

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.

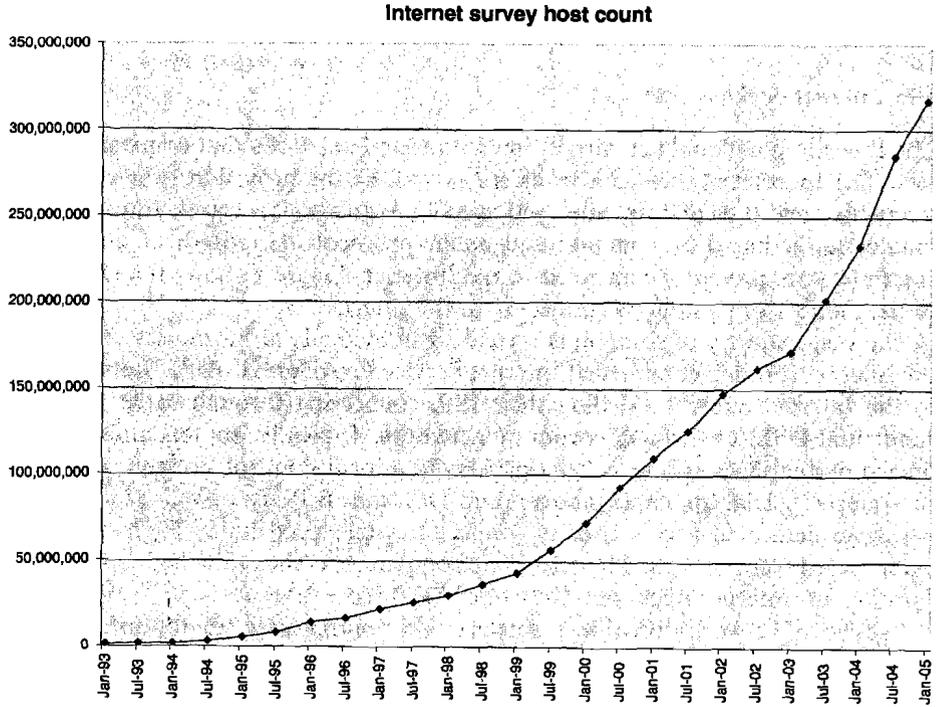


Fig. 1.

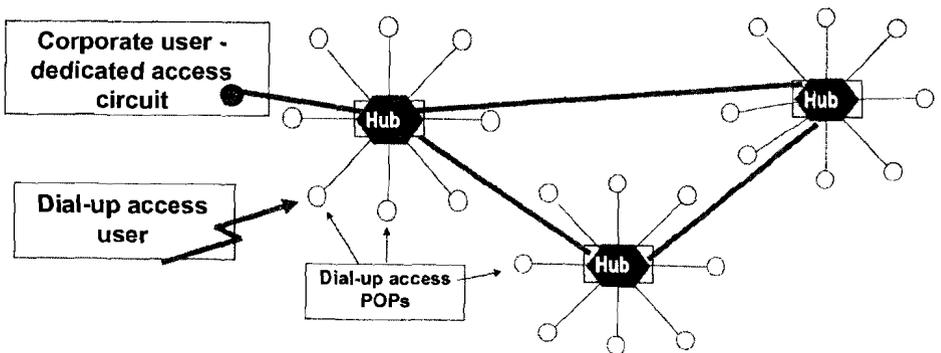


Fig. 2.

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

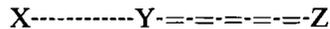
Internet networks have contracts that govern the terms under which they pay each other for connectivity. Payment takes two distinct forms: (i) payment in dollars for

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.

