

THE ECONOMICS OF THE INTERNET BACKBONE

NICHOLAS ECONOMIDES

New York University, New York

Contents

1. Competition among Internet backbone service providers	375
1.1. Internet backbone services	375
1.2. Interconnection	375
1.3. The transit and peering payment methods for connectivity	379
1.4. Conduct of Internet backbone service providers	382
1.4.1. Pricing of transport services in the backbone networks	382
1.4.2. ISP multihoming; Additional demand responsiveness to price changes	383
2. Structural conditions for Internet backbone services; Negligible barrier to entry and expansion	385
2.1. The markets for raw transport capacity and other inputs to Internet transport services	385
2.2. Ease of expansion and entry	386
2.3. Public standards and protocols on the Internet	386
3. Potential for anticompetitive behavior on the Internet backbone	388
4. Network externalities and the Internet	388
4.1. Procompetitive consequences of network externalities	390
4.2. Conditions under which network externalities may inhibit competition	391
5. Network externalities and competition on the Internet	392
5.1. Conditions necessary for the creation of bottlenecks fail on the Internet	392
5.2. Bottlenecks such as the ones of the local exchange telecommunications network do not exist on the Internet	393
6. Strategies that a large IBP might pursue	394
6.1. Raising the price of transport	394
6.2. Discriminatory price increases directed simultaneously against all backbone rivals	397
6.3. Raising rivals' costs and degrading connectivity	398
6.3.1. Terminating interconnection simultaneously with all rivals (refusal to deal)	398

6.3.2. Degrading interconnection simultaneously with all rivals	399
6.3.3. Sequential attacks on rivals	400
7. Conclusions	407
Appendix	407
References	410

1. Competition among Internet backbone service providers

1.1. Internet backbone services

The Internet is a global network of interconnected networks that connect computers. The Internet allows data transfers as well as the provision of a variety of interactive real-time and time-delayed telecommunications services. Internet communication is based on common and public protocols. Hundreds of millions of computers are presently connected to the Internet. Figure 1 shows the expansion of the number of computers connected to the Internet.

The vast majority of computers owned by individuals or businesses connect to the Internet through commercial Internet Service Providers (ISPs)¹. Users connect to the Internet either by dialing their ISP, connecting through cable modems, residential DSL, or through corporate networks. Typically, routers and switches owned by the ISP send the caller's packets to a local Point of Presence (POP) of the Internet². Dial-up, cable modem, and DSL access POPs as well as corporate networks dedicated access circuits connect to high-speed hubs. High-speed circuits, leased from or owned by telephone companies, connect the high-speed hubs forming an 'Internet Backbone Network.' See Figure 2.

Backbone networks provide transport and routing services for information packets among high-speed hubs on the Internet. Backbone networks vary in terms of their geographic coverage. *Boardwatch magazine* has listed the following national backbones³ in Table 1. Market shares of national backbones are listed in Table 2 based on a 1999 projection. In papers filed in support of the merger of SBC and AT&T as well as the merger of Verizon with MCI, there was mention of two recent traffic studies by RHK. These studies showing traffic for 2004, summarized in Table 3, show a dramatic change in the ranking of the networks, with AT&T now being first and MCI fourth. They also show that now a much bigger share of traffic (over 40 percent) is carried by smaller networks. These latest traffic studies show that the concern of the EU and the USDOJ that the Internet backbone market would tilt to monopoly were proved to be overstated.

1.2. Interconnection

There is wide variance of ISPs in terms of their subscriber size and the network they own. However, irrespective of its size, an ISP needs to interconnect with other

¹ Educational institutions and government departments are also connected to the Internet but do not offer commercial ISP services.

² Small ISPs may not own routers and switches, but rather just aggregate traffic at modem banks and buy direct access to a larger ISP.

³ See <http://www.boardwatch.com/isp/summer99/backbones.html>. *Boardwatch magazine* also lists 348 regional backbone networks.

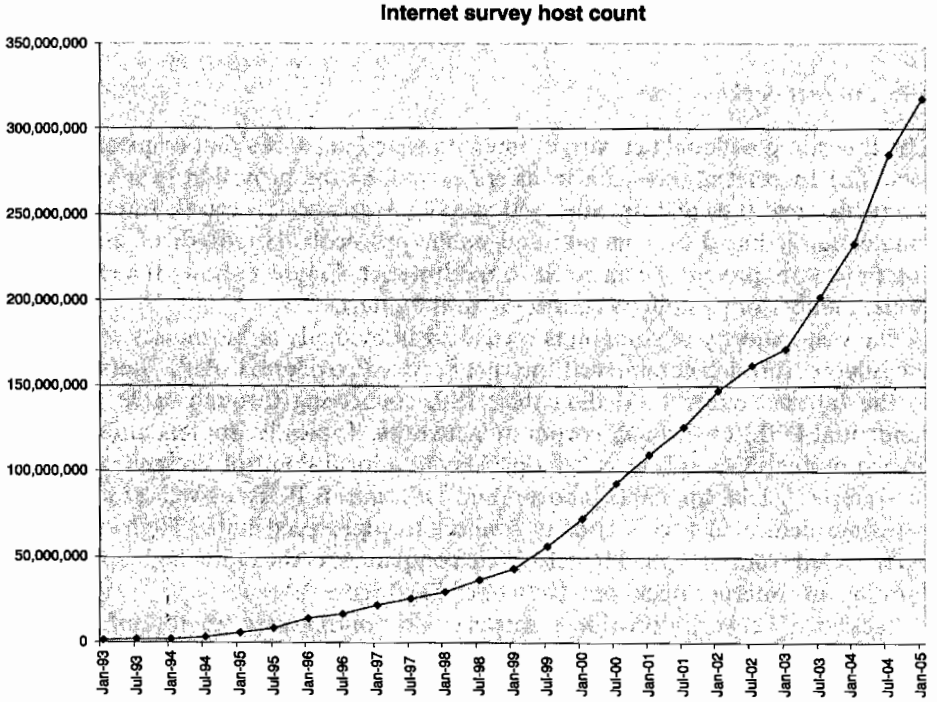


Fig. 1.

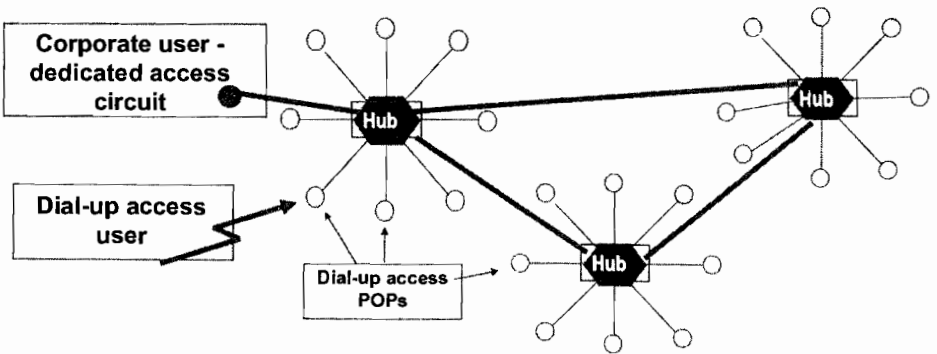


Fig. 2.

Table 1
Partial list of national Internet backbones

@Home Network	Intermedia Business Internet
1 Terabit	Internet Access/GetNet
Abovenet	Internet Services of America
Apex Global Information Services (AGIS)	IXC Communications, Inc
AT&T Networked Commerce Services	Level 3
Cable & Wireless, USA	MCI WorldCom—Advanced Networks
CAIS	MCI WorldCom—UUNET
Concentric	NetRail
CRL Network Services	PSINet, Inc.
Digital Broadcast Network Corp.	Qwest/Icon CMT
Electric Lightwave	Rocky Mountain Internet/DataXchange
EPOCH Networks, Inc.	Savvis Communications Corporation
e.spire	ServInt
Exodus	Splitrock Services
Fiber Network Solutions	Sprint IP Services
Frontier Global Center	Teleglobe
Globix	Verio
GTE Internetworking	Visinet
GST Communications	Vnet
IBM Global Services	Winstar/Broadband
ICG/Netcom Online	ZipLink
IDT Internet Services	

Table 2
Market shares of national Internet backbones

Market Share	1997	1999	2001 (projected in 1999)	2003 (projected in 1999)
MCI WorldCom	43%	38%	35%	32%
GTE-BBN	13%	15%	16%	17%
AT&T	12%	11%	14%	19%
Sprint	12%	9%	8%	7%
Cable & Wireless	9%	6%	6%	6%
All Other	11%	21%	22%	19%
Total	100%	100%	100%	100%

Note: *Hearing on the MCI WorldCom-Sprint Merger Before the Senate Committee on the Judiciary*, Exhibit 3 (Nov 4, 1999) (Testimony of Tod A. Jacobs, Senior Telecommunications Analyst, Sanford C. Bernstein & Co., Inc.), Bernstein Research, *MCI WorldCom* (March 1999) at p. 51.

ISPs so that its customers will reach all computers/nodes on the Internet. That is, interconnection is necessary to provide universal connectivity on the Internet, which is demanded by users. Interconnection services at Network Access Points

Table 3
Carrier traffic in petabytes per month in 2004

Company	Traffic				Market share among all networks
	1Q2004	2Q2004	3Q2004	4Q2004	4Q2004
A (AT&T)	37.19	38.66	44.54	52.33	12.58%
B	36.48	36.50	41.41	51.31	12.33%
C	34.11	35.60	36.75	45.89	11.03%
D (MCI)	24.71	25.81	26.86	30.87	7.42%
E	18.04	18.89	21.08	25.46	6.12%
F	16.33	17.78	17.47	19.33	4.65%
G	16.67	15.04	14.93	15.19	3.65%
Total traffic top 7 networks	183.53	188.28	203.04	240.38	57.78%
Total traffic all networks	313	313	353	416	100%

Note: Data from *RHK Traffic Analysis – Methodology and Results*, May 2005. The identities of all networks are not provided, but it is likely that B, C, E, and F are Level 3, Quest, Sprint, and SBC in unknown order.

(NAP) and Metropolitan Area Exchanges (MAEs)⁴ are complementary to Internet transport. In a sense, the Internet backbone networks are like freeways and the NAPs like the freeway interchanges.

Internet networks in two ways:

1. Private bilateral interconnection; and
2. Interconnection at public NPAs.

Private interconnection points and public NAPs are facilities that provide collocation space and a switching platform so that networks are able to interconnect. Network Access Points' services are not substitutes for ISP, or for transport services. Rather, they are a complement to ISP services and to transport services. The NAPs allow networks to interconnect more easily by providing the necessary space and platform.

Interconnection at NAPs is governed by bilateral contracts of the parties. Some NAPs, such as the London Internet Exchange (LINX) facilitate such negotiations by posting a set of common rules and standard contracts, which may be used by its members in their bilateral negotiations. Interconnection of two networks X and Y at a NAP is governed by a contract between networks X and Y. Other NAPs such as the ones owned by MCI do not dictate the terms of contracts between third-party networks⁵.

