merger proceeding, there is a requirement that a peering candidate exchange a certain minimum of data at the beginning of the peering relationship.[25] An MCI spokesperson was quoted as saying that, for this reason, Level 3 was denied peering by MCI.[26] One justification given by the larger backbones is that it is difficult and costly to allocate necessary resources to potential peers with low current volumes that may or may not grow rapidly in the future. Nevertheless, this requirement may place backbones with low volumes in a Catch 22; without a large number of customers, it is not possible to get peering with the large backbones, but without peering, it may be difficult to gain a large number of customers. In order to determine whether the latter statement is valid, one must examine the implications of not being able to peer.

Larger backbones are not refusing to interconnect with smaller backbones; rather, instead of peering with these smaller backbones, they offer them a transit arrangement.[27] With transit, if backbone *A* is refused peering by backbone *B*, then backbone *A* must use a transit arrangement in order for its customers to have access to backbone *B*'s customers. Backbone *B* could take transit directly from backbone *A*, or it could become a transit customer of a third backbone *C* that is interconnected with backbone *A*. As a transit customer, it may be possible to grow and later qualify to peer with backbones that initially refused peering, including the transit supplier. Nevertheless, a backbone may prefer peering rather then being a paying transit customer, for either quality or cost reasons.

---

[24] UUNET argues that companies which provide "web server farm" services and request peering with UUNET are "seeking to use UUNET's network for free, after UUNET has spent hundreds of millions of dollars to create its infrastructure." UUNET May 12, 1997 Press Release.

[25] UUNET expected a peering candidate "to exchange at least 40 Megabits of traffic total average utilization." UUNET peering policy. MCI imposed a "minimum traffic requirement of 20 Mbps per pair of direct DS-3 peering connections," and that "each individual DS-3 connection must carry a minimum of 5 Mbps." MCI Peering Policy.

[26] Joan Engebretson, "Level 3: whiner or visionary?"

[27] *See*, *e.g.*, Rob Frieden, "Without Public Peer: The Potential Regulatory and Universal Service Consequences of Internet Balkanization," Virginia Journal of Law and Technology, Fall 1998, Vol. 3, No. 8, 1522-1687, at para. 16.

In a transit relationship, one backbone must pay another for access to the Internet. For instance, at the time that UUNET changed its peering policy in 1997, it announced that wholesale connectivity started at $2,000 per month for a T-1 connection and $6,000 for a fractional T-3 connection.[28] Transit does not involve the same service as does peering, and therefore, refusing peering in favor of transit is not a means of charging for a service that was otherwise provided free of charge. As noted above, unlike peering, transit gives a backbone access to the entire Internet, not just the customers of the peering partner. In order to provide transit customers with access to the entire Internet, the transit provider must either maintain peering arrangements with a number of other backbones or in turn must pay for transit from another backbone. In other words, a backbone providing transit services is providing access to a greater number of customers than a peering partner, and incurring a correspondingly higher cost.

In terms of the quality of transit, Level 3 has suggested that, as a transit customer of another backbone, Level 3 would depend on the supplying backbone for delivery of IP traffic, at the very least placing Level 3 at a marketing disadvantage. On the other hand, at least one backbone, SAVVIS, initially relied only on transit connections and not peering, and has been very competitive in terms of quality.[29] As a result, it is difficult to conclude that paying for transit puts a backbone at a disadvantage from a quality point of view.

Finally, there is an implication that, having denied peering to smaller backbones, the larger backbones providing transit may increase the cost of transit in order to squeeze out the smaller rivals. Given the current competitive situation among the large backbones, it seems unlikely that any of the large backbones will unilaterally raise transit prices. The first reason is traditional, that the large backbones will compete for the transit business of smaller backbones in order to increase their revenues. The second reason is unique to the Internet, however. In negotiating peering, one important bargaining chip is the number of customers that a backbone brings to the table. Therefore, the larger backbones will compete with each other to win transit

---

[28] UUNET May 12, 1997 Press Release. A T-1 connection is a digital transmission link with a capacity of 1.544 Mbps. A fraction T-3 connection is a portion of a T-3 (44.7364 Mbps) digital transmission link. We note that transit agreements are not disclosed, making it easier for suppliers to discriminate against customers without detection. Letter from Terrence J. Ferguson, Senior Vice President and General Counsel, Level 3 Communications, to Michelle Carey, Common Carrier Bureau, August 7, 1998 (Level 3 Aug. 7 Letter), Attach. at 4.

[29] Doug Mohney, "SAVVIS Shifts Gears and Ownership," Boardwatch, April 1999. <http://boardwatch.internet.com/mag/99/apr/bwm66.html>. Since 1997, SAVVIS Internet has been rated #1 by

customers in order to be able to negotiate peering relationships with other backbones. The larger backbones are also unlikely to be able to tacitly agree to coordinate a transit price increase for the simple reason that, as with most agreements to tacitly collude, they would each have an incentive to undercut the others to gain customers, and, given the widespread use of non-disclosure agreements, any deviation from collusion would be difficult to detect.

