*** Pages 27 and 28 are missing from original**

PRICING THE INTERNET: HOW TO PAY THE TOLL FOR THE ELECTRONIC SUPERHIGHWAY

BRETT A. PERLMAN

95-01 **MARCH 1995**

CITATION AND REPRODUCTION

This document appears as Discussion Paper 95-01 of the Center for Science and International Affairs. CSIA Discussion papers are works in progress. Comments are welcome and may be directed to the author in care of the Center.

This paper may be cited as: Brett A. Perlman. " Pricing the Internet: How to Pay the Toll for the Electronic SuperHighway." CSIA Discussion Paper 95-01, Kennedy School of Government, Harvard University, March 1995.

The views expressed in this paper are those of the authors and publication does not imply their endorsement by CSIA and Harvard University. This paper may be reproduced for personal and classroom use. Any other reproduction is not permitted without written peg Mission of the Center for Science and International Affairs, Publications, 79 JFK Street, Cambridge, MA 02138, telephone (617) 495-3745 or telefax (617) 495-8963.

TABLE OF CONTENTS

I. Introduction: Building the Internet's Commercial Information Infrastructure 1

II. History of the Internet and the Growth of the Commercial Internet 3

A. The Internet's Origins.....4
B. NSFNet Expands the Internet for Academic and Research Users 4

1. Engineering Implications of NSFNet 6
2. Managerial Structure of the Internet 7
3. Economic and Regulatory Aspects of the Internet ...8

C. First Steps in the Transition to Commercial Internet Applications 9

D. CIX Enters the Fray 13

E. The ANS-CIX Dispute 14

F. A Congressional Hearing on the Internet 15

G. ANS and CIX Reach a Temporary Agreement 16

II. The Impact of Commercialization on the Internet's Continuing Structural Evolution 17

III. Pricing The Internet: How to Pay the Toll 19

A. Internet Pricing and the Telephone Model 19
B. Settlements v. the Free Market : How to Price the Internet? 23

1. An Economic Framework for Internet Connectivity.. 23
2. Will the Free Market Maximize Internet Interconnections? 24

a. The Transit Network Problem 25
b. Network Externalities 27

V. Conclusion 28

I. Introduction: Building the Internet's Commercial Information Infrastructure

By almost any calculation, the Internet has been a huge success. The Internet started as government-sponsored research and development network and has rapidly evolved within the last five years into a "network of networks" linking not only academic and government institutions, but also commercial users. While the Internet continues its rapid growth, the commercial Internet applications have grown fastest.[1]

The Internet is now making the transition from a government-funded research and education network to a non-subsidized commercial network.

Privatization and commercialization are the two drivers of the Internet's future commercial development. As the Internet becomes "commercialized," users will have the benefit of using and accessing commercial Internet facilities and services. At the same time, the Internet will become increasingly "privatized" with the federal government no longer subsidizing network service.[2]

The Internet's transition has not been without problems. Both users and providers have had to adapt in the transition from a federally subsidized monopoly service to a market-based commercial network.

---

[1] "Booming Commercial Use Changes Face of the Internet," InfoWorld April 12, 1993, p.1.

[2] See Kahin, Summary Report on the Commercialization of the Internet (unpublished paper, Harvard University 1990) (description of issues relating to the Internet's transition to a private, commercial network) (hereinafter cited as "Commercialization Report").

This paper recounts one story of the Internet's transitional growing pains. In 1991, a dispute arose between a consortium of commercial networks represented by the Commercial Internet Exchange ("CIX") and the government-contract provider of network backbone services, Advanced Network Services ("ANS").[3] The dispute involved whether ANS had abused its position to gain an unfair competitive advantage.

The dispute arose when ANS formed a commercial subsidiary, called ANS CO+RE (Commercial + Research) Systems, Inc., which competed with the other private network vendors. As a result of its government contract, ANS provided the NSFNet backbone service. ANS implemented the federal government's Acceptable Use Policy ("AUP"), which restricted the NSFNet backbone service from carrying commercial traffic. CIX members charged that this requirement provided ANS with an unfair commercial advantage. ANS claimed that its commercial operations served the public interest by subsidizing research and education users from the fees generated by commercial use of the network.