⁴ The NAPs run by MCI are called Metropolitan Area Exchanges (MAEs).

⁵ In particular, interconnection at a NAP owned or controlled, for example, by MCI, does not imply or require a barter (peering) or transit arrangement between UUNET and networks X and Y.

Table 4
MAEs' capacity growth and utilization

	Capacity (Gbps)			Sales (Gbps)
	1997	1999	January 2000	January 2000
MAE-East	7.6	11.2	19.9	11.4
MAE-West	4.3	11.2	19.9	11.8
MAE-Dallas	N/A	7.5	7.5	2.6

Recently, there has been a significant increase in the number of NAPs as well as expansion and renewal of preexisting NAPs. In 1995, there were only 5 NAPs, MAE East, MAE West, NY (Sprint), Chicago (Ameritech), and Palo Alto (PacBell). In 1999, there were 41 NAPs in the United States (including 5 MAEs), and 40 European NAPs (including 2 MAEs) and 27 Asia-Pacific NAPs⁶. Table 4 shows the capacity expansion of NAPs from 1997 to January 2000. The fifth column of Table 4 shows capacity in January 2000. It is evident that there is very significant spare capacity. A partial list of NAPs in North America and the rest of the world is provided by the Exchange Point Network at <http://www.ep.net/ep-main.html>⁷.

1.3. The transit and peering payment methods for connectivity


Internet networks have contracts that govern the terms under which they pay each other for connectivity. Payment takes two distinct forms: (i) payment in dollars for

⁶ Source <http://www.ep.net>.

⁷ The exchange point information net at http://www.ep.net/naps_na.html lists the following NAPs in North America: East Coast: ATL-NAP Atlanta; BNAP – Baltimore NAP; Louisville-nap.net; MAGPI – a Mid Atlantic Gigapop for Internet2; MassachusettsIX; NY6IX – A New York IPv6 exchange; NYIIX – New York International Internet Exchange (Telehouse); Nashville Regional Exchange Point; Nap of the Americas; MetroIX; Philadelphia Internet Exchange; Pittsburgh Internet Exchange; Research Triangle Park; Sprint NAP (Pennsauken NJ); Vermont ISP Exchange; Blacksburg Electronic Village – VA. West Coast: AMAP – Anchorage Metropolitan Access Point; Ames Internet Exchange; COX – Central Oregon Internet Exchange; HIX – Hawaii Internet Exchange; LAIIX – Telehouse Los Angeles; LAAP – A Los Angeles Exchange, includes MAE-LA; Northwest Access Exchange – Portland; OIX – Oregon Internet Exchange; PACIFIC WAVE – Pacific Wave Exchange; SBC-Oakland; SD-NAP – San Diego (Caida); SIX – Seattle Internet Exchange. The South: New Mexico Internet Exchange; IX New Mexico; TTI – The Tucson Interconnect; Yellowstone RIE. The Middle American Exchange Points: CMH-IX – Columbus Internet Exchange; D-MIX – Dayton OH; DIX – Denver Internet Exchange; IndyX – Indianapolis Data Exchange; Nashville CityNet; Ohio Exchange; RMIX Rocky Mountain Internet eXchange; SBC-Chicago STAR TAP (12 GigaPOP); St. Louis, Mo.; Utah REP. Canada: BC Gigapop; CA/NAP Canada/Toronto Exchange; CANIX: Originally CA*net Sponsored; MIX – Montreal Internet Exchange; The Nova Scotia Internet eXchange; Ottawa Internet eXchange; Toronto Internet Exchange.

“transit”; and (ii) payment in kind (i.e., barter, called ‘peering’). Connectivity arrangements among ISPs encompass a seamless continuum, including ISPs that rely exclusively on transit to achieve connectivity, ISPs that use only peering to achieve connectivity, and everything in between. Although there are differences between transit and peering in the specifics of the payments method, and transit includes services to the ISP not provided by peering, it should be made clear that these two are essentially alternative payment methods for connectivity⁸. The transport and routing that backbone networks offer do not necessarily differ depending on whether cash (transit) or barter (peering) is used for payment. The same transport and routing between customers of the two networks can be obtained by purchase, or through barter for other transport services.

Under transit, a network X connects to network Y with a pipeline of a certain size, and pays network Y for allowing X to reach all Internet destinations. Under transit, network X pays Y to reach not only Y and its peers, *but also any other network*, such as network Z by passing through Y, as in the diagram below.



Under peering, two interconnecting networks agree not to pay each other for carrying the traffic exchanged between them as long as the traffic originates and terminates in the two networks. Referring to the diagram above, if X and Y have a peering agreement, they exchange traffic without paying each other as long as such traffic terminating on X originates in Y, and traffic terminating on Y originates in X. If Y were to pass to X traffic originating from a network Z that was not a customer of Y, Y would have to pay a transit fee to X (or get paid a transit fee by X, i.e., it would not be covered by the peering agreement between X and Y).

Although the networks do not exchange money in a peering arrangement, the price of the traffic exchange is not zero. If two networks X and Y enter into a peering agreement, it means that they agree that the cost of transporting traffic from X to Y and vice versa that is incurred within X is roughly the same as the cost of transporting traffic incurred within Y. These two costs have to be roughly equal if the networks peer, but they are not zero.

The decision as to whether interconnection takes the form of peering or transit payment is a commercial decision. Peering is preferred when the cost incurred by X for traffic from X to Y and Y to X is roughly the same as the cost incurred by Y for the same traffic. If not, the networks will use transit. As is explained below, the decision of whether to peer or not depends crucially on the geographic coverage of the candidate networks.

Generally, peering does not imply that the two networks should have the same size in terms of the numbers of ISPs connected to each network, or in terms of the

⁸ Transit customers receive services, such as customer support, DNS services, etc., that peering networks do not receive.

traffic that each of the two networks generate⁹. If two networks, X and Y, are similar in terms of the types of users to whom they sell services, the amount of traffic flowing across their interconnection point(s) will be roughly the same, irrespective of the relative size of the networks. For example, suppose that network X has 10 ISPs and network Y has 1 ISP. If all ISPs have similar features, the traffic flowing from X to Y is generally equal to the traffic flowing from Y to X¹⁰.

What determines whether a peering arrangement is efficient for both networks is the *cost* of carrying the mutual traffic within each network. This cost will depend crucially on a number of factors, including the geographic coverage of the two networks. Even if the types of ISPs of the two networks are the same as in the previous example (and therefore the traffic flowing in each direction is the same), the cost of carrying the traffic can be quite different in network X from network Y. For example, network X (with the 10 ISPs) may cover a larger geographic area and have significantly higher costs per unit of traffic than network Y. Then network X would not agree to peer with Y. These differences in costs ultimately would determine the decision to peer (barter), or receive a cash payment for transport.

Where higher costs are incurred by one of two interconnecting networks because of differences in the geographic coverage of each network, peering would be undesirable from the perspective of the larger network. Similarly, one expects that networks that cover small geographic areas will only peer with each other. Under these assumptions, who peers with whom is a consequence of the extent of a network's geographic coverage, and may not have any particular strategic connotation¹¹.

In summary, whether two interconnecting networks use peering (barter), or cash payment (transit) does not depend on the degree of competition among backbone services providers. In particular, the presence of peering is not necessarily a sign of intense or weak competition, nor would the replacement of peering by cash pricing necessarily be a sign of diminished or increased competition. Moreover, as the analysis above shows, generally, an ISP's decision not to peer reflects

⁹ For example, MCI WorldCom has peering arrangements with a number of smaller networks. See Letter from Sue D. Blumenfeld, Attorney for Sprint Corporation, and A. Richard Metzger, Jr., Attorney for MCI WorldCom, Inc. to Magalie Roman Salas, FCC, CC Docket No. 99-333 (dated January 14, 2000) at p. 20.

¹⁰ Suppose the larger network has 10 ISPs with 10 Websites per ISP and a total of 1000 users, and it interconnects with a smaller network with 1 ISP with 10 Websites and a total of 100 users. For simplicity, suppose that every user visits every Website. Then the smaller network transmits $100 \times 10 \times 10 = 10,000$ site-visits to the larger network, and the larger network transmits $1000 \times 1 \times 10 = 10,000$ site-visits to the smaller network. Thus, the traffic across networks of different sizes is the same if the types of ISPs and users are the same across networks.

¹¹ Milgrom et al. (2000) shows how peering (with no money changing hands) can emerge under some circumstances as an equilibrium in a bargaining model between backbones.

its assessment that the average costs of transport within one network are larger than the average costs of transport within the other network. Thus, refusal to peer is not inherently an anticompetitive act; it can be a consequence of some networks being much larger than others in terms of geographic coverage.

1.4. Conduct of Internet backbone service providers

1.4.1. Pricing of transport services in the backbone networks

The author first discusses business conduct of Internet backbone service providers. Structural conditions for Internet backbone services (discussed in the next section) ensure negligible barriers to entry and expansion and easy conversion of other transport capacity to Internet backbone capacity. As discussed in the next section, raw transport capacity as well as Internet transport capacity has grown dramatically in the last four years. Transport capacity is a commodity because of its abundance.

The business environment for Internet backbone services is competitive. Generally, ISPs buying transport services face flexible transit contracts of relatively short duration. Backbones do not impose exclusivity of service on their customers. For example, UUNET (MCI) does not require that it be the exclusive Internet transport provider to its ISP customers.

Often an ISP buys from a backbone bandwidth of a certain capacity that allows it to connect to the whole Internet (through a 'transit' payment). The bandwidth capacity and speed of the connecting pipe vary widely and depend on the demand for transport that an ISP wants to buy from a particular backbone. Price lists for various bandwidth capacities are printed in *Boardwatch magazine*. The strength of competition among the various backbone providers is evidenced in the small, or nonexistent differences in the prices for various bandwidth capacities. For example, Table 5 shows the prices for AT&T and UUNET (MCI) for various bandwidth capacities as reported by the latest edition of *Boardwatch magazine* (August 1999). Despite the fact that AT&T's backbone business was significantly smaller than UUNET's, their prices are identical for most bandwidths, and when they differ, the differences are very small. Many other providers of various sizes have very similar prices as reported in *Boardwatch magazine*¹².

As the expected growth of the Internet in the mid to late 1990s of 400 percent a year in terms of bits transferred was not realized in the post 1999 period, and instead a growth of only about 100 percent a year was realized, transit prices fell. As an example, Table 6 compares the AT&T prices for the same connectivity in 1999 and 2001.

¹² As *Boardwatch Magazine* reports in the 1999 and subsequent editions, prices for the same connectivity were very comparable for a large array of services among large IBPs.