Therefore, we find that in the current status quo, there is little chance that any of the large backbones may have either the incentive or the ability to act in an anticompetitive manner. On the other hand, there is legitimate concern that, if a single backbone became dominant, it would be able to engage in a number of anticompetitive actions. First, by definition a dominant firm has the ability to unilaterally affect prices. In addition, a dominant backbone would have both the ability and the incentive to stop cooperating with smaller backbones. Failure to cooperate could take a number of forms, explored below, including refusing to interconnect at all, executing a price squeeze, or degrading the quality of interconnection by not upgrading the capacity of connections with smaller backbones. It is worth noting that this situation characterizes local telephony, where incumbent local exchange carriers (LECs) must be compelled by statute to interconnect with smaller competitive LECs.[30]

An established result in the network externalities literature is that a larger network is less likely to want to become compatible with a smaller network, as customers of the smaller network have more to gain from being able to communicate with those of the larger network than vice versa.[31] In the context of the Internet, if larger backbones refused to interconnect with smaller ones, the customers of the smaller backbone would be inclined to switch to the larger network in order to enjoy the network externalities associated with the larger backbone's customer base. Although customers of the larger network would be hurt by losing access to the smaller network's customer base, they are unlikely to respond by switching to the smaller network. As a result, the smaller backbone is unlikely to provide any significant competition, reinforcing the dominance of the largest backbone.

Keynote Systems. Jack Rickard, Editor's Notes, *Boardwatch Magazine*, May 1998. SAVVIS created private NAPs, bought transit from the largest backbones, and didn't peer at all. Rickard notes that this is expensive, however.

30 This is required in section 251 of the Telecommunications Act of 1996. 47 U.S.C. § 251.

31 Katz and Shapiro (1985); Jacques Cremer, Patrick Rey, and Jean Tirole, "Connectivity in the Commercial Internet," May 1999, mimeo.

A dominant backbone may choose to interconnect with smaller backbones, but is in a position to set up a price squeeze. In a price squeeze, a vertically integrated firm with market power over an essential upstream input raises the price of this input to rivals competing in downstream markets. The increased cost of this essential input forces downstream rivals to raise their retail prices. The vertically integrated firm, on the other hand, does not pay more for this input, and can then undercut the downstream rivals in retail markets and thereby increase market share and profits. In the backbone example, the essential input that smaller backbones must purchase from larger ones is transit services; all firms in turn compete to sell retail services to ISPs or directly to end-users. Having weakened smaller rivals, the dominant backbone can raise retail prices and prevent entry of new backbones that could constrain the dominant firm's market power.

Another action that a dominant backbone might take was raised by GTE during the MCI WorldCom merger proceeding, when GTE suggested that the merged firm may engage in non-price discrimination.[32] GTE presented a scenario in which a dominant backbone may degrade interconnections with other backbones in order to win customers from that backbone. This could most easily be done by "slow rolling" the increases in the capacity of interconnection with other backbones that are necessary to keep pace with the rapid growth in demand that the Internet has experienced. Under this scheme, a backbone, *A*, may degrade a connection with a smaller backbone, *B*. As *B*'s customers begin to feel the effects of this degradation when communicating with customers of backbone *A*, they may be likely to switch to backbone *A* in order to experience better connections with customers of backbone *A*. It should be noted that the customers of backbone *A* will also be affected by this degradation when they attempt to communicate with the customers of backbone *B*. Therefore, *A* must be significantly larger than *B* so that its customers feel less effects when communicating with backbone *B* and don't themselves switch. In order to further limit the effects on its own customers, GTE argued that backbone *A* would engage in what they called "serial degradation" and only target one smaller backbone at a time. This strategy again depends on one backbone already being dominant.[33]

---

[32] *See* Cremer, Rey, and Tirole (1999). These authors advised GTE on the non-price discrimination issue during the MCI WorldCom proceeding, and have formalized their analysis in this paper.

[33] GTE and its experts note that in a situation in which there is no dominant backbone, there would be no incentive for any backbone to attempt a serial degradation strategy in order to become dominant. *Id*., Section 6.

Thus, we find that a dominant backbone may be able to act in an anticompetitive manner by engaging in price and non-price discrimination. Our analysis suggests that it is unlikely in the current environment that a firm may become dominant by discriminating against competitors. The easiest and most likely means for a backbone to grow to a dominant size is a merger. There is precedent to suggest that antitrust laws may be sufficient to keep the backbone competitive, notably in the MCI WorldCom merger case, when the European Union and the U.S. Department of Justice made the divestiture of MCI's Internet business a condition for merger approval. As long as antitrust authorities adhere to the precedent set in the MCI WorldCom merger, it is therefore unlikely that any dominant backbone will emerge.