The dispute erupted publicly when the parties aired their grievances before the U.S. House Subcommittee on Science on March 12, 1992.[4] Following the public hearing, ANS and CIX reached a trial interconnection agreement and decided to discuss how ANS

---

[3] "Battle Is On for Computer Network Right of Way: New Limits to Commercial Activity on Internet Stir Criticism," Houston Chronicle, December 22, 1991, Sec. A, p. 21.

[4] Hearing on the National Science Foundation Network, before the Subcommittee on Science, Space and Technology, March 12, 1992 (hereinafter cited as "House Hearing").

should charge for network backbone service.

This paper is based on a series of meetings held at the John F. Kennedy School of Government at Harvard University in which representatives from ANS, CIX, American Telephone and Telegraph (AT&T) , Bellcore and others discussed the economic, engineering and public policy issues underlying the ANS-CIX dispute.

The purpose of this paper is to understand these economic, engineering and regulatory issues and to examine the dispute's implications for the Internet's commercial development. The paper will discuss how the Internet's history has impacted the network's commercial transition. The paper will then evaluate whether free market pricing will foster the Internet's commercial growth. Finally, the paper will discuss the public policy issues regarding the Internet's transition to a commercial network.

II. History of the Internet and the Growth of the Commercial Internet

An understanding of the Internet's history helps to explain the origins of the ANS-CIX dispute and the implications of that dispute on the development of commercial Internet services.

Various factors, including network engineering and federal government regulation, contributed to CIX-ANS dispute. While these factors encouraged the Internet's initial success as a government-funded research and education network, they have inhibited the Internet's transition to a commercial network.

A. The Internet's Origins

Although the Internet has grown explosively over the last five years, the network's origins date back over twenty years to the

ARPANET.[5] The federal government's Defense Advanced Research Projects Agency (DARPA) formed ARPANET to facilitate communications between government-funded researchers. This interconnection of experimental and research networks was originally called the ARPA Internet, with the name later shortened to the "Internet."[6]

B. NSFNet Expands the Internet for Academic and Research Users

The National Science Foundation's networking effort expanded over time. NSF's original intention for the network was to allow researchers at remote university locations to utilize NSF-funded supercomputing centers. NSF originally planned to provide supercomputing access through direct point-to-point connections between universities and the supercomputing centers.

Over a short period of time, however, it became clear that connections could be more efficiently provided through a network that linked universities to more than one supercomputing center and to each other as well. By 1986, this network had evolved into a three-tier structure:

- Local area networks that linked computing resources over a single campus.
- Regional "midlevel" networks that connected traffic traveling between different campus networks within a geographically compact region.

---

[5] A complete history of the Internet's development is contained in Kahin and McConnell, "Towards a Public Metanetwork: Interconnection, Leveraging and Privatization of Government-Funded Networks in the United States." (unpublished paper, Harvard University 1992) (hereinafter cited as "Metanetwork").

[6] LaQuey, The Internet Companion: A Beginner's Guide to Global Networking 3 (1993).

- A network "backbone" which connected the midlevel networks and carried traffic between the regional midlevel networks.[7]

By 1987, the rapid growth of the network had resulted in congestion of the existing network backbone and led NSF to solicit offers for a new, high-speed network backbone, which became known as the NSFNet backbone.

In November 1987, NSF awarded a cooperative agreement to a consortium consisting of Merit, a part of the University of Michigan, IBM and MCI to provide backbone services. NSF had contracted with Merit to construct the NSFNet backbone because NSF did not have sufficient managerial resources to build and maintain network services. Merit, in turn, subcontracted with IBM to provide router support and maintenance and with MCI to provide the physical circuits.

Although this structure has had enormous engineering, managerial, economic and regulatory implications for the network's commercial development, the network's vast potential for connecting more than supercomputing centers had not yet been recognized. During this initial period, NSF viewed the Internet's mission as the relatively limited one of providing access to universities for high performance computing.

The next sections of this paper describe the network's various structural elements that arose from the Internet's early mission and how they have impacted commercial development.

---

[7] Mandelbaum and Mandelbaum, "The Strategic Future of MidLevel Networks" in Kahin, Building Information Infrastructure (1993).