Table 5
Comparison of early 1999 monthly prices of AT&T and UUNET (MCI) for U.S. DS3s (T3s)

Service	AT&T	UUNET	Price difference = UUNET-AT&T
Burstable 0–6 Mbps	\$12,500	\$12,000	\$500
Burstable 6.01–7.5 Mbps	\$14,000	\$14,000	\$0
Burstable 7.51–9 Mbps	\$17,000	\$17,000	\$0
Burstable 9.01–10.5 Mbps	\$19,000	\$19,000	\$0
Burstable 10.51–12 Mbps	\$22,000	\$22,000	\$0
Burstable 12.01–13.5 Mbps	\$26,000	\$26,000	\$0
Burstable 13.51–15 Mbps	\$29,000	\$29,000	\$0
Burstable 15.01–16.5 Mbps	\$32,000	\$32,000	\$0
Burstable 16.51–18 Mbps	\$37,000	\$37,000	\$0
Burstable 18.01–19.5 Mbps	\$43,000	\$43,000	\$0
Burstable 19.51–21 Mbps	\$48,000	\$48,000	\$0
Burstable 21.01–45 Mbps	\$55,000	\$55,500	\$500

Note: *Boardwatch Magazine's* Directory of Internet Service Providers, 11th Edition, 1999.

Table 6
Comparison of 1999 and 2001 monthly prices of AT&T for U.S. DS3s (T3s)

Service	Year 1999	Year 2001	Percentage price difference ($P_{2001} - P_{1999}$)/ P_{2001}
Burstable 0–6 Mbps	\$12,500	\$6550	-47.60%
Burstable 6.01–7.5 Mbps	\$14,000	\$8150	-41.79%
Burstable 7.51–9 Mbps	\$17,000	\$9250	-45.59%
Burstable 9.01–10.5 Mbps	\$19,000	\$10,150	-46.58%
Burstable 10.51–12 Mbps	\$22,000	\$11,050	-49.77%
Burstable 12.01–13.5 Mbps	\$26,000	\$11,950	-54.04%
Burstable 13.51–15 Mbps	\$29,000	\$12,850	-55.69%
Burstable 15.01–16.5 Mbps	\$32,000	\$13,600	-57.50%
Burstable 16.51–18 Mbps	\$37,000	\$14,350	-61.22%
Burstable 18.01–19.5 Mbps	\$43,000	\$15,100	-64.88%
Burstable 19.51–21 Mbps	\$48,000	\$15,850	-66.98%
Burstable 21.01–45 Mbps	\$55,000	\$31,050	-43.55%

Notes: *Boardwatch Magazine's* Directory of Internet Service Providers, 11th and 13th Edition, 1999 and 2001.

1.4.2. ISP multihoming; Additional demand responsiveness to price changes

Internet Service Providers are not locked-in by switching costs of any significant magnitude. Thus, ISPs are in good position to change providers in response to any increase in price, and it would be very difficult for a backbone profitably to increase price. Moreover, a large percentage of ISPs has formal agreements that allow them to route packets through several backbone networks and are able to control the way the traffic will be routed (multihoming). Table 7 shows that, in

Table 7
Additional backbone connections held by multihoming ISPs

Year	# ISPs	Number of backbone connections sold to ISPs	Share of additional connections sold to multihoming ISPs
1997	4354	5739	24%
1998	4470	5913	24%
1999	5078	8950	43%

Note: *Boardwatch Magazine's* Directory of Internet Service Providers, Fall 1997, p. 6. *Boardwatch Magazine's* Directory of Internet Service Providers, Winter 1998, p. 5. *Boardwatch Magazine's* Directory of Internet Service Providers, 11th Edition, 1999, p. 4. The last column is calculated as the difference between the third and the second columns divided by the third column, for example, for 1999, $(8950-5078)/8950 = 43.26\%$ rounded to 43%.

1999, additional (i.e., second or subsequent) connections sold to multihoming ISPs amounted to 43 percent of all ISP connections to backbones. One of the reasons for the increase in multihoming is likely the decrease in the cost. The cost of customer routers that are required for ISP multihoming has decreased from \$10,000 to \$2000–\$3000¹³. An additional reason for an ISP to multihome is that it increases the ability of the ISP to route its traffic to the lowest-priced backbone, as discussed in the next section.

When an ISP reaches the Internet through multiple backbones, it has additional flexibility in routing its traffic through any particular backbone. A multihoming ISP can easily reduce or increase the capacity with which it connects to any particular backbone in response to changes in prices of transit. Thus, multihoming increases the firm-specific elasticity of demand of a backbone provider. Therefore, multihoming severely limits the ability of any backbone services provider to profitably increase the price of transport. Any backbone increasing the price of transport will face a significant decrease in the capacity bought by multihoming ISPs.

Large Internet customers also use multiple ISPs, which is called 'customer multihoming.' They have chosen to avoid any limitation on their ability to switch traffic among suppliers even in the very shortest of runs. Customer multihoming has similar effects as ISP multihoming in increasing the firm-specific elasticity of demand of a backbone provider and limiting the ability of any backbone services provider to profitably increase the price of transport.

New technologies of content delivery that utilize distributed storage of Web-based content on various locations on the Internet reduce the need for backbone network transport. 'Caching' stores locally frequently requested content. 'Mirroring' creates a replica of a Website. Intelligent content distribution,

¹³ Source: *Boardwatch Magazine's* Directory of Internet Service Providers, 11th Edition, 1999.

implemented among others, by Akamai Technologies¹⁴, places its servers closest to the end users inside an ISP's network. Intelligent content distribution technology assesses the fastest route on the Internet for content access, and delivers content faster to end users. Placing content delivery close to end users and optimizing content delivery through intelligent content distribution, caching, and mirroring reduces in effect the demand for Internet transport services and the ability of backbone providers to affect the transit price.

2. Structural conditions for Internet backbone services; Negligible barrier to entry and expansion

2.1. The markets for raw transport capacity and other inputs to Internet transport services

Almost all Internet transport uses fiber-optic transmission capacity, which is based on a well-known and easily available technology¹⁵. There are no significant barriers to entry in the supply of additional raw transmission capacity. Fiber transmission capacity is essentially fungible, and the same physical networks can be used for the transmission of voice, Internet traffic, and data by using different protocols.

Fiber that will not be needed by an Internet transport supplier can be leased, or sold for nonInternet uses. The same fiber and electronics are used for both circuit switched and packet switched networks, which can each transport both voice and data. Before construction, the operator has a completely open choice between creating either a circuit switched or a packet switched network. Only the interface differs between voice and data applications. Once capacity is in place, there are small costs of converting from one use to the other. Moreover, capacity can be upgraded in small steps so that fiber networks can respond flexibly to increasing capacity requirements.

Fiber capacity has grown rapidly and is expected to grow for the indefinite future. Because there is always new capacity in the planning stage, no operator needs to consider switching the use of existing capacity. As a result, fiber capacity is not in any way a barrier to entry in Internet transport¹⁶.

¹⁴ Akamai was founded in 1998 and made a \$234M initial public offering in October 1999. Akamai has industry relationships with AT&T, BT plc, DIGEX, Global Center, GTEI, Lycos, Microsoft, PSINet, Qwest, Real Networks, Telecom Italia, Teleglobe, Universo Online, UUNET, and Yahoo!, among others.

¹⁵ The transport and switching technologies are available from firms that do not sell backbone transport or ISP services.

¹⁶ In the early stages of Internet expansion and given the explosive growth that was anticipated then, the possibility of a future backbone capacity shortage may have bid up the value of firms with installed Internet backbone capacity and may explain the price that WorldCom paid for MFS and implicitly UUNET. This should be seen in the context of a real options analysis. See Economides (1999a,b) and Hubbard and Lehr (2000).

In order to build or expand Internet backbone capacity, besides fiber-optic cable, networks need routers and switches. Routers and switches are readily available from a variety of third-party suppliers. Fiber capacity can be leased, and there is no shortage of capacity that would constrain the ability of smaller networks, or new entrants to expand capacity or enter the market. Fiber networks can add leased capacity, or increase their capacity by deploying new technologies such as Dense Wave Division Multiplexing (DWDM). The construction of fourth-generation fiber-optic networks, deploying the latest technology, promises an abundance of capacity that appears to be able to accommodate the very rapid growth in capacity demand that has been the hallmark of the Internet market to date.

2.2. *Ease of expansion and entry*

National, international, and regional long-haul fiber-optic transmission capacity has increased very rapidly, both as a result of expansion of networks of incumbents, such as AT&T, MCI, Sprint, and GTE but also as a result of entry of a number of carriers that created new networks, including Quest, Level 3, Williams, and others. The FCC's *Fiber Deployment Update* reports that total fiber system route miles of interexchange carriers increased by two-thirds between 1994 and 1998¹⁷. After 1998, the FCC discontinued the publication of this report. However, data reported by Besen and Brenner (2000)¹⁸ and Hogendorn (2004) supports the conclusion that the capacity of long-haul fiber is increasing in an accelerated rate.

As evidence of ease of entry, the number of North American ISPs more than tripled in the years 1996–1999, and has continued thereafter. The number of North American backbone providers has grown almost fivefold in the same period. These statistics are shown in Tables 8 and 9.

Bandwidth and equipment costs have decreased and continue to decrease. Hence, access to fiber capacity is unlikely to be an impediment to sellers wishing to upgrade their networks, or to new competitors wishing to enter the market.

2.3. *Public standards and protocols on the Internet*

In markets, where the incumbent has a proprietary standard and an entering rival must promote an incompatible alternative standard—as in operating systems for personal computers—standards can be used to create a barrier to entry. However, in markets where all rivals use the same public standard, no such barrier exists, or

¹⁷ See Jonathan M. Kraushaar, *Fiber Deployment Update: End of Year 1998*, FCC, Industry Analysis Division, Common Carrier Bureau, Table 1.

¹⁸ See Declaration of Stanley Besen and Steven Brenner, March 20, 2000.

Steering Group, and conducted by the Internet Engineering Task Force. In considering changes in standards, these groups require mandatory disclosure of any proposed change before it gets considered, so no proprietary standard can be introduced²⁰.

3. Potential for anticompetitive behavior on the Internet backbone

Some have proposed²¹ that the existence of network effects creates a grave danger that the Internet backbone will quickly become monopolized once the largest Internet backbone provider becomes 'large enough.' Various theories have been proposed of how this could be done. The author first discusses the general context in which network effects affect competition on the Internet, and subsequently discusses the specific theories.

4. Network externalities and the Internet

Like any network, the Internet exhibits network externalities. Network externalities are present when the value of a good or service to each consumer rises as more consumers use it, everything else being equal²². In traditional telecommunications networks, the addition of a customer to the network increases the value of a network connection to all other customers, since each of them can now make an extra call. On the Internet, the addition of a user potentially

1. adds to the information that all others can reach;
2. adds to the goods available for sale on the Internet;
3. adds one more customer for e-commerce sellers; and
4. adds to the collection of people who can send and receive e-mail, or otherwise interact through the Internet.

Thus, the addition of an extra computer node increases the value of an Internet connection to each connection.

In general, network externalities arise because high sales of one good make complementary goods more valuable. Network externalities are present not only in traditional network markets, such as telecommunications, but also in many other markets. For example, an IBM-compatible PC is more valuable if there are

²⁰ "No contribution that is subject to any requirement of confidentiality or any restriction on its dissemination may be considered in any part of the Internet Standards Process, and there must be no assumption of any confidentiality obligation with respect to any such contribution." *Id.*, Section 10.2.