Although it may be unlikely that any backbone or group of backbones is able to act in an anti-competitive fashion, the current environment of peering among the top tier of backbones may nevertheless not be stable in the long run. Given the rapid growth in demand that the Internet is experiencing, capacity must constantly be increased at all levels: the capacity of the backbone itself, the peering capacity between backbones, and the access capacity between the backbone and customers. In terms of choosing between upgrading peering capacity or access capacity, backbones may concentrate their resources on the latter at the expense of the former, with both direct and indirect benefits. First, backbones are in the business of selling access; peering is simply a means of supplying the desired product. Therefore backbones have a direct incentive to increase access capacity to increase their revenues. Second, backbones may indirectly profit from not upgrading the capacity of peering interconnections. In a variant of the serial degradation argument discussed above, by slow rolling such upgrades, backbones may increase their revenues. If backbone *A* does not upgrade its interconnections with *B* to keep pace with demand, customers of backbone *B* will experience degradation of access to customers of backbone *A*, and may choose to directly or indirectly become a paying customer of backbone *A* in order to increase access speeds.[34] As a result, the backbone may "balkanize;" with most customers maintaining direct connections to all of the largest backbones in order to ensure high-

---

[34] Customers of *B* may choose to directly connect with *A* as well as *B*, in which case they are "multihomed," or they may choose as their primary backbone a third backbone such as AboveNet that maintains direct connections to all backbones.

speed connections to the customers of each.[35] This situation is likely to be exacerbated by the quality of service issues raised in the next section.

B. Quality of Service

Trends in the industry are likely to change the status quo that currently prevails. Specifically, as customers become more sensitive to interconnection quality, quality degradation is likely to become more of an issue. New services are continually being added to the Internet. Many of these services, including IP telephony and video-conferencing, are real-time applications that are very sensitive to any delays in transmission. As a result, quality of service (QoS) is becoming a critical issue for backbones and ISPs.

First we consider whether two or more backbones may decide to establish high-quality interconnections over which they could guarantee QoS to their customers wishing to communicate with customers of the other backbone. Backbones face a number of private economic and technical considerations in making such decisions. The lack of such interconnection may have public consequences, however, as the Internet may fracture with each backbone only supporting new services between its own customers.

We first examine the economic incentives and effects of interconnection. The decision to interconnect for the provision of QoS services would appear to be relatively similar to the one that backbones currently make when deciding whether to peer with one another. The backbones each calculate whether the benefits of interconnecting with one or more other backbones would outweigh the costs. The benefits of interconnecting to exchange QoS traffic would flow from increasing the network of customers who can be contacted using services that rely on QoS; this would help attract new users, and encourage usage from existing users. The cost would come from a competitive network externality, as defined above; interconnection makes the other backbone more attractive to customers.

There is, however, an economic difference between current interconnection and interconnection for QoS purposes. The Internet services that interconnection enables today, such as email and Web access, are universally available, and no backbone or ISP could differentiate itself based on its provision of these services. Universal connectivity, however, is a legacy of the

---

35 *See* Frieden (1998). Professor Frieden raises issues similar to the ones dealt with in this paper, with a focus

cooperative spirit that characterized the Internet in its early days. In the commercial spirit that pervades the Internet today, backbones and ISPs may view the new services that rely on QoS as a means to differentiate themselves from their competitors. A firm that introduces these new services is going to be less willing to share the ability to provide these services with competitors, as it reduces the ability of the firm to charge a premium to its own customers. As a recent example, America Online (AOL) has been trying to defend its Instant Messaging service as a service available only to its own customers.[36]

There is a further problem associated with interconnection for QoS services. Assuming that QoS interconnection is implemented under the current settlement-free peering system, backbones will divide interconnection traffic into two types: originating traffic which they are paid by their customers to carry, and terminating traffic which they are not paid to carry. Under this system, backbones can send traffic to other backbones for termination free of charge, while in turn they must terminate traffic from other backbones without compensation. As a result, backbones will have no disincentive to sending traffic to other backbones, and also have no incentive to increase capacity to receive traffic from other backbones. As a result, QoS traffic will face congestion and may not provide satisfactory quality. Of course, similar problems exists today with the current peering system, as described above, leading to the current congestion, but given the premium that can be charged for QoS traffic and the high volume characterizing such services, the problem will be exacerbated. One solution to this problem is to implement a traffic-sensitive settlement system for such traffic.

If backbones are not able to overcome the economic, administrative, and technical hurdles enabling interconnection that supports QoS, then the Internet faces the risk of further balkanization. Each backbone will be its own network, only supporting certain services between its own customers. The result would be that network externalities, once taken for granted, would suddenly play a major role for consumers of Internet services. In the current environment of universal connectivity, consumers who simply want to send and receive email and surf on the Web can choose any retail provider without worrying about the choices of other consumers. With QoS, consumers may need to be aware of the choices of those with whom they wish to

---

on the effects of such balkanization on the universal access to the Internet, notably in rural areas.

[36] "Microsoft-AOL Instant Messaging Battle Continues," *Communications Daily*, August 19, 1999, v. 19, No. 160, at 3.

communicate when making their choice of Internet provider. A consumer who wishes to make IP telephony calls to friends and family, for instance, would need to be sure that the intended receivers are all connected to the same backbone in order to make high quality phone calls. Likewise, a business that wishes to use the Internet for videoconferencing must make sure that all relevant branches, customers, and suppliers are connected to the same backbone. Therefore, network externalities are present, as the backbone choice of each consumer influences the choices of other consumers.