1. Engineering Implications of NSFNet

From an engineering standpoint, the Internet's three tiers can be described as a "hierarchical" network structure. The NSFNet backbone sat at the apex of the hierarchy and connected non-overlapping midlevel university consortiums.[8] Mid-level networks consisted of a series of leased telephone lines connecting local area networks (LANs) located on university campuses.

This hierarchical structure simplified network engineering since this arrangement required that only the network backbone contain full routing information.[9] At each level of the hierarchy, the network routed traffic to network addresses within its own network and sent traffic destined for another network to the next level of the hierarchy. As a result, a midlevel network would route to the backbone only those messages intended for other midlevel networks and route internally those messages intended for its own network and for campus LANs connected to the midlevel network.

Moreover, the early routing technology constrained the network and meant that messages between regional networks had to cross the NSFNet backbone.[10] An early version of the TCP/IP routing protocol would only permit one gateway between the networks. Although TCP/IP eventually allowed more flexible routing, the Internet retained a hierarchical engineering structure.

---

[8] Presentation to the Harvard "Settlements" Workshop on May 6, 1993 (hereinafter cited as "Harvard Settlements Presentation").

[9] Harvard Settlements Presentation.

[10] Harvard Settlements Presentation.

This hierarchical engineering structure may have been appropriate for the network's original limited purpose of connecting supercomputing centers since this purpose did not require that the network architecture accommodate a high volume of network traffic. As network usage increased with the rapid growth of the midlevel networks, this structure became a factor inhibiting the network's development.

2. Managerial Structure of the Internet

In contrast to the Internet's strict engineering hierarchy, the Internet has a long tradition of managerial decentralization.

Other than the NSFNet backbone, the Internet consists of computing resources owned by separate regional consortia. These regional networks were responsible for providing physical connections and technical support to university campuses. Each regional consortium remained free to develop a managerial and organizational structure best suited for its members. For example, BARRNet, the San Francisco Bay Area Regional Research Network, left the ownership and maintenance of the network routers to universities, while others, such as NYSERNet, the New York State Education and Research Network, contract with local telephone companies to provide network services.[11]

The Internet's technical activities are loosely coordinated by the Internet Society and its overall direction is coordinated by the Federal Networking Council. The Internet Engineering Task Force consists entirely of volunteers who meet at least three times

---

[11] Mandelbaum and Mandelbaum, in Kahin at 67-8.

a year to discuss technical issues and to appoint task forces to deal with these problems. The Federal Networking Council consists of five federal agencies, DOE, DOD, NSF, NIH and NASA, that coordinate network planning for the network.

This decentralized control structure served the network well during its early mission of providing links to a relatively small number of academic researchers. Decentralization allowed each midlevel network to respond to the needs of its constituents and the use of loose coordination allowed broad participation in the network's development.

As will be seen, this combination of decentralization and loose overall coordination may have inhibited the Internet’s continued growth and transition to commercialization.

3. Economic and Regulatory Aspects of the Internet

As long as the network remained a partially federally-subsidized research and education facility with a relatively limited mission, the NSFNet backbone architecture did not impose capacity constraints to the Internet's growth. NSF's subsidization of regional mid-level research and education users resulted in the network's rapid development as more users took advantage of a "free" resource.

NSF, however, managed Internet use by limiting the NSFNet backbone service to "acceptable" uses under its Acceptable Use Policy (AUP). NSF believed that such policies were required by NSF’s authorizing statute, which directed the NSF to “foster and support the development and use of ... technologies, primarily

for research and education in the sciences and engineering."[12] Moreover, NSF wanted to avoid any appearance of competing with existing commercial networks.

As the network's mission changed, the AUP constrained the Internet's commercial growth, because the AUP prohibited commercial use.

C. First Steps in the Transition to Commercial Internet Applications

Each of these factors posed little problem before the "commercialization" and "privatization" of the network. As pressure built on NSF to open the network up to commercial services, however, these factors worked to distort the emerging commercial marketplace.

By 1991, many regional networks began to understand that the network had the potential that vastly exceeded the network's relatively limited initial mission. University research parks had long offered accounts to commercial users on the university local are area network and guest accounts on university computers frequently permitted.

Several regional networks realized that the network had commercial potential and that such commercial services might help to subsidize the use of high-speed data networks by the research and education community even if federal government subsidies ceased.