²¹ See Cremer, Rey, and Tirole (1998, 2000).

²² See Economides (1996a, b), Farrell and Saloner (1985), Katz and Shapiro (1985), and Liebowitz and Margolis (1994, 2002).

more compatible PCs sold because, then, there will be more software written and sold for such computers.

In networks of interconnected networks, there are large social benefits from the interconnection of the networks and the use of common standards. A number of networks of various ownership structures have harnessed the power of network externalities by using common standards. Examples of interconnected networks of diverse ownership that use common standards include the telecommunications network, the network of fax machines, and the Internet. Despite the different ownership structures in these three networks, the adoption of common standards has allowed each one of them to reap huge network-wide externalities.

For example, users of the global telecommunications network reap the network externalities benefits, despite its fragmented industry structure. If telecommunications networks were not interconnected, consumers in each network would only be able to communicate with others on the same network. Thus, there are strong incentives for every network to interconnect with all other networks so that consumers enjoy the full extent of the network externalities of the wider network.

The Internet has very significant network externalities. As the variety and extent of the Internet's offerings expand, and as more customers and more sites join the Internet, the value of a connection to the Internet rises. Because of the high network externalities of the Internet, consumers on the Internet demand universal connectivity (i.e., to be able to connect with every Website) on the Internet and to be able to send electronic mail to anyone. This implies that every network must connect with the rest of the Internet in order to be a part of it.

The demand for universal connectivity on the Internet is stronger than the demand of a voice telecommunications customer to reach all customers everywhere in the world. In the case of voice, it may be possible but very unlikely that a customer may buy service from a long-distance company that does not include some remote country because the customer believes that it is very unlikely that he/she would be making calls to that country. On the Internet however, one does not know where content is located. If company A did not allow its customers to reach region B or customers of a different company C, customers of A would never be able to know or anticipate what content they would be missing. Thus, consumers' desire for Internet universal connectivity is stronger than in voice telecommunications. Additionally, because connectivity on the Internet is two-way, a customer of company A would be losing exposure of his/her content (and the ability to send and receive e-mails) to region B and customers of company C. It will be difficult for customer A to calculate the extent of the losses accrued to him/her from such actions of company A. Thus, again, customers on the Internet require universal connectivity²³.

The existence of common interconnection standards and protocols in the telecommunications and the network of fax machines have guaranteed that no service

provider or user can utilize the existence of network externalities to create and use monopoly power. Similarly, the existence of common and public interconnection standards on the Internet guarantees that no service provider or user can utilize the existence of network externalities to create and use monopoly power based on proprietary standards. With competitive organization of the Internet's networks, the rising value is shared between content providers and telecommunications services providers (in the form of profits) and end users (in the form of consumer surplus).

4.1. Procompetitive consequences of network externalities

The presence of network externalities does not generally imply the existence of monopoly power. Where there are network externalities, adding connections to other networks and users adds value to a network, so firms have strong incentives to interconnect fully and to maintain interoperability with other networks. Thus, network externalities can act as a strong force to promote competition for services based on interconnected networks²³. For example, various manufacturers compete in producing and selling fax machines that conform to the same technical standards, and are connected to the ever-expanding fax network. It would be unthinkable that a manufacturer, however large its market share, would decide to produce fax machines for a different fax network that would be incompatible with the present one. In contrast, firms would like to conform to existing standards and fully interconnect to a network so that they reap the very large network externalities of the network.

The incentive to interconnect and to conform to the same standard applies similarly to competitive firms as it applies to firms with market power. Although, as in other markets, firms involved in network businesses may sometimes have

²³ If universal connectivity were not offered by a backbone network, a customer or its ISP would have to connect with more than one backbone. This would be similar to the period 1895–1930 when a number of telephone companies run disconnected networks. Eventually most of the independent networks were bought by AT&T, which had a dominant long-distance network. The refusal of AT&T to deal and interconnect with independents was effective because of three key reasons: (i) AT&T controlled the standards and protocols under which its network ran, (ii) long-distance service was provided exclusively by AT&T in most of the United States, and (iii) the cost to a customer of connecting to both AT&T and an independent was high. None of these reasons apply to the Internet. The Internet is based on public protocols. No Internet backbone has exclusive network coverage of a large portion of the United States. Finally, connecting to more than one backbone (multihoming) is a common practice by many ISPs and does not require big costs. And ISPs can interconnect with each other through secondary peering as explained later. Thus, the economic factors that allowed AT&T to blackmail independents into submission in the first three decades of the 20th century are reversed in today's Internet, and therefore would not support a profitable refusal to interconnect by any backbone.

²⁴ See also Faulhaber (2004).

market power, that power does not arise automatically from the network, even in the presence of externalities.

4.2. Conditions under which network externalities may inhibit competition

In markets with network externalities, firms may create bottleneck power by using proprietary standards. A firm controlling a standard needed by new entrants to interconnect their networks with the network of the incumbent may be in a position to exercise market power²⁵. Often a new technology will enter the market with competing incompatible standards. Competition among standards may have the snowball characteristic attributed to network externalities.

For example, VHS and Beta, two incompatible proprietary standards for video cassette recorders (VCRs), battled for market share in the early 1980s. Because Sony, the sponsor of the Beta standard, chose a pricing and licensing strategy that did not trigger the snowball effect; VHS was the winner. In particular, Sony refused to license its Beta standard, while VHS was widely licensed. Even though VHS was the winning standard, the market for VCRs did not become a monopoly since there are a number of suppliers of VHS-type video equipment. Thus, a standard may be licensed freely or at a low cost, and therefore the existence of a proprietary standard does not preclude competition. Moreover, in many cases a sufficiently open licensing policy will help to win the standards battle, and may therefore be in the interest of the owner of the standard to freely license even its proprietary standards²⁶.

Economics literature has established that using network externalities to affect market structure by creating a bottleneck requires three conditions²⁷:

1. Networks use proprietary standards;
2. No customer needs to reach nodes of or to buy services from more than one proprietary network;
3. Customers are captives of the network to which they subscribe and cannot change providers easily and cheaply.

First, without proprietary standards, a firm does not have the opportunity to create the bottleneck. Second, if proprietary standards are possible, the development of proprietary standards by one network isolates its competitors from network benefits, which then accrue only to one network. The value of each proprietary network is diminished when customers need to buy services from more than one network. Third, the more consumers are captive and cannot easily and economically change providers, the more valuable is the installed base to any

²⁵ See Economides (2003).

²⁶ See Economides (1996b).

²⁷ See Economides (1996a, 1989), Farrell and Saloner (1985), and Katz and Shapiro (1985).

proprietary network. The example of snowballing network effects just mentioned—VHS against Beta—fulfills these three conditions. The next section shows that these conditions fail in the context of the Internet backbone.

5. Network externalities and competition on the Internet

5.1. Conditions necessary for the creation of bottlenecks fail on the Internet

The Internet fails to fulfill any of the three necessary conditions under which a network may be able to leverage network externalities and create a bottleneck. First, there are no proprietary standards on the Internet, so the first condition fails. The scenario of standards wars is not at all applicable to Internet transport, where full compatibility, interconnection, and interoperability prevail. For Internet transport, there are no proprietary standards. There is no control of any technical standard by service providers and none is in prospect. Internet transport standards are firmly public property²⁸. As a result, any seller can create a network complying with the Internet standards—thereby expanding the network of interconnected networks—and compete in the market.

In fact, the existence and expansion of the Internet and the relative decline of proprietary networks and services, such as CompuServe, can be attributed to the conditions of interoperability and the tremendous network externalities of the Internet. America On Line (AOL), CompuServe, Prodigy, MCI, and AT&T folded their proprietary electronic mail and other services into the Internet. Microsoft, thought to be the master of exploiting network externalities, made the error of developing and marketing the proprietary Microsoft Network (MSN). After that product failed to sell, Microsoft relaunched MSN as an ISP, adhering fully to the public Internet standard. This is telling evidence of the power of the Internet standard and demonstrates the low likelihood that any firm can take control of the Internet by imposing its own proprietary standard.

Second, customers on the Internet demand *universal connectivity*, so the second condition fails. Users of the Internet do not know in advance what Internet site they may want to contact, or to whom they might want to send e-mail. Thus, Internet users demand from their ISPs and expect to receive universal connectivity. This is the same expectation that users of telephones, mail, and fax machines have: that they can connect to any other user of the network without concern about compatibility, location, or, in the case of telephone or fax, any concern about the manufacturer of the appliance, the type of connection (wireline or wireless), or the owners of the networks over which the connection is made. Because of the users' demand for universal connectivity, ISPs providing services

²⁸ See Kahn and Cerf (1999) and Bradner, *The Internet Standards Process*, revision 3, Network Working Group (<http://ftp.isi.edu/in-notes/rfc2026.txt>), Section 1.2.

to end users or to Web sites must make arrangements with other networks so that they can exchange traffic with *any* Internet customer.

Third, there are no captive customers on the Internet, so the third condition fails, for a number of reasons:

1. ISPs can easily and with low cost migrate all, or part of their transport traffic to other network providers;
2. Many ISPs already purchase transport from more than one backbone to guard against network failures, and for competitive reasons (ISP 'multihoming');
3. Many large Web sites/providers use more than one ISP for their sites ('customer multihoming');
4. Competitive pressure from their customers makes ISPs agile and likely to respond quickly to changes in conditions in the backbone market.

5.2. Bottlenecks such as the ones of the local exchange telecommunications network do not exist on the Internet

There are significant differences between local telephone networks and the Internet, which result in the existence of bottlenecks in local telephone markets and lack of bottlenecks on the Internet. Until the passage of the Telecommunications Act of 1996, the local telephone company had a legal franchise monopoly over local telephony in its territory in most States. Most importantly, the local telephone company monopolizes the fixed wireline connection to customers, especially the residential ones, thereby controlling the bottleneck for access to customers. Such a bottleneck does not exist on the Internet backbone. A number of reasons contribute to this:

1. the cost of connecting an ISP to the rest of the Internet is very low compared to the cost of connecting every house to local telephone service;
2. the location of an ISP is not predetermined, but can be placed most conveniently within a geographic area;
3. the elasticity of supply for Internet transport services is high (i.e., there are no barriers to expansion);
4. there are negligible barriers to entry on the Internet; and
5. Internet demand growth and expansion are exponential, driven by expanding market and geographic penetration, and by the introduction of new applications²⁹.

The only bottleneck in the Internet arises out of the control of the first/last mile of the local telecommunications network, by incumbent local exchange

²⁹ Demand grows yearly at about 100 percent. The number of North American ISPs more than tripled in 3 years. Up to 2003, demand growth has been overestimated as 400 percent per year both by the U.S. government and most providers of backbone connectivity, including MCI-WorldCom.

carriers, since this first/last mile is used by the majority of users to connect to the Internet.