As a result of the introduction of network externalities considerations, unless they can coordinate on the choice of one backbone, customers wishing to communicate with a wide variety of others may end up "subscribing" to competing backbones. This raises the specter of the early days of telephony, when competing telephone companies refused to interconnect, resulting in multiple telephones in many businesses and even some homes.[37] While users will not need two or more computers at their desks, they may need to arrange access to more than one backbone. Such balkanization may lead to calls for some form of regulation, such as implementing interconnection obligations on backbones. We argue below that such calls should not be honored.

IV. The Backbone as an Unregulated Information Service

For over thirty years, the Federal Communications Commission has maintained a categorical distinction between regulated communications services and unregulated computer-based services.[38] To understand why Internet backbone services are, and should continue to be, treated as unregulated enhanced services, it is important to understand two basic policies. First, it is important to understand the basic policies behind common carrier regulation, whereby carriers agree to the requirement that they provide service to all who make a reasonable demand for it in exchange for limitations on liability and, in some cases, a government-granted monopoly. For the telecommunications network, like the railroad and the telegraph before it, to grow into a healthy and vibrant universally available network, early regulatory invention was necessary. The

[37] As an example of this phenomenon, in 1910 Louisville, Kentucky was served by two local telephone companies; the Bell-licensed Cumberland Telephone and the independent Home Telephone company. More than 75% of the large businesses in Louisville had duplicate subscriptions, and 9% of homes had two telephones. See Milton L. Mueller, Jr., *Universal Service: Competition, Interconnection, and Monopoly in the Making of the American Telephone System*, MIT Press and AEI Press, 1997, at 83.

[38] *See also* Unregulation of the Internet.

telecommunications network was extraordinarily expensive to built and maintain, and the public benefits of ensuring that it was available for all to use meant that striking a "common carrier" bargain with telephone companies was a beneficial governmental intervention. Second, it is important to understand why enhanced services are not regulated as common carrier services. Soon after their introduction, the FCC determined that the competitive computer-based services market would remain open so long as the only essential input to such services – telecommunications capability – were available to providers of such services on a nondiscriminatory basis. Thus it was not necessary to impose common carrier regulations on the users of those telecommunications services as well as the providers. Such intervention, with the onerous implications of common carrier regulation, is only necessary in the rarest of circumstances. The Internet backbone market does not warrant such heavy-handed government regulation.

A. Principles of Common Carriage

The telephone network, much like the Internet, is only as useful to the end user as the number of fellow subscribers that end user can contact. In the early days of the telephone network, thirty-four states determined that mandating interconnection obligations was the best way to resolve disputes between the Bell System, the largest telephone company at the time, and smaller independent telephone companies, that had arisen between 1894 and 1906. It was not until 1910 that the Mann-Elkins Act extended the jurisdiction of the Interstate Commerce Commission to include telephone, telegraph, and cable companies.[39] In 1934, Congress established the Federal Communications Commission to regulate telecommunications common carriers.[40] The Communications Act of 1934 (the Communications Act) defined a common carrier as "any person engaged as a common carrier for hire."[41]

What is a common carrier, and what is special about such carriers that subjects them to special treatment under the law? Historically, common carriage obligations were imposed on shipowners, stable keepers, and inn owners pursuant to common law. In England, common carriers were defined as those entities that set out to serve the public and could not refuse to

---

[39] Mann-Elkins Act, Pub. L. No. 61-218, 36 Stat. 539 (1910).

[40] The Communications Act of 1934, the FCC's enabling statute, created the Commission "for the purpose of regulating interstate and foreign commerce in communication by wire and radio . . ." and established "the public interest, convenience, and necessity" as the principles that guide Commission decision making.

provide the service they held out to provide without justification. Thus, an innkeeper could not deny a room to a traveler if accommodations were available. At common law, common carriers were, in exchange for the requirement that they serve all upon reasonable demand, granted a limitation of liability arising from those services.

In this country, common carrier obligations were initially applied to the railroad, and subsequently to other transportation networks. Even before the passage of the Communications Act, the Supreme Court ruled that telegraph companies had a duty – arising out of the common law – to serve all customers in a nondiscriminatory manner as a common carrier.[42] Today, pursuant to the Communications Act, communications common carriers must offer service on demand to the public at large without unreasonable discrimination. Common carrier status for telephone companies meant the important protection from liability and, often, a government grant of monopoly (the monopoly granted local telephone companies, for example, lasted for nearly a century and was only lifted by Congress in 1996). For the public, such regulation brought the beneficial construction of a nationwide network of information highways while ensuring access to that network for all users.

Common carrier regulation is also an extraordinarily heavy government regulation. Not only must common carriers publish all of their rates in public tariffs from which they are generally forbidden to waiver, they must interconnect with all other common carriers who desire to exchange traffic (essentially a requirement to do business with any and all competitors), and must seek permission from the FCC before constructing or purchasing new interstate communications facilities.[43] Common carrier regulations such as these serve to protect the quality of the telephone network, ensure nondiscriminatory access to end users, and protect against anticompetitive behavior by telecommunications providers with market power.

B. Basic versus enhanced services

The absence of market power in the computer services industry led the FCC to conclude that imposing common carrier regulations would discourage innovation and distort the nascent data marketplace. While the government maintained a watchful regulatory eye over the

---

41 47 U.S.C. § 153(h) (1994).