NYSERNet formed Performance Systems International (PSI), a

---

[12] Kahin and McConnell, Metanetwork at 21 n. 40 (citing Section 3(a)(4) of the National Science Foundation Act of 1950).

for-profit corporation. PSI took over the provision of network services and provided commercial services unconstrained by the limitations imposed by the AUP.[13]

Similarly, the Merit/IBM/MCI consortium saw a similar opportunity.[14] In September 1990, the three announced the formation of Advanced Network Services, Inc., a not for-profit corporation. Merit then negotiated NSF's approval for it to transfer operations of the NSFNet backbone service to ANS, with Merit continuing to be directly responsible to the NSF. This provided a vehicle for IBM and MCI to invest substantially in network infrastructure. IBM and MCI agreed to invest $10 million in order to establish a new network backbone, called ANSNet, which would provide network services for NSFNet traffic. The renegotiation of the backbone agreement also benefited NSF which saw a way to increase network capacity at a rate far in excess of that permitted by its budget.

As a result, ANS could handle the traffic of its customers on ANSNet, constrained only by its obligations to continue switching the NSFNet backbone service and by its own non-profit mission. In June 1991, ANS then formed ANS CO+RE, a for-profit subsidiary, that used ANSNet's own network resources to provide commercial network services that might fail to meet AUP guidelines.

The regulatory, engineering and economic characteristics of

---

[13] Mandelbaum and Mandelbaum, in Kahin at 74.

[14] "NSF Embroiled in Commercialization Debate," Common Carrier Week February 10, 1992.

the ANSNet and the NSFNet backbone service, however, caused difficulties in the transition to commercial service. The same NSF policies that have made the Internet so successful in attracting research and education users now inhibited the Internet's commercial growth.

For example, the AUP represented a barrier to the commercialization of the network because it prohibited "commercial" traffic on the network backbone. The AUP might appear as a reasonable policy for a federally-subsidized network, but in reality, midlevel research and educational networks had always exchanged traffic with industrial users involved with university research and development projects. NSF, in fact, had encouraged this process by writing the AUP in such a way so as to permit such activities. NSF applied the AUP to the purpose for which work was performed, not to specific identity of users. This interpretation, however, meant the AUP was honored only in its breach since NSF had no realistic mechanism to distinguish "commercial" and "noncommercial" traffic and did nothing to enforce the AUP, except after the fact.

On the other hand, the AUP thwarted the use of the NSFnet backbone for purely "commercial" services. By forming the CO+RE subsidiary, ANS bypassed this problem.

The network's tradition of subsidization also constrained the Internet's commercial growth. ANS's original pricing policies for commercial service reflected the long tradition of government subsidization. These policies required that commercial users would

pay the average cost of connection to the NSFNet backbone. Since the average costs charged to commercial users exceeded ANS's marginal cost, ANS avoided a situation in which NSF funds subsidized commercial traffic.

ANS had at least three slightly different pricing schemes for commercial traffic.[15] ANS required that mid-level networks agree to a pricing scheme by signing the ANS CO+RE Gateway agreement before ANS would agree to carry the network's commercial traffic. ANS's proposed (but never implemented) "combits" pricing scheme best shows how ANS intended to subsidize research and education users. Under the combits pricing scheme, commercial users received a network identification address which allowed ANS to measure the number of commercial bits ("combits") over the network. The commercial identification numbers allowed ANS to mark up the number of "combits" as a multiple of the network's charge for research/education traffic. ANS placed the additional revenues collected in a network infrastructure "pool" for division among the regional research and education networks. ANS claimed that by directing funds back to research and education users, it contributed to the "enriching of the network infrastructure" since commercial revenues subsidized the research and education portions of the network.

---

[15] In addition to combits, ANS offered a flat fee arrangement based on the number of commercial sites within a midlevel network and a fee arrangement based on the fraction of membership fees collected from commercial customers by a midlevel network. **See** Exhibit 2 (Almes, An Inventory of Current Internet Pricing Models). The common characteristic of each fee models involved a payment by the midlevel network to subsidize research and education networks.