In summary, an analysis of network externalities shows that network effects cannot create barriers to entry for new networks on the Internet or barriers to expansion of existing ones. It is also showed that network effects on the Internet do not create a tendency to dominate the market or tip it toward monopoly. On the contrary, network effects are a *procompetitive force* on the Internet, providing strong incentives to incumbents to interconnect with new entrants. In the next sections is discussed in detail the competition on the Internet.

6. Strategies that a large IBP might pursue

There are two main ways in which a large Internet backbone connectivity provider could attempt to exercise market power and harm consumers:

1. **Price increases.** It could raise the price of network services across-the-board to all customers, including replacing peering with transit sold at a high price; alternatively, it could selectively increase price to one or few networks;
2. **Raising rivals' costs or degrading interconnection without changing price(s).** It could selectively degrade the quality of interconnections with competing networks, in an effort to make their networks less attractive and divert traffic to itself.

As is explained in the next section, neither of these courses of action is likely to be profitable on the Internet backbone.

6.1. Raising the price of transport

The simplest exercise of market power by a large firm would be to raise the price of its transport services. In addition, the company might refuse to continue peering with some networks and to charge them transit fees instead. The ability of a company to profitably depeer other networks is equivalent to the ability of a company to increase the price of transport. Depeering does not mean cutting off a customer from the network or charging an infinite price to the customer; it does not mean refusal to deal. A price increase would create profit opportunities for the large internet backbone provider's (IBP's) rivals in the transport market, and is also likely to induce entry.

Internet backbone providers sell transport as a bandwidth of a certain capacity that allows an ISP to connect to the whole Internet. If a large Internet backbone connectivity provider were to increase the prices it charges to ISPs for such capacity, ISPs would promptly switch to other backbone providers. Thus, an

increase in transit price by a large IBP would decrease its sales sufficiently to make such a price increase unprofitable.

ISP connections to multiple backbones are very common. Forty-three percent of all ISP connections to backbones were sold as *additional* connections to ISPs who connected to more than one backbone. A multihoming ISP can easily and at a low cost limit the size of its purchases from an IBP that increases the price of transport. Thus, the presence of multihoming increases the firm-specific elasticity of demand of IBP transport services and creates a bigger demand response to IBP price increases. This makes it even more likely that the firm-specific demand response to a price increase will be sufficiently negative to render a contemplated price increase unprofitable.

If the large Internet backbone connectivity provider's strategy were to impose equal increases in transport costs on all customers, the response of other backbone providers and ISPs will be to reduce the traffic, for which they buy transit from the large IBP and to instead reroute traffic and purchase more transit from each other. Thus, in response to a price increase by the large Internet backbone connectivity provider, other IBPs and ISPs reduce the traffic for which they buy transit from the large IBP down to the minimum level necessary to reach ISPs that are *exclusively* connected to the large IBP. All other IBPs and ISPs exchange all other traffic with each other bypassing the large IBP network.

Figures 3 and 4 show the typical reaction of an increase in the price of a large IBP, and illustrate why the strategy of increasing price is unprofitable. Consider, for example, a situation where, prior to the price increase, 4 ISPs (1 to 4) purchase transit from IBP 0, which considers increasing its price. Two of these ISPs (ISP 2 and ISP 3) peer with each other. This is illustrated in Figure 3. Internet Service Provider 1 and ISP 4 buy transit capacity for all their traffic to IBP 0, and the other 3 ISPs, whereas, ISP 2 and ISP 3 buy transit capacity for all their traffic to ISP 0, ISP 1, and ISP 4.

Now suppose that, IBP 0 increases its transit price. In response, ISP 1 and ISP 4 decide to reduce the traffic for which they buy transit from IBP 0, and instead to reroute some of their traffic and purchase more transit from ISP 2 and ISP 3, respectively. See Figure 3. Because of the peering relationship between ISP 2 and ISP 3, all traffic from ISP 1 handed to ISP 2 will reach ISP 3 as well as ISP 4, who is a customer of ISP 3. Similarly, by purchasing transit from ISP 3, ISP 4 can reach all the customers of ISP 1, ISP 2, and ISP 3. Thus, in response to the price increase of IBP 0, each of the ISPs 1, 2, 3, and 4 will reduce the amount of transit purchased from the IBP 0. Specifically, each of the ISPs buys from IBP 0 only capacity sufficient to handle traffic to the customers of network 0. This may lead to a considerable loss in revenues for IBP 0, rendering the price increase unprofitable. The big beneficiaries of the price increase of IBP 0 are peering ISPs 2 and 3, who now start selling transit to ISPs 1 and 4, respectively, and become larger networks.

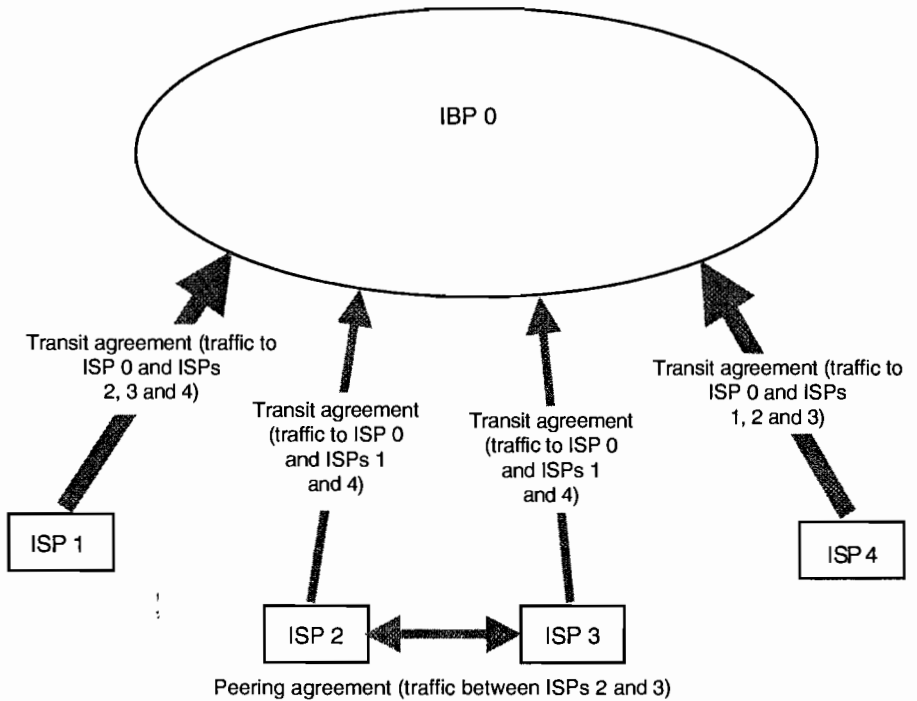


Fig. 3.

In response to a price increase by the large IBP, rivals would be able to offer their customers universal connectivity at profitable prices below the large IBP's prices. In the scenario described in the example above, market forces, responding to a price increase by a large network, reroute network traffic so that it is served by rival networks, except for the traffic to and from the ISPs connected exclusively with the large network. The rivals purchase the remaining share from the large IBP in order to provide universal connectivity. Thus, the rivals' blended cost would permit them to profitably offer all transport at prices lower than the large IBP's prices, but above cost.

A direct effect of the increase in price by the large network is that: (i) ISPs who were originally exclusive customers of the large IBP would shift a substantial portion of their transit business to competitors and (ii) ISPs that were not exclusive customers of the large IBP would also shift a significant share of their transit business to competitors' networks, keeping the connection with the large

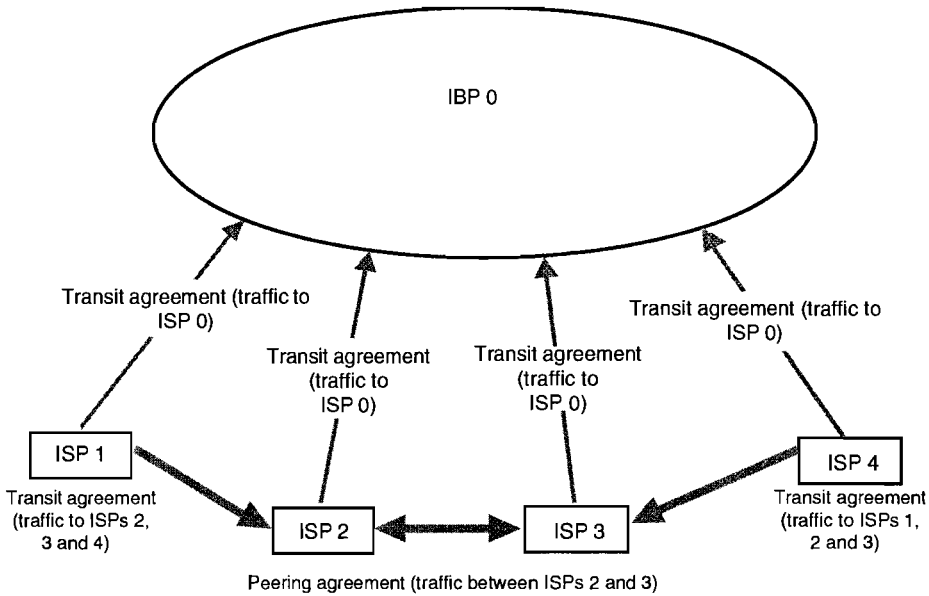


Fig. 4.

IBP only for traffic for which alternate routes do not exist, or for cases of temporary failure of the rivals' networks.

6.2. *Discriminatory price increases directed simultaneously against all backbone rivals*

Here I consider the possibility that a large IBP might try to displace its rivals by charging them more than it charges ISPs who are not rivals in the transport business. It is believed that this form of price discrimination is particularly unlikely. Of all customers, rivals in the transport business—major backbones and smaller regional networks—are the best positioned to avoid the use of the large IBP's network if it is more expensive than the alternatives. Even the smaller rivals are large enough that the transactions costs of establishing alternative connections are unimportant in relation to the cost increases for transport that could be avoided by making new deals.

6.3. Raising rivals' costs and degrading connectivity

Alternatively one may consider the possibility that the large IBP would find it profitable to raise the nonprice costs of rivals by reducing the connectivity it provides with other IBPs³⁰. 'The first observation regarding the 'raising rivals' cost' or 'degradation' strategy applied to clients is that as a matter of economics, it is *always* preferable to a firm to increase price rather than increase the nonprice costs of rivals. A firm can choose a price increase that will have the same effect as increasing the costs (or reducing the benefits) of its clients, and it is able to collect extra revenue through the price increase while, if it just degrades the product it receives no extra revenue. Another difference from the traditional raising rivals cost theory is that on the Internet backbone often imposing a quality decrease on a rival simultaneously results in a quality decrease on the perpetrator. This is because the quality degradation affects both the services demanded by the target network (its clients connecting to the perpetrators network) as well as the services demanded by perpetrators' network (its clients connecting to the target networks). As discussed in detail in the next section, because of the network feature of connectivity, such a degradation cannot be confined to the target, but it also simultaneously affects negatively the perpetrator.