42 *See Western Union Tel. Co. v. Call Publishing Co.*, 181 U.S. 92, 99-104 (1901).

43 Although this last requirement has been eased somewhat through the use of blanket licenses pursuant to section 214 of the Communications Act.

underlying transport, ensuring it was available to all wishing to use it, the competitive enhanced services market was able to flourish without onerous regulations impeding its growth.

Beginning with the *Computer Inquiry* proceeding in 1966, the FCC sought to avoid imposing unnecessary common carrier regulations on providers of computer services that relied on the nation's telecommunications infrastructure for transmission of those services, but did not themselves provide telecommunications services to the public. The result: the creation of "basic" and "enhanced" service categories, and the ability of enhanced service providers to avoid onerous regulation. As for the Internet backbone, immunity from regulation has allowed that market to grow since privatization in 1995 into a competitive market with a multitude of providers. Yet the backbone market is only four years old, and questions are arising as to whether the government can, or should, maintain a fully hands-off approach to backbone providers. Given the growing trend towards convergence of voice and data networks, it is increasingly difficult to maintain a distinct line between regulated and unregulated networks. An examination of the policy reasons behind the initial decision of the FCC not to regulate computer networks, and an examination of what imposition of those regulations would have meant for the Internet backbone, suggests that common carrier regulation of the backbone is neither warranted, nor appropriate, in the current environment. Rather, the nation's antitrust laws provide adequate protection against any future anticompetitive practices in the competitive backbone market.

In 1966, the FCC opened its first inquiry into the marriage of computers and communications, seeking to explore the regulatory and policy issues raised by the interdependence of the two technologies. In the opening lines of the inquiry, the Commission foreshadowed the incredible attributes of computer networks that would make the Internet such a valuable tool.

> The modern day electronic computer is capable of being programmed to furnish a wide variety of services, including the processing of all kinds of data and the gathering, storage forwarding, and retrieval of information – technical, statistical, medical, cultural, among numerous other classes. With its huge capacity and versatility, the computer is capable of providing its services to a multiplicity of users at locations remote from the computer. Effective use of the computer is, therefore, becoming increasingly dependent upon communication common carrier facilities and services by which the computers and the user are given instantaneous access to each other.[44]

---

44 *First Computer Inquiry*, 7 FCC 2d 11, para. 1.

The Commission determined that the policies and objectives of the Communications Act would best be served by allowing computer services to operate in an environment free from government regulation. The Commission devised rules, in effect to this day, that require common carriers to grant nondiscriminatory access to their networks to enhanced service providers. Mandating such nondiscrimination, the Commission concluded, was necessary because the computer-based service industry "cannot survive, much less develop further, except through reliance upon and use of communications facilities and services."[45]

In order to facilitate the implementation of its policies towards the nascent computer services industry, the Commission created the categories of "basic" and "enhanced" services. The basic services category denotes common carrier services subject to Title II of the Act.[46] The enhanced services category denotes those services

> offered over common carrier transmission facilities used in interstate communications, which employ computer processing applications that act on the format, content, code, protocol, or similar aspects of the subscriber's transmitted information, provide the subscriber additional, different, or restructured information; or involve subscriber interaction with stored information.[47]

Thus a basic service is a communications pathway, like a telephone line, while an enhanced service is a computer-enhanced service that operates via that communications pathway. Present day examples of enhanced services include voicemail services, gateway services and electronic publishing, Internet access services, and the Internet backbone.

The distinction between basic and enhanced services has permitted Internet backbone providers to escape regulation as a Title II common carrier. As discussed in greater detail below, the absence of common carrier regulations on Internet backbone providers is a deliberate policy in support of innovation in the data marketplace. So long as common carrier regulations assure a functional telecommunications network, there is no reason to impose such regulations on the users of that network. Internet backbone providers, like other enhanced service providers, purchase telecommunications capacity from a common carrier, add computer processing and other enhanced capabilities, and offer a service free from regulation intended to

---

45 *First Computer Inquiry*, Final Decision and Order, 28 FCC 2d at 7.

46 Basic services are defined as a common carrier offering of a pure "transmission capacity for the movement of information." *Amendment of Section 64.702 of the Commission's Rules and Regulations (Second Computer Inquiry)*, 77 FCC 2d 384, 419 (1980) (*Computer II Final Decision*).

ensure only that the underlying telecommunications network is available on a nondiscriminatory basis to those who wish to use it.

Why is the Internet backbone an enhanced service? The Internet backbone operates as an enhancement to the communications infrastructure. Backbone providers offer service via telecommunications, that is, they use telecommunications capability to send packets of customer-generated data across the backbone network, but they do not themselves offer that raw transmission capability. Rather, backbone providers either lease telecommunications capability from long distance telecommunications providers, or from affiliated telecommunications entities. In the case of MCI WorldCom, for example, the UUNET data subsidiary that operates MCI's large Internet backbone leases telecommunications transport capability from MCI's telephone company operations. The "enhancement" added to those telecommunications services by the backbone providers are "computer processing applications that act on the format, content, code, protocol, or similar aspects of the subscriber's transmitted information." For example, backbone operators utilize routing tables that act on the addressing information contained in packet headers. Backbone providers also conduct protocol conversion to ensure that information using different communications protocols can traverse disparate networks.