ANS's pricing policy raised several problems. While ANS stated that the "goal is to insure that government is not subsidizing commercial traffic,"[16] many commercial users saw this policy as an abuse of ANS's position as the bottleneck supplier of network backbone services. From an engineering standpoint, the Internet's hierarchical structure was optimized by having traffic going from one regional network to another cross the NSFnet backbone. Although regional traffic did not have to go through the backbone, the backbone provided the most efficient means of transport. As a result, ANS's control of the network backbone and its insistence that commercial users have to "use our [pricing] model" in order to make use of the T-3 backbone,[17] appeared to many commercial networks as a misuse of ANS's position.

Another problem was that some believed that ANS had conflicting interests. While ANS operated the NSFNet backbone, it also offered a commercial service which competed directly with other commercial network service providers. As a result, some thought that ANS controlled access while it competed in the commercial service marketplace.

D. CIX Enters the Fray

The Commercial Internet Exchange (CIX) became the primary commercial network advocate in the controversy surrounding ANS's decision to enter the commercial Internet business and to set

---

[16] "Data Network Raises Monopoly Fear," New York Times, December, 19, 1991, Sec. D, p,.7.

[17] "NSF Embroiled in Commercialization Debate", Common Carrier Week February 10, 1992.

prices for commercial access to the NSFNet backbone.

The CIX had been formed in 1991 as a trade association open to all commercial Internet providers. These commercial providers sought to provide commercial network access to companies with links to the university research community. CIX consists of a router in Northern California connected by private links to routers belonging to its members. By sending traffic through this central router, CIX members could exchange both commercial and research/education traffic.

CIX represented an example of commercial maturation of the Internet. CIX provided an alternative to the Internet's hierarchical structure and provided competition to ANS for the transit of Internet messages. On the other hand, CIX's impact was limited since CIX traffic constitutes only a small part of the network.

While CIX competed with ANS as a traffic transit point, it differed in several significant respects. Unlike ANS, CIX members exchanged traffic at fixed costs through unmetered service. In addition, CIX members exchanged traffic without crossing the NSFnet backbone and were therefore free from the NSF's AUP restrictions.

E. The ANS-CIX Dispute

In December 1991, ANS and CIX became embroiled in a pricing dispute. ANS claimed that the CIX's refusal to accept ANS's pricing scheme meant that ANS was not being fully compensated for its investment in the network infrastructure. CIX claimed that ANS was misusing its position by forcing commercial network vendors to

choose either to subsidize Internet research and education members or not to communicate through the NSFNet backbone. CIX Chairman, Mitch Kapor, stated that ANS "should not be permitted to indulge in ephemeral cost-sharing . in return for a monopolistic right to sell commercial access".[18]

The ANS-CIX dispute highlighted a major problem in the Internet's commercial future. Under the NSF's subsidization policy, interconnection had been encouraged and subsidized. On the other hand, the decentralization of the Internet's structure implied a possibility of network fragmentation if midlevel networks did not have the proper incentives to interconnect.

An incident involving Dialog Information Services,a commercial information services vendor, exemplified the fragmentation problem. Dialog had bought commercial service from ANS CO+RE to obtain access to the entire Internet. ANS, however, had placed filters on network backbone so that traffic from commercial customers would not reach midlevel networks that had not signed ANS connectivity agreements.[19] Dialog initially found that it had access to only 25 percent of the network because 75 percent of the midlevel research and education networks at first refused to sign ANS connectivity agreements.[20]

F. A Congressional Hearing on the Internet

---

[18] Cook, "Affordable Access is Key to NREN's Success," Network World March 1, 1993, p.26.

[19] Cook, "A National Network That Isn't," Computerworld, March 9, 1992, p.91.

[20] Id.

During March 1992, the House Science Subcommittee aired the grievances that had arisen between commercial network providers and ANS.

The Committee heard testimony from William Schrader, president of PSI, a commercial network vendor, regarding the "distortions" caused by NSF's Acceptable Use Policy and ANS's pricing policy. Schrader testified that ANS had abused its position as the bottleneck supplier of NSFNet backbone services to gain a competitive position over commercial competitors.[21] Similarly, Mitch Kapor testified that the AUP had prevented the commercialization of the Internet because NSF permitted only ANS to carry commercial traffic over the network.[22]

In response, Stephen Wolff, director of NSF Internet activities, testified that NSF was concerned that commercial service providers were using the NSFNet backbone as a transit network to interconnect fee-paying customers in violation of the AUP.[23]

G. ANS and CIX Reach a Temporary Agreement

In June 1992, ANS and CIX reached an agreement to interconnect networks for a "provisional period" in an effort to settle the

---

[21] Testimony of William Schrader, Hearing on the National Science Foundation Network before the Committee on Science, Space and Technology, March 12, 1992.