6.3.1. Terminating interconnection simultaneously with all rivals (refusal to deal)

The author first considers the extreme case in which a large IBP terminates interconnection with all rivals. This setup is equivalent to the large IBP increasing the rivals' price to infinity if they were to interconnect with it.

Termination by the large IBP of interconnection with a network customer has a bilateral effect. It prevents the other network's customers from reaching any customer of the large IBP, and it prevents the large IBP's customers from reaching any customers of the other network. Whatever the relative sizes of the two networks, customers of both networks are harmed. If the large IBP's network has more customers than the interconnecting network, then the termination strategy will affect the large IBP's network as much, or more than the interconnecting network.

Termination of interconnection would deny the large IBP's customers the universal connectivity sought by every customer, and would have devastating effects for the large IBP. Its customers—larger Web sites and the ISPs specializing in end-user services and Web hosting—would seek new transport providers to make up for the large IBP's inability to deliver universal connectivity. The loss of business is likely to make termination of interconnection highly unprofitable.

³⁰ Hausman (2000) at paragraph 53.

This is a good demonstration of the procompetitive effects of network externalities in the Internet. Each network, including a large network, has a more valuable product if it interconnects with other networks. Termination of interconnection would severely lower the value of the large network's service because it would shrink the connectivity the company offered.

6.3.2. *Degrading interconnection simultaneously with all rivals*

Alternatively, it has been suggested that a large IBP would degrade interconnection with all rivals without terminating service³¹. However, a large IBP could always make more profit by charging more for interconnection than by offering poor service. There is always a price level that has the equivalent harmful effect on customers as a program of degradation. The higher charge puts money in the seller's pocket immediately; degradation does not. Because, as the author has concluded, a large IBP would not find it profitable to raise transport charges; it follows immediately that it would suffer even more from degrading service.

In a similar vein to the earlier discussion, even if a large IBP decided to degrade interconnections rather than raise price, degrading interconnections would impose a cost on it that is comparable to the cost imposed on the rivals. In total, the large IBP's customers would experience the same level of degradation in terms of the traffic sent to, or received from, the other networks as would the other networks' customers.

Some have argued that the effects of degraded interconnections would be less severe for a large IBP than for the other networks because of the large IBP's size. In this line of argument, if the traffic is isotropic³², a large number of Internet interactions will be within the network of the large IBP, and these interactions will be unaffected by degradation of interconnection. According to this theory, the rest of the Internet networks (with the smaller total number of customers if the large IBP has more than 50 percent of Internet customers) will suffer more than the larger network; it follows that the large IBP can then attract the customers of other networks³³.

This argument is based on the assumption that Internet users do *not* require universal connectivity. This, however, is factually incorrect. Internet users demand to be able to reach every node of the Internet, in a similar way that telecommunications customers demand that they be able to reach anyone connected to the telecommunications network, no matter where the receiving party is located, which local exchange carrier he/she subscribes to, and who carries the long-distance call.

³¹ Hausman (2000) at paragraph 53.

³² Isotropic traffic is generated when every user initiates the same number and type of Internet interactions with every other user.

³³ Hausman (2000) at paragraph 53.

Since users demand universal connectivity on the Internet, no network, however large, can afford not to offer universal connectivity. Therefore, no network would decide to degrade connections with the rest of the Internet networks unless the degrading network was certain that *all* ISPs *not* connected to it would immediately react to the degradation by instantaneously switching to the degrading network. This instantaneous switching is extremely unlikely to happen. Instead, many ISPs would reduce rather than increase use of a network that is degrading the quality of interconnections for a significant amount of Internet traffic. And, as long as there are ISPs who have not switched to the degrading network, all customers of the degrading network suffer. Each one of these customers of the degrading network is receiving connectivity significantly below his expectations of universal connectivity, and is now willing to pay less for it. Thus, the loss in value from degradation is comparable on both sides of the degraded interconnections, and can in fact be higher for the larger network. This means that a large network can only harm its rivals by harming itself by just as much or more.

Degradation of interconnections, like termination of interconnections, sacrifices the benefits of network externalities. It would result in a loss of value in the large IBP's Internet businesses because it would limit its customers' ability to interact with the rest of the Internet. A rational business would not take this step. Because there are limited switching costs and negligible barriers to expansion and entry, transport customers would switch to other networks or new entrants rather than tolerate a degraded interconnection and alienate their customers. Networks monitor the quality of service aggressively on behalf of their end users and Website customers, and they are able to identify and react to problems that would result from deliberate degradation of interconnection.

6.3.3. *Sequential attacks on rivals*

Some authors have claimed that although a raising-rivals'-costs strategy is unprofitable against all rivals; it would be profitable if applied sequentially to one rival at a time³⁴. In this line of thought, a large IBP would degrade interconnections by targeting rivals and ISP customers one after the other. Cremer et al. (2000) raise a number of anticompetitive concerns for networks that obey the following assumptions:

1. consumers do *not* demand universal connectivity;
2. there is an installed base of clients (ISPs) of Internet backbone networks who cannot migrate to other providers.

Under these assumptions, Cremer et al. (2000) argue: (a) a large IBP network has an incentive to introduce incompatibilities and to degrade interconnection with one rival, but not with all rivals, (b) even small differences in network size will

³⁴ See, Hausman (2000) at p. 54 and Cremer et al. (2000).

Table 10
 Contrast between the assumptions of Cremer et al. (2000) and Internet facts

Issue	Assumptions of Cremer et al. (2000)	Internet facts
Consumers' preferences for connectivity	Consumers do not demand universal connectivity	Consumers demand universal connectivity ^a
Consumers' willingness to switch Internet provider	No switching by existing customers	Easy customer migration
Effects of congestion on network performance	Interconnection is totally degraded when capacity is slightly exceeded	Networks have spare capacity; in situations of congestion, quality falls proportionally with congestion

^aFor example, the European Union Commission, in its Statement of Objections to the MCI-WorldCom merger recognized the lack of validity of Cremer et al.'s (2000) first assumption in stating that 'the demand for Internet connectivity continues to be universal in scope' (at p. 81).

lead to a spiral of ever-increasing dominance by a larger IBP network, since dominance is defined by size, (c) large IBP networks will refuse to cooperate with small networks, and (d) in the case, where switching costs are low, large IBP networks will still be able to dominate small networks³⁵.

The Internet violates the assumptions of Cremer et al. (2000), as is described later. And, since the fundamental assumptions of Cremer et al. (2000) diverge in fundamental ways from the reality of the Internet, the conclusions of Cremer et al. (2000) do not apply to competition on the Internet. These differences are summarized in Table 10.

The claim of Cremer et al. (2000) that a large IBP network will degrade interconnection with a targeted rival, is empirically invalid as explained in next section. The fact that such behavior has not occurred on the Internet backbone despite significant differences in market shares among the various backbone providers should be sufficient proof that Cremer et al. (2000) are discussing a different network from the Internet. Moreover, on the Internet we have observed a trend in the opposite direction (i.e., toward interconnection and full compatibility). Proprietary networks that preceded the commercial Internet, some dominant in their realm, such as AOL, CompuServe, Prodigy, MCI Mail, AT&T Mail, and MSN chose not to remain incompatible networks, but instead accepted full compatibility as parts of the Internet.

The results of Cremer et al. (2000) are indeed extremely sensitive to variations in the assumptions made. The assumption that ISPs are not allowed to migrate to

³⁵ Cremer et al. do not identify the structural conditions under which they expect the anticompetitive behavior described in the earlier paragraph would occur.

other backbones, which is presented by Cremer et al. (2000) as 'conservative,' is not only unsubstantiated but also critical to support its claim of dominance of a large IBP network and of degradation of interconnection. When consumer migration is allowed within the framework of the Cremer et al. (2000), there is no dominance or 'snowballing.' It is shown in the Appendix that, with exactly the same assumptions of Cremer et al. (2000), except now allowing customer migration, the market equilibrium shows no network dominance by any firm, and no network has an incentive to degrade interconnection.

Cremer et al. also state that multihoming will not diminish the incentives or ability of a dominant firm to engage in serial degradation. They base this on an unrealistic network model setup. In Cremer et al.'s (2000) targeted degradation model, a large IBP network (network 1) cuts interconnection with network 3, while the only other remaining IBP network (network 2) interconnects with both networks 1 and 3, but is prevented (by assumption) from offering transit to the targeted network 3³⁶. Cremer et al. (2000) allow for multihoming only between networks 1 and 3; they do not allow multihoming across other networks. Thus, multihoming *a-la*-Cremer et al. (2000) shields some customers of network 3 from the effects of targeted degradation but has no other effect. In the reality of the Internet, multihoming is available to customers of all networks, and large percentages of customers of all networks utilize it. If the interconnection between networks 1 and 3 were severed: (i) customers of network 1 that multihome with network 2 would shift their traffic to network 2 to gain access to network 3, thus reducing the capacity they would buy from network 1, causing the targeted degradation to be even less profitable for network 1, (ii) customers that multihome with all three networks would also shift their traffic to network 2, since network 2 is the only one that provides universal connectivity, and (iii) customers of network 3 that multihome with network 2 would increase the capacity of transit they buy from network 2, so that they are able to gain access to network 1. Thus, in the real Internet, the existence of multihoming: (i) makes targeted degradation even less profitable for the targeting network since it results in a steeper demand response and (ii) makes the nondegraded network(s) stronger competitors of the targeting network. In conclusion, the presence of multihoming makes it even less likely that targeted degradation will ever occur. Table 11 summarizes the differences in the results when customer migration is allowed.

A key conclusion of Cremer et al. (2000) is that the largest network will use targeted degradation of rival networks. But, targeted degradation is unprofitable for a large network that would initiate it because:

1. ISP clients of the targeted network are likely to switch to third IBP networks that are unaffected by the degradation; it is very unlikely that any will switch to

³⁶ The assumption of Cremer et al. that network 2 (or other third networks in a more general setting), will not sell transit to the targeted network is totally unreasonable.

- the degrading IBP network because it is itself degraded, and cannot offer universal connectivity; there is no demand reward to the large IBP network;
2. Degradation of interconnection hurts all the ISP customers of the targeting IBP network as well, since they lose universal connectivity; these customers of the large network would now be willing to pay less to the large network; this leads to significant revenue and profit loss;
 3. After losing universal connectivity, customers of the large IBP network are likely to switch to other networks that are unaffected by degradation and can provide universal connectivity; this leads to even further revenue and profit loss for the degrading network;
 4. Multihoming ISPs would purchase less capacity from the large IBP network, or even terminate their relationship with the large network, which, through its own actions sabotages their demand for universal connectivity; this further reduces demand and profits for the degrading network; the same argument applies to multihoming customers of ISPs;
 5. As the large IBP network pursues target after target, its customers face continuous quality degradation while the target's customers face only temporary degradation; this would result in further customer and profit losses for the large IBP network;
 6. Prospective victims would seek alternative suppliers in advance of being targeted by the large IBP network; the scheme cannot play out the way it is proposed;
 7. The degradation scheme is implausible in its implementation. How large do networks need to be to become serial killers? Why have we not observed this behavior at all?
 8. There is no enduring change to the number of competitors in a market caused by serial degradation in a market with negligible entry barriers; the eliminated rival is likely to be replaced by another.