Of particular note for this discussion of the economics of the Internet backbone is the Commission's determination that there were, in the early 1970s, "no natural or economic barriers to free entry into the market for [enhanced] services."[48] So long as the underlying telecommunications network was available to enhanced service providers, there was no legal or economic reason to regulate enhanced services. We find that today the Commission's determination that enhanced services should be unregulated applies more than ever to Internet backbone services, given the recent increase in the number of facilities-based carriers capable of providing the underlying telecommunications transport capability necessary to operate a backbone.

C. Economics and Law Intersect

---

[47] 47 C.F.R. sec. 64.702(a).
[48] *First Computer Inquiry*, Tentative Decision, 28 FCC 2d at para. 18.

We find that it would be unwise and unjustified to impose any common carrier regulations on Internet backbone providers, given the implications of such regulation. The challenge for the common carrier to make a profit is clear: it cannot, due to nondiscrimination requirements, offer differentiating prices and service quality levels to different customers, nor can the carrier refuse any reasonable request for service. Rather, as a general matter the common carrier is obligated to serve all comers at the same price and quality levels. Unable to avoid unprofitable customers or to tailor packages to lower costs and maximize profits, common carriers often relied on government grants of monopoly to survive.

Internet backbone providers, if subjected to common carrier regulations, would be unable to offer the innovative data transport arrangements that are the hallmark of the Internet. For example, backbone providers would be unable to peer with each other, because nondiscriminatory pricing requirements would prevent the providers from negotiating private carriage pricing deals with customers. Backbone providers would similarly be unable to enter into private carriage arrangements that carried quality of service guarantees unavailable to other customers.

Is common carrier regulation of backbone providers necessary? An examination of the most basic of common carriage obligations – interconnection – suggests that the traditional justification for mandatory interconnection is not present in Internet backbone providers. Interconnection among carriers in a monopoly environment generally will not take place absent government regulation. Incumbents have no incentive to encourage competition by granting easy access to their networks to nascent competitors. The asymmetry of bargaining power requires government intervention to ensure the development of competition: the growing complexity of interconnection regulations formulated by the FCC and state commissions is a testament to the difficulty in fostering such regulations. If left entirely to its own devices, the monopolist has the ability and the incentive to thwart competition by refusing to interconnect pursuant to reasonable terms and conditions.

It is difficult to justify the imposition of interconnection requirements in a fully competitive marketplace. As discussed above, given the current competition in Internet backbone services, no single backbone would have the incentive to refuse to interconnect with any other backbone. This is, of course, only true in a fully competitive market. For example, Congress determined in 1996 that the long monopoly of incumbent local exchange carriers should end. In order to facilitate the development of competition in the local exchange market, Congress implemented several provisions to ensure that the monopolist local providers could not block competitors from offering service. Perhaps the most important of those regulations was the requirement, found in section 251 of the 1996 Act, that incumbent LECs offer nondiscriminatory interconnection terms and conditions to competitors. Absent such a requirement, incumbents would be able to deny competitors access to the nationwide network already built and controlled by the monopolists, access to which is crucial to the ability of competitors to offer service. Section 251 also requires incumbent LECs to lease parts of their networks to competitors at cost-based rates. These requirements are crucial to the development of a competitive telecommunications network, and the FCC rules implementing these requirements may some day be relaxed in the face of a fully competitive market.

What is clear, however, is that common carrier obligations impose huge costs on service providers, and cannot be applied selectively in a marketplace only to specific services or providers. In the early 1970s, for example, the FCC began a proceeding to attempt to reduce some of the onerous burdens placed on carriers by common carrier regulation. The FCC began by modifying its definition of common carrier services, which allowed carriers to offer more services outside of Title II regulation, offering the carriers flexible pricing and service deployment options. The District of Columbia Court of Appeals later ruled[49] that the FCC lacked the discretion to modify or waive the requirements of Title II, a holding affirmed by the Supreme Court.[50] Common carrier regulations, once applied to the Internet backbone, would impose requirements not necessary in a competitive marketplace.

D. Consumer Protection in the Backbone Market

Because the average consumer has only an indirect relationship with the backbone provider, it may be difficult for consumers to gauge discriminatory prices or terms and conditions that the backbone provider may be imposing. Although the backbone market is competitive today, the market may change in the future, and consumers need to be protected in the event unfair practices in the backbone market raise prices to monopoly levels. An adequate

---

49 *AT&T v. FCC*, 978 F.2d 727 (D.C. Cir. 1992).

50 *MCI v. AT&T*, 114 S.Ct. 2223 (1994) (FCC does not have authority to lift tariffing obligations from nondominant carrier).

enforcement measure, short of common carriage obligations, exists in the antitrust laws to guard against anticompetitive behavior. Because Internet backbone providers are “consumers” of telecommunications services, and because providers of those telecommunications services are subject to the nondiscrimination common carrier obligations discussed above, backbone providers should have little difficulty purchasing telecommunications capacity necessary for the provision of their services. But should backbone providers themselves collude in an anticompetitive way to foreclose competition in any of the ways outlined below, the solution to such problems should be found in the antitrust laws.