[22] Testimony of Mitchell Kapor, Hearing on the National Science Foundation Network before the Committee on Science, Space and Technology, March 12, 1992

[23] Testimony of Stephen Wolff, Hearing on the National Science Foundation Network before the Committee on Science, Space and Technology, March 12, 1992

dispute and to allow connectivity between CIX and ANS network users.[24] The agreement allowed ANS and CIX to inter-exchange traffic from midlevel research and education networks to commercial providers. ANS and CIX also agreed to study "equitable arrangements" for permanent interconnection through meetings at a planned public workshop sponsored by the Harvard University Information Infrastructure Project.

II. The Impact of Commercialization on the Internet's Continuing Structural Evolution

Commercialization is rapidly changing the Internet. The Internet is rapidly evolving from a hierarchical structure to a "modified mesh" architecture in which midlevel networks communicate without transport across a single network backbone.

Some evidence shows that the Internet is rapidly moving toward a "modified mesh" structure. Exhibit 2 contains a diagram of the "modified mesh" architecture. One expert has estimated that approximately 80 to 90 percent of the CIX network is presently interconnected to the rest of the Internet without routing through the NSFNet backbone.[25]

This changing architecture represents a third wave in the Internet's economic structure. While NSF first provided free research and education transit and ANS sought a consistent pricing model, midlevel networks will now have increasing choices for interconnection.

---

[24] Communications Daily, June 10, 1992.

[25] Rick Adams at Harvard Settlements Presentation.

Even today, midlevel networks have choices regarding transport. First, they can sign "gateway agreements" with ANS. Under such agreements, ANS will carry their commercial traffic across the NSFNet backbone and connect with other midlevel networks or with CIX members. For example, Nearnet, the New England Academic and Research Network, uses ANS to transport its commercial traffic to CIX commercial networks.

Alternatively, midlevel networks can join the CIX and use the CIX to hook up their commercial traffic with other CIX members. While a CIX membership does not provide connections to non-member midlevel networks, networks can gain such connections by arranging for transport by ANS or other alternative carriers For example, BARRNet, the Bay Area Regional Research Network, is a CIX member. BARRNet's CIX membership allows it only to transport commercial traffic over the CIX router to other CIX members.

While these arrangements reflect the network's increasing interconnectedness and complexity, the NSFNet backbone and the AUP still remain as significant impediments to the free flow of commercial traffic throughout the Internet.

NSF has recently taken steps to eliminate these problems. The NSF has endorsed the "modified mesh" concept in its recent solicitation to replace the NSFNet backbone. Instead of funding the NSFNet backbone, NSF will directly subsidize midlevel networks, leaving the mid-levels free to use these NSF-allocated funds to

forge their own high speed connections.[26] The NSF intends the solicitation to avoid the problems inherent in the NSFNet backbone service. Commercial providers will build Network Access Points (NAPs) which will create a proliferation of points where regional providers and other commercial vendors could interconnect and share traffic.

III. Pricing The Internet: How to Pay the Toll

The CIX-ANS debate highlights the tension created between the Internet's commercial transition and its tradition of providing free subsidized interconnections. On one hand, subsidization encouraged network interconnections but avoided economic efficiency considerations regarding such interconnections. In contrast, Internet commercialization will require a close examination of how to achieve economic efficiency while promoting network connectivity. As the Internet becomes commercialized, network users will be forced to address the issue of whether the free market can provide an optimal method of network connections.

This section first shows how the Internet's development paralleled and contrasted to the history of the telephone network interconnection and subsidization policies. The paper then examines whether the Internet should follow telephone network's model of "settlements"-based pricing policies or whether a free market model is more appropriate.

A. Internet Pricing and the Telephone Model

---

[26] "NSF Changes Course in Its Internet Plan," Network World December 21, 1992, p.1.