Table 11
Contrast between the results of Cremer et al. (2000) and results when customer migration is allowed

Issue	Claims by Cremer et al. (2000)	Results when customer migration is allowed
Strategic power	Dominance by 'large' network	Equal bargaining power among networks
Dynamic effects	'Snowballing' or 'tipping' leading to monopoly	Equilibrium at equal market shares; no 'snowballing' or 'tipping'
Willingness of providers to interconnect	Even a slightly larger network will refuse to interconnect with other networks	Network externalities and demand for universal connectivity force networks to interconnect

The reasons why the strategy of targeted degradation would be self-defeating: first, degrading interconnections with networks that have an alternative way to send and receive traffic through a second network connection with another network would lead to a quick response by the rivals of routing almost all of their traffic through the second network, and would therefore be undesirable to a large network. Figures 3 and 4 above illustrated the rerouting of traffic in response to a price increase by a large IBP. The response of competitors and clients of an IBP that degraded interconnection would be very similar to the responses of rivals and clients to a price increase by the large IBP as shown in Figures 3 and 4. Moreover, a target network is likely to enter into new peering and transit arrangements with other networks that would further divert traffic from the degrading IBP. The target network could buy transit from other networks, whose connectivity with the large IBP's network is intact, and avoid all degradation problems. Thus, in response to degradation, traffic is routed away from the degrading IBP so that the culprit loses customers, traffic, and profits.

Second, as explained earlier, inequality in size does not imply inequality in the value of the damage sustained by two interconnecting networks as a result of a degraded interconnection. Suppose that the large IBP degraded its interconnection with a much smaller network. If traffic were spread evenly across all customers (end users and Websites), the reduction in service quality experienced by each of the large IBP's customers may be smaller than the reduction in service quality experienced by each of the smaller rival's customers. Some argue that this implies that ISPs connected to the targeted rival would then switch to the large IBP, and therefore the degradation strategy is 'successful' in attracting customers to the large IBP. This argument is based on the assumption that Internet users do *not* require universal connectivity, an assumption that is factually incorrect. Since Internet users demand universal connectivity, no network would decide to degrade a target network unless the degrading network was certain that *all* ISPs of the target network would immediately react to the degradation by instantaneously switching away from the target network. This instantaneous switching is extremely unlikely to happen. The target network is likely to establish new peering and transit relationships with other networks and utilize its multihoming arrangements to divert traffic away from the degraded interconnection and minimize the effect on its customers. After all, since the target network is the only one with degraded connectivity to the large IBP's network, the target network can easily buy transit service from other networks, which have full connectivity to the large IBP's network and avoid all degradation problems. And, as long as there are ISPs of the target network who have not switched to the degrading network, the users of the ISPs connected to the large IBP will suffer significantly as a result of the degradation. If the large IBP were to degrade its interconnection to a target network, the customers of the large IBP will be willing to pay less for the degraded service, and the large IBP would lose profits, even if the degradation strategy were 'successful' in attracting customers to it. After all, a larger number of customers of

the large IBP would experience a reduced service quality than the potential number of customers that the large IBP could attract from the small target ISP³⁷. Thus, the commercial impact of the serial degradation on the large IBP in terms of profit loss would be significant.

Third, the large IBP's customers are anything but captives. Business and individual end users and Website operators are sensitive to the quality of the service they receive. The large IBP could not use its customer base as a tool for harming rivals because it would lose the customer base in the process. Customers would switch to another network in response to a reduction in service quality. A degraded interconnection reduces the quality of the service that the large IBP's customers receive, and if they could not get reliable and quick access to popular Websites served by the network rival whose connection was degraded, these customers would move to other networks whose connection with the victimized network was unimpaired. Therefore, picking rivals one by one would not reduce the damage of this strategy to the large IBP.

Fourth, as already discussed, a significant number of end-user service providers have connections with more than one transport provider and most large content providers have connections with a number of networks. Even if the serial killer argument were correct for traffic that went to ISPs that were exclusively connected with the large IBP, and somehow the large IBP benefited from degradation of quality to these ISPs, the degradation of quality of the large IBP network would lead multiple connection ISPs to move traffic away from the large IBP and terminate their relationship with the large IBP.

Fifth, by targeting rivals sequentially (rather than all at once), the large IBP might limit the size of the damage to itself at any point in time, but it would be just as large in total. Moreover, over a period of time, the serial degradation strategy hurts more a customer of the large IBP than a customer of any targeted network.

If the serial degradation strategy is pursued, the large IBP's customers would experience constant problems in connecting to Websites not served by the company, while each victim would face only temporary quality degradation. For example, suppose that, over a period of a year, a large network sequentially degrades interconnections for 4 months for each of three smaller competitors. Then customers of the larger network will experience degradation over the course of all 12 months, but customers of each of the smaller networks will not experience degradation for 8 months of the year. The continuous quality degradation experienced by customers of the larger network is at least as great as that occasionally experienced by customers of smaller (target) networks.

³⁷ As explained earlier, even if the merged company is 'successful' in making customers leave the target network, it is likely that most of the customers leaving the target will not switch to the merged company because of the merged company's network also faces a quality degradation.

Sixth, the serial killer scenario assumes that the purchasers of Internet transport services have a passive response to the plan as it unfolds. After each victim falls, they switch their transport business to the predator, knowing perfectly well that the ultimate result will be higher prices for transport services. In fact, the rational response would be the opposite. As the plan developed, the prospective victims would take action to avoid becoming victims at all. They would seek alternative suppliers for the majority of their Internet connectivity, cutting back purchases from the large IBP to the bare minimum.

Seventh, the 'serial killer' scenario is totally implausible in its implementation. Its proponents have left a number of key questions unanswered. For example, for how long will the large IBP target a network before switching to its next victim? How does the large IBP hide from its customers the increasing degradation in its service to them? How large do networks need to be to find it desirable to be serial killers? Why have we not observed this behavior at all? How do the proponents of the serial killer theory explain why the degradation of connectivity would happen in the future but has never happened up to now?

Eighth, the serial degradation strategy would be impossible to execute in practice, because new networks are coming into existence all the time. By the time that the large IBP had degraded interconnection with one network, the number of alternatives will have multiplied. In a market with negligible barriers to entry, there is no gain to eliminating one set of rivals because they will be replaced by another.

Ninth, the role of customer mobility that helps in maintaining competition in Internet transport. Larger customers already have multiple connections to the Internet and all customers can switch suppliers easily. Many ISPs have multiple connections to IBPs. Advocates of the serial killer scenario have suggested that customer mobility may contribute to the potential success of the serial killer strategy, because the customers of the targeted IBP will abandon that IBP quickly and fully³⁸. This theory is incorrect because it disregards the incentives of multihoming customers and of other customers of the large IBP to switch their traffic away from the large IBP in response to the degradation.

A multihoming ISP who is a customer of the large IBP (which initiates the connectivity degradation of the small IBP in the serial killer scenario) will also observe the degradation. Such an ISP will have an incentive to switch most of its traffic away from the two affected IBPs (large and small) to a third network. The ISP that switches traffic to a third network will now buy less transit from the large IBP. This provides incentives for the large IBP not to engage in degradation. The existence of multihoming implies that ISPs can easily reduce the amount of transit they buy from the large IBP in response to even small degradation of

³⁸ See, *id* at p. 57.

quality. Thus, multihoming decreases the incentive for a large IBP to degrade connectivity.

In conclusion, serial degradation is no more likely than simultaneous degradation. It would lower, not raise, the large IBP's profits.

7. Conclusions

The commercial Internet is one of the most important innovations in telecommunications and computing of the last 50 years. This ubiquitous data network based on low-level public technical standards has displaced well-established sophisticated high-level networks and has grown to reach a very large percentage of computers worldwide. At the core of the ability of the Internet to provide transport services lie the Internet backbones. The Internet backbone market has quickly grown to extremely high capacity of transmission and has surpassed the transmission capacity of the traditional long-distance network. Despite ups and downs, including the dot com boom and bust, and the WorldCom accounting scandal and bankruptcy, the Internet backbone market has shown robust competition. The dire predictions of the European Union Competition Authority in 1998 and 2000, that the Internet would be dominated by a single firm that would impose its own standards and refuse to interconnect rival backbones, have failed to materialize.

Appendix

A. 1. Duopoly

Reexamine the duopoly model of Cremer et al. (2000), Section 4, keeping all the assumptions of the model, except one: Allow customers in the installed base of each network to migrate to the other network if price and quality considerations so warrant. Thus, in the modification, the size of the network (sales) and the installed base coincide. All the symbols are the same as in Cremer et al. (2000) except that now output of firm i is q_i rather than $q_i + \beta_i$

In particular, as in Cremer et al. (2000), assume two interconnected Internet backbone networks, $i = 1, 2$ with θ in $[0, 1]$ being the quality of interconnection between the networks and v signifying the importance of connectivity. Cremer et al. (2000) assume that backbone i has an installed base of captured customers β_i , who do not respond to prices and would not sign up with a backbone other than i at any price. Backbone i also has q_i customers, who respond to prices. Assuming that the quality of interconnection within a backbone is $\theta = 1$, Cremer et al. (2000) define the 'quality' of service of backbone i corresponding to its ability to reach customers as

$$s_i = v[(\beta_i + q_i) + \theta(\beta_j + q_j)] \quad (1)$$

As mentioned earlier, in the analysis in Cremer et al. (2000), customers β_i and β_j are not allowed to change providers. Cremer et al. (2000) show (Proposition 1) that under these conditions, for some parameter values, the larger of the two backbones chooses a lower interconnection quality (θ) than its rival, and that the quality of interconnection that the larger backbone chooses decreases in the difference between the captured customers of the larger and the smaller networks who are not allowed to change providers.

If, alternatively, all customers are allowed to buy service from a competing backbone, that is the number of captive customers is zero, $\beta_i = \beta_j = 0$, then Cremer et al.'s (2000) Equation (1) that defines the quality of good i becomes

$$s_i = v(q_i + \theta q_j).$$

Equation (2) remains as in Cremer et al. (2000). Equations (3) and (4) defining willingness to pay for each backbone become

$$q_i + q_j = 1 - (p_i - s_i) = 1 - (p_j - s_j),$$

$$p_i = 1 - (q_i + q_j) + s_i = 1 + v(q_i + \theta q_j) - (q_i + q_j), \quad i = 1, 2.$$

Profits of firm (backbone) i are

$$\Pi_i = (p_i - c)q_i = [1 + v(q_i + \theta q_j) - (q_i + q_j) - c]q_i.$$

Maximization with respect to q_i results in the best response of firm i to the sales of the opponent:

$$q_i = R_i(q_j) = \frac{1 - c - q_j(1 - v\theta)}{2(1 - v)}.$$

Cournot equilibrium sales (network sizes) are then

$$q_i^* = q_j^* = \frac{(1 - c)}{3 - v(2 + \theta)}.$$

Notice that, for all θ , that is whatever the degree of network interconnection quality, both networks have exactly the same size, $q_i^* = q_j^*$. Thus, when customers in the installed base are allowed to migrate across networks, contrary to the results of Cremer et al. (2000), *there is no network dominance at the market equilibrium*.