One area of potential unlawful behavior is in joint ventures. Backbone providers could find economies of scale by sharing transport capacity, routing functionality, and other aspects of service provision. Joint ventures can run afoul of antitrust laws, however, if the services offered in a joint venture would have been offered separately by the venture’s participants in competition with one another. Significant efficiencies could be realized by the parties to the joint venture, but such economizing could be an insufficient justification in the face of competition-reducing pricing schemes.

Further consolidation of the telecommunications service provider industry could lead to concerns over vertical arrangements. The backbone market today is fully competitive, with backbone providers utilizing their own telecommunications transport capability in conjunction with capacity leased from competitors. Such capacity is readily available today from a wide variety of providers, making entry into the backbone market relatively easy. Antitrust laws will serve to prevent an unlawful consolidation of the telecommunications transport market that could lead telecommunications providers to vertically integrate backbone services with common carrier services. In addition, existing common carrier regulations would prevent telecommunications carriers from refusing to provide transport sought by competing backbone providers.

Exclusive dealing arrangements are another potential threat to the competitive vibrancy of the backbone market. Backbone providers could enter into exclusive dealing arrangements with local Internet access providers, such that the backbone provider would provide only one access provider in a particular market with its services. Such an arrangement could potentially foreclose other access providers from competing, as they would have no access to the global Internet. In today's marketplace, however, such a scenario is unlikely. Because both the Internet access market and the backbone market are competitive, it is not likely that an exclusive arrangement between one backbone provider and one access provider would reduce competition, because other access providers would have the ability to seek backbone capacity from any of the numerous other providers in the market. In addition, the possibilities of entry into both the access market and the backbone market would not be foreclosed by such an arrangement.

E. Competition, not regulation

The Internet backbone market is a competitive market, and the assignment of common carrier regulations to that market would be destructive. Common carrier regulations serve their purpose when there is a risk to the public that services deemed of public necessity will not be available on a nondiscriminatory basis. Common carrier rules have helped make the public telecommunications network in this country widely available and competitively priced. As a result, that network is available to end users who desire to use it, and to other service providers who desire to combine its capacity with computer-based capabilities. Internet backbone providers take advantage of the open telecommunications network to purchase the communications services they need for transport and combine those services with computer processing functionality. The existence of multiple competitive providers of long distance telecommunications services – Qwest, Level 3, NEXTLINK, to name a few – suggests that the telecommunications component of those services is available from a wide variety of sources. The Internet backbone market is competitive in large part for that very reason: the necessary input – telecommunications – is available from a multitude of providers, making it difficult for backbone providers to foreclose entry into their market. Should backbone providers act in an anticompetitive manner, either by restricting access providers' ability to purchase backbone services or by squeezing competitors out of the market, the nation's antitrust laws are sufficient to protect against such behavior. It is clear that classifying Internet backbone services as common carrier services would serve none of the policy goals generally advanced by common carriage, and such unnecessary government regulation would cause significant harm to the backbone market.

V. Marketplace Solutions

We find that, in general, the marketplace has provided solutions quickly and effectively to challenges arising on the Internet, and that this is likely to continue. For instance, a number of solutions have arisen to congestion problems. As discussed above, backbones have moved from multilateral NAPs to bilateral private interconnections in order to avoid the congestion of the NAPs. Web providers have begun to use mirror sites and caching in order to spread traffic over a wider area, and move sites closer to end-users. In order to increase the speed of residential users' Internet access, firms have begun to provide high-speed access services such as xDSL and cable modems.

In terms of future challenges, we are confident that the marketplace will provide necessary solutions. As discussed above, we find that the Internet may begin to fracture in the future. Backbones and ISPs are beginning to provide QoS services to their customers. These services enable their providers to differentiate themselves from their competitors, and therefore, the providers may not be eager to interconnect with their competitors for purposes of making such services available to their competitors' customers. As these Internet QoS services begin to compete with traditional, regulated, telephony services, there may be a call to regulate these services by, for instance, mandating interconnection between backbone providers, as mentioned above. We argue that this should be avoided, in favor of allowing the marketplace to dictate interconnection and provision of such services. In order to explain this argument, we first examine the traditional rationale behind the regulation of network services.

On the supply-side, the provision of service in network industries such as telecommunications was traditionally considered to embody almost overwhelming economies of scale.[51] Economic theory and practice suggests that a natural monopolist is likely to arise in such industries, and this is considered efficient because duplicative facilities are not installed. However, without competitors, a natural monopoly can raise prices and/or reduce service quality. Governments have generally chosen to regulate such monopolists, in order to benefit from the efficiencies inherent in having one provider, while not incurring the corresponding costs that could be inflicted on consumers. As a result, rate and service regulations, such as the common carrier regulations described above, have traditionally been placed on natural monopolies in network industries. In addition, where the natural monopoly is not geographically pervasive, such as local telephony in the United States,[52] or where the services of the natural monopoly complement services provided in competitive markets, such as wireless telephony, interconnection is regulated so that the natural monopoly does not leverage market power further than necessary to provide the basic services. Such regulation is the norm not just for telecommunications, but other network industries such as water and electricity.