It is no coincidence that the NSF and ANS drew on telephone industry antecedents for developing the Internet's economic policies. Both voice and data networks have many similarities. Consequently, the voice network provided an ready analogy upon which to base the Internet's development. Subsidization and interconnection were issues both in the development of the telephone network and the Internet.

First, the NSF's policy of promoting interconnections through subsidization is similar to the policy of "universal service" in the telephone industry. The Communications Act of 1934 required that the telephone network provide for "universal service" so that all citizens wanting telephone connections could obtain them. As a result, average telephone network penetration rates in the U.S. exceed 95 percent.[27]

Like the telephone network, NSF promoted universal Internet interconnections between regional networks by providing free transit across the NSFNet backbone. NSF's backbone subsidy resulted in a form of universal service by encouraging interconnection between autonomous regional networks.

NSF's subsidization policy was even more important in building the data communications network than in the telephone network since Internet consists of nothing more than a "loose amalgam" of

---

[27] Johnson, Competition and Cross-Subsidization in the Telephone Industry 77 (Rand Corp. 1982). Telephone penetration varies from 72 percent in Alaska to essentially 100 percent for several states.

independent regional data networks.[28] These networks have voluntarily agreed to interconnect because the substantial benefits that accrued from communicating with the other networks' users. Moreover, interconnection also had no costs because NSF subsidized the network.

NSF's subsidization policy, like universal service, created broad access to the Network and has been a major factor responsible for the Internet's growth. The federal government has heavily subsidized access to the network by underwriting the costs of creating the NSFNet backbone and subsidizing the research and academic users of the network. NSF's subsidy stimulated the Internet’s growth for research and development purposes by providing free backbone service to NSF-funded midlevel networks.[29]

Second, the telephone network provided an analogy in developing the mechanism to fund the Internet's subsidy.

Prior to the breakup of AT&T, the telephone industry relied on a very complicated subsidization mechanism, called "separations and settlements." Under separations, regulators made arbitrary allocations of network costs and placed additional revenues collected from overpriced long distance services into a settlement pool which was then paid to subsidize local telephone service.[30] "Separations and settlements" implemented universal service by

---

[28] LaQuey, the Internet Companion, at 1

[29] Kahin, Commercialization Report at 3

[30] Grace, Access and the Demise of Separations and Settlements, Public Utilities Fortnightly, September 1, 1983, p. 17

subsidizing high cost users in rural or less densely populated areas that arguably would not have received telephone service absent the subsidy.

It is no coincidence that ANS referred to its pricing policy as "settlements" since the policy had a clear telephone industry heritage. While the NSF originally subsidized the network by providing free backbone services, ANS's "combits" pricing scheme closely resembled some aspects of traditional telephone industry separations and settlements. ANS attempted to charge a higher education users and to use the excess revenue to subsidize research and education users through the infrastructure "pool."

The telephone industry's transition to deregulation and competition provides an third analogy for the Internet.

The opening of the long distance telephone market to competition destroyed the ability of regulators to "settle" by overpricing long distance service in order to subsidize local service. In order to continue to subsidize local service, which regulators believed important both to keep telephone penetration rates and to avoid political pressure, a system of access charges replaced the separations and settlements procedures.

The commercialization of the Internet inhibited the ability to subsidize regional users through either free backbone services or a settlements mechanism in a competitive market environment.

First, free backbone service clearly caused economic distortions. Even before Internet commercialization, it was well

recognized that NSF's subsidy of free backbone services resulted in inefficient network use since, with a zero price, no incentives existed to use the network in an economically efficient manner.[31] Since users viewed the Internet as a free, common resource, it would inevitably have been subject to overuse, the classic tragedy of the commons. Internet commercialization would have exacerbated this problem. In the absence of a market price, networks would have incentives to forego upgrading their own networks and instead take a free ride on other networks' transport.

Second, ANS's "settlements" scheme could not withstand competition because it imposed higher costs on competitive providers and encouraged bypass of the network. In this light, the CIX represented an attempt to bypass the high cost NSFNet backbone.