Equilibrium profits of the two networks are equal:

$$\Pi_i^* = (1 - v)(q_i^*)^2 = (1 - v)(q_j^*)^2 = \Pi_j^*.$$

It follows that both networks have the same incentive to increase quality:

$$\frac{d\Pi_i^*}{d\theta} = \frac{d\Pi_j^*}{d\theta} > 0.$$

Thus, contrary to Proposition 1 of Cremer et al. (2000), both networks have equal and positive incentives to maintain a high quality of interconnection between them.

A. 2. Merger analysis

Examine the merger analysis model of Cremer et al. Section 6, keeping all the assumptions of the model, except one: Allow customers in the installed base of each network to migrate to the other network if price and quality considerations so warrant. As in the duopoly model above, the size of the network (sales) and the installed base coincide so that all the symbols are the same as in Cremer et al. (2000) except that now output of firm i is q_i rather than $q_i + \beta_i$.

Cremer et al. (2000) start with four networks of equal sizes. In the original equilibrium, all networks have equal sizes and profits. After a merger between two of them, there are three networks in the market. In the ‘targeted degradation’ scenario of Cremer et al. (2000), network 1 severs its interconnection to network 3, while maintaining full interconnection to network 2. Networks 2 and 3 are fully interconnected, but network 3 is not allowed to use network 2 for transit to network 1. Cremer et al. (2000) show that, for some parameters, the merged firm will prefer to follow the ‘targeted degradation’ strategy (Proposition 6).

However, if one alternatively assumes that the installed base of each network is allowed to migrate to the other network if price and quality considerations so warrant, the ‘targeted degradation’ result of Cremer et al. (2000) is reversed. Specifically, calling q_i the sales of firm I , after a merger between two of the four networks, there are now three networks in the market. Assuming no degradation, their prices and profits are

$$p_i = 1 - (q_1 + q_2 + q_3) + v(q_1 + q_2 + q_3), \quad \Pi_i = (p_i - c)q_i, \quad i = 1, 2, 3,$$

Equilibrium quantities, prices, and profits without degradation are

$$q_i^* = \frac{(1 - c)}{4(1 - v)}, \quad p_i^* = \frac{(1 + 3c)}{4}, \quad \Pi_i^* = \frac{(1 - c)^2}{16(1 - v)}.$$

Now consider the ‘targeted degradation’ scenario of Cremer et al. as described above. In this scenario, network 1 severs its interconnection to network 3, while maintaining full interconnection to network 2. As in Cremer et al. (2000), although networks 2 and 3 are fully interconnected, it is assumed that network 3 is not allowed to use network 2 for transit to network 1. Then prices and profits are:

$$p_1 = 1 - (q_1 + q_2 + q_3) + v(q_1 + q_2), \quad p_2 = 1 - (q_1 + q_2 + q_3) + v(q_1 + q_2 + q_3),$$

$$p_3 = 1 - (q_1 + q_2 + q_3) + v(q_1 + q_2), \quad \Pi_i = (p_i - c)q_i.$$

Using superscript 'd' to denote degraded interconnection, equilibrium quantities, prices, and profits are

$$q_1^d = q_3^d = \frac{(1-c)}{2(2-v)}, \quad q_2^d = \frac{(1-c)}{2(2-v)(1-v)},$$

$$p_1^d = p_3^d = \frac{1+3c-v(1+c)}{2(2-v)}, \quad p_2^d = \frac{(1+3c-2vc)}{2(2-v)},$$

$$\Pi_1^d = \Pi_3^d = \frac{(1-c)^2(1-v)}{4(2-v)^2}, \quad \Pi_2^d = \frac{(1-c)^2}{4(1-v)(2-v)^2}.$$

Now compare profits of network 1 with and without targeted degradation of interconnection. It is easy to show that profits of network 1 without degradation are higher than profits of the same network with targeted degradation:

$$\Pi_1^* = \Pi_{1:}^d = \frac{(1-c)^2(4-3v)v}{16(2-v)^2(1-v)} > 0,$$

since $0 < v < 1$. Therefore, contrary to Cremer et al. (2000), if all consumers are allowed to change providers if prices and qualities so warrant, network 1 (the largest one) finds it profitable *not* to use 'targeted degradation.'

References

- Boardwatch Magazine, 1999, Directory of Internet Service Providers, 11th Edition.
 Boardwatch Magazine, 1997, Fall.
 Cremer, J., P. Rey and J. Tirole, 2000, Connectivity in the commercial internet, *Journal of Industrial Economics*, 48, 433-472.
 Cremer, J., P. Rey and J. Tirole, 1998, The Degradation of Quality and the Domination of the Internet, Appendix 5 in the Submission of GTE to the European Union for the merger of MCI with WorldCom.
 Besen, S. M. and S. R. Brenner, 2000, Declaration to the FCC, March 20.
 Economides, N., 1989, Desirability of compatibility in the absence of network externalities, *American Economic Review*, 79(5), 1165-1181.
 Economides, N., 1996a, The economics of networks, *International Journal of Industrial Organization*, 14(2), 675-699, lead article. Prepublication electronic copy available at <http://www.stern.nyu.edu/networks/94-24.pdf>.
 Economides, N., 1996b, Network externalities, complementarities, and invitations to enter, *European Journal of Political Economy*, 12, 211-232. Prepublication electronic copy available at http://www.stern.nyu.edu/networks/Networkexternalities_EJPE_1996.pdf.
 Economides, N., 1999a, The Telecommunications Act of 1996 and its impact, *Japan and the World Economy*, 11, 455-483, lead article. Prepublication electronic copy available at <http://www.stern.nyu.edu/networks/98-08.pdf>.

- Economides, N., 1999b, Real Options and the Costs of the Local Telecommunications Network, in J. Alleman and E. Noam, eds., *Real Options: The New Investment Theory and its Implications for Telecommunications Economics*, Regulatory Economics Series, Boston: Kluwer Academic Publishers. Prepublication electronic copy available at <http://www.stern.nyu.edu/networks/real.pdf>.
- European Union Commission, 1998, Statement of Objections to the MCI WorldCom Merger.
- Farrell, J. and G. Saloner, 1985, Standardization, compatibility, and innovation, *Rand Journal of Economics*, 16, 70–83.
- Faulhaber, G., 2004, *Bottlenecks and Bandwagons: Access Policy in the New Telecommunications*, Handbook of Telecommunications, Amsterdam: Elsevier Publishers.
- Hausman, J., 2000, Declaration to the FCC.
- Hogendorn, 2004, Excessive Entry of National Telecom Networks, 1990–2001, NET Institute Working Paper, available at <http://www.netinst.org/Hogendorn.pdf>.
- Hubbard, G. and W. Lehr, 2000, in I. Vogelsang and B. M. Compaine, eds., *Telecommunications, the Internet, and the Cost of Capital in The Internet Upheaval—Raising Questions*, in *Seeking Answers in Communications Policy*, Cambridge, MA and London: MIT Press.
- Kahn, R. E. and V. G. Cerf, 1999, What Is The Internet (And What Makes It Work), December 1999, at http://www.wcom.com/about_the_company/cerfs_up/internet_history/whatIs.phtml.
- Katz, M. and C. Shapiro, 1985, Network externalities, competition and compatibility, *American Economic Review*, 75(3), 424–440.
- Liebowitz, S. J. and S. E. Margolis, 1994, Network externality: An uncommon tragedy, *The Journal of Economic Perspectives*, 133–150.
- Milgrom, P., B. M. Mitchell and P. Srinagesh, 2000, Competitive Effects of Internet Peering Policies, in I. Vogelsang and B. M. Compaine, eds., *The Internet Upheaval—Raising Questions*, *Seeking Answers in Communications Policy*, Cambridge, MA, and London: MIT Press.
- RHK, 2005, *Internet Traffic Analysis—Methodology and Results*, May 2005.

Further reading

- Armstrong, M., 1998, Network interconnection in telecommunications, *Economic Journal*, 108, 545–564.
- Bradner, S., The Internet Standards Process, revision 3, Network Working Group at <ftp://ftp.isi.edu/in-notes/rfc2026.txt>.
- Carter, M. and J. Wright, 1999, Interconnection in network industries, *Review of Industrial Organization*, 14, 1–25.
- Competition Policy in Network Industries: An Introduction, in Dennis, Jansen, ed., *The New Economy and Beyond: Past Present and Future*, Edward Elgar (2006), pre-publication copy at http://www.stern.nyu.edu/networks/Competition_Policy.pdf.
- Economides, N., 1998, The incentive for nonprice discrimination by an input monopolist, *International Journal of Industrial Organization*, 16, 271–284. Prepublication electronic copy available at <http://www.stern.nyu.edu/networks/1136.pdf>.
- Economides, N., 2000, Declaration to the FCC, May 21.
- Economides, N., G. Lopomo and G. Woroch, 1996, Regulatory pricing policies to neutralize network dominance, *Industrial and Corporate Change*, 5(4), 1013–1028, Prepublication electronic copy available at <http://www.stern.nyu.edu/networks/96-14.pdf>.
- Kende, M., 2005a, Declaration of Michael Kende to the FCC in the Verizon-SBC merger.
- Kende, M., 2005b, Reply Declaration of Michael Kende to the FCC in the Verizon-SBC merger.
- Laffont, J.-J., P. Rey and J. Tirole, 1998a, Network competition: I. Overview and nondiscriminatory pricing, *RAND Journal of Economics*, 29, 1–37.
- Laffont, J.-J., P. Rey and J. Tirole, 1998b, Network competition: II. Price discrimination, *RAND Journal of Economics*, 29, 38–56.

- Laffont, J.-J., S. Marcus, P. Rey and J. Tirole, 2001, Internet interconnection, and the off-net cost pricing principle, *American Economic Review*, May 2001.
- Laffont, J.-J. and J. Tirole, 2000, *Competition in Telecommunications*.
- Liebowitz, S. J. and S. E. Margolis, Network Effects, in M. Cave, S. Majumdar, and I. Vogelsang, eds., *Handbook of Telecommunications Economics*, pp. 75–96, Amsterdam: Elsevier Publishers.
- Noam, E. M., 2001, *Interconnecting the Network of Networks*, Cambridge, MA: MIT Press.
- Odlyzko, A. M., 2003, Internet traffic growth: Sources and implications, at <http://www.dtc.umn.edu/~odlyzko/doc/itcom.internet.growth.pdf>.
- Schwartz, M., 2005a, Declaration of Marius Schwartz to the FCC in the SBC-AT&T merger.
- Schwartz, M., 2005b, Reply Declaration of Marius Schwartz to the FCC in the SBC-AT&T merger.