On the demand-side, however, telephony differs significantly from other regulated industries such as water and electricity. The difference is that consumers of telephony services enjoy network externalities, and therefore their benefits from owning a telephone increase, the larger the network of connected end-users. Consumers of water and electricity, on the other hand, do not enjoy such benefits. Given a choice of providers, or the option of self-provision, consumers of these services base their sourcing decision on cost and quality, rather than the size of a provider's network. It is important to note that regulation is generally based on the cost structure of network industries, rather than any of the demand-side network effects described here.

There are many examples of products that provide both direct and indirect network externalities that are not regulated. For instance, almost every consumer electronics product consists of a hardware/software system with indirect network externalities. The usefulness of compact disk players, personal computers, web browsers, and video cassette recorders depends to a great degree, if not totally, on the availability of compatible "software." The greater the number of users of the relevant "hardware," the more software will be available. Likewise, fax machines and email involve direct communications between end users with corresponding direct network externalities. In every example listed here, with the notable exception of personal computers, an unregulated marketplace led to a single standard.

The marketplace is quite successful at choosing standards allowing the products and services of different companies to be compatible with one another. Often, this is accomplished by a standards battle, such as the one waged between Betamax and VHS for the video cassette recorder standard. In other cases, one firm may create an adapter that enables its

---

[51] Lately, new technologies are reducing the economies of scale in a number of network industries such as telephony, leading to the pro-competitive, deregulatory provisions of the 1996 Act.

[52] The original Bell System, prior to divestiture in 1984, did not provide local service in many rural areas, as well as several cities and states, including most of Connecticut.

products to be compatible with the products of another firm. As an example, recently Microsoft provided software enabling its online customers to communicate via "instant messaging" with AOL's customers, a move that AOL attempted to rebuff.[53] One factor that influences firms to provide compatible standards is consumer demands. On the Internet in general, consumers have a high expectation of universal compatibility, for browsers, email, and now instant messaging, and firms that attempt to impose unilateral standards do so at their risk. Indeed, both AOL and Microsoft are now supporting a movement towards developing industry wide standard for instant messaging.[54]

Although the marketplace is remarkably successful at generating widely-adopted standards, it would be a mistake to conclude that this process is costless for consumers or firms. Purchasers of Sony's Betamax VCRs found it impossible to rent or buy movies after the VHS standard won the standards battle, while Sony was force to concede and begin selling its competitors' standard. The fax machine market was very slow to mature without a fixed standard, causing consumers to hold off on purchases while firms waited to enter the market. In each case, a government could theoretically have chosen a standard, thereby avoiding these costs. Nevertheless, governments, consumers, and firms rarely, if ever, call for government intervention in these cases.

The marketplace is the preferred means for setting standards in most industries and for most products for a variety of reasons. First, an open marketplace for standards leads to healthy competition, and often "second-mover" standards are able to overcome an industry leader by embodying their standards in better products, or more creative marketing of these products. For instance, VHS was able to overcome the Betamax lead to become the industry standard. Second, innovators such as Apple Computer may target new standards at niche markets, with consumers benefiting from the resulting variety. Such variety and innovation may not occur if a standard is chosen by non-market means. Therefore, while the marketplace may increase short-run costs involved with adopting new standards, the long run marketplace benefits of competition and innovation more than make up for any short-run costs.

[53] Instant messaging software enables users to "chat" with one another online by sending instantaneous messages. *See* Wylie Wong and Sandeep Junnarkar, "Net messaging standards war brewing?," CNET News.com, July 23, 1999 http://www.news.com:80/News/Item/0,4,39649,00.html?st.ne.ni.rel.

[54] "Microsoft-AOL Instant Messaging Battle Continues."

We conclude that the marketplace will continue to provide solutions to any and all challenges that arise as the Internet develops current services and extends to new services. Consumers, with expectations of universal connectivity and basic compatibility, are likely to ensure that backbones continue to interconnect and provide ever-increasing functionality. Although firms may limit offerings of new QoS services to their own customers, other marketplace solutions are available to ensure that consumers can remain connected to the full Internet. For instance, INTERNAP is a firm that connects with all major backbones, enabling its customers to communicate directly with all customers of the major backbones without themselves becoming customers of each major backbone. Firms like this may alleviate the costs of any Internet balkanization, if not preventing it altogether.

VI. Conclusion

As technologies converge, the Internet is increasingly being used to provide services that resemble traditional telecommunications services currently provided by regulated common carriers. As a result, there may increasingly be calls to impose common carrier regulations, including interconnection obligations, on Internet service providers and backbones. Fueling such demands, Internet backbones may increasingly not be interconnected in a fashion that would enable these services to be provided between networks, for a variety of reasons, including technical or competitive reasons. This paper argues that any regulation of Internet backbones is unnecessary and would harm the vibrant Internet industry. The Internet backbone, riding on a competitive network of facilities, is itself competitive, as opposed to the monopolistic networks that historically warranted common carrier regulation. The famously innovative nature of the Internet will find solutions to any technical roadblocks that may arise on the way to high-speed interconnection, and consumers, accustomed to universal connectivity, are bound to demand and receive interconnection if competitive considerations are in the way.