B. Settlements v. the Free Market: How to Price the Internet?

On one level, the NSF's decision to fund mid-level networks directly instead of using a non-market "settlements" pricing scheme resolved the ANS-CIX dispute. On another level, a question still remains regarding whether a market mechanism can assure a highly interconnected network or whether some sort of "settlements" is required to encourage interconnections as the Internet becomes commercialized. Because the Internet contains a highly decentralized structure, the proper economic incentives are crucial to promote interconnectivity.

1. An Economic Framework for Internet Connectivity

---

[31] Kahin, Commercializatioin Report at p. 3-4.

A few fundamental economic principles provide the framework for analyzing whether the free market will promote Internet connectivity.

Economists view the free market as the most efficient mechanism for promoting economic welfare because it maximizes the number of voluntary exchanges between parties.[32] Economists favor voluntary free market exchanges because such exchanges permit individuals to maximize their individual preferences. Economists define an optimum situation (a Parieto optimum) as that where individuals have maximized those voluntary exchanges by pricing goods or services at their marginal cost.

Economics recognizes, however, marginal cost pricing may not represent economic efficiency where supply and demand do not internalize all of the market preferences. In such situations, the market must be corrected either through voluntary agreements that internalize the externalities or through some type of regulation.[33]

2. Will the Free Market Optimize Internet Interconnections?

The goal of any Internet pricing mechanism should be to optimize the number of economically efficient interconnections. On one hand, it is important to distinguish the optimum efficient number of Internet connections from the goal of universal Internet interconnectivity.

---

[32] For a general discussion of welfare economics in the telecommunications industry, see Chapter 2 (Economic Efficiency, Competition and Cartels) in Wender, The Economics of Telecommunications (1987).

[33] Id. at 29.

While it would be possible to have full Internet interconnectivity for all users, such a goal would not be economically efficient. There are clear substitutes for Internet communications that may be more efficient for some users. For example, it may be more efficient for an individual to make a telephone call, send a letter or walk down the hall than to provide Internet connections between every computer. Similarly, it may be more efficient to have some regional networks connect directly to each other than to join a Internet exchange point. As long as each regional network remains free to make those decisions (and assuming no externalities), then the free market will correctly determine the most efficient level of Internet connections.

a. The Transit Network Problem

Interconnection problems might arise as the Internet becomes commercialized because each network constitutes a separately owned resource. Under the hierarchical model, NSF encouraged connectivity by providing a subsidy for transit traffic through free use of the NSFNet backbone. With the commercialization of the Internet, the NSF network backbone subsidy will cease and will be replaced by a direct subsidy to midlevel networks. Moreover, the NSF draft solicitation contemplates that commercial vendors will build several Network Access Points (NAPs) in order to transport traffic across the Internet.

Under these circumstances, networks must be compensated for the use of their facilities. Therefore, the issue becomes how to compensate networks for traffic transport in an economically

efficient manner.

The combination of the Internet's new mesh architecture and the transition to commercialization appear to provide the correct incentives because each regional network can choose how to interconnect. In such an environment, voluntary bilateral interconnection agreements appear to be one way to achieve the goal of economic efficiency. In order to communicate with CIX members, a regional network must either agree to join the CIX by signing an interconnection agreement, which mandates zero-settlements, or choose one of several alternative methods of connecting with the CIX router. For example, a regional network could choose to pay ANS to transport its commercial traffic to the CIX router, it could choose a third party, such as UUNet, or it could choose to lease its own lines. Alternatively, it could choose not to connect to the CIX. Under any of these circumstances, the regional network's voluntary decision how and whether to connect maximizes economic efficiency as long as multiple routes exist and the players can bargain over the cost of transporting traffic across the network.

In comparison, "settlement" methods for compensating transit networks usually distort free market forces because settlements use some factor other than marginal cost to determine access charges. For example, U.S. international telecommunications carriers settle international telephone traffic with foreign government- controlled telephone monopolies (PTTs) by dividing revenue based on an artificial accounting rate instead of a voluntarily agreed free market price. PTTs prohibit foreign telephone companies from owning

function and to replace Internet's original monopolistic pricing and hierarchical network structure.

Even though the market appears to be functioning, enlightened public policy must still play a role in the Internet's development in two areas. First, NSF should directly subsidize research and education users to promote network connectivity. Second, NSF should determine whether the network externality justifies any further governmental role.