

Pricing the Internet

by

Jeffrey K. MacKie-Mason

Hal R. Varian

University of Michigan

April 1993

Current version: February 10, 1994

Abstract. This paper was prepared for the conference “Public Access to the Internet,” JFK School of Government, May 26–27, 1993. We describe the technology and cost structure of the NSFNET backbone of the Internet, and discuss how one might price Internet access and use. We argue that usage-based pricing is likely to be necessary to control congestion on the Internet and propose a particular implementing of usage-based pricing using a “smart market”.

Keywords. Networks, Internet, NREN, NII.

Address. Hal R. Varian, Jeffrey K. MacKie-Mason, Department of Economics, University of Michigan, Ann Arbor, MI 48109-1220. E-mail: jmm@umich.edu, halv@umich.edu.

Pricing the Internet

Jeffrey K. MacKie-Mason

Hal R. Varian

On December 23, 1992 the National Science Foundation (NSF) announced that it will cease funding the ANS T3 Internet backbone in the near future. This is a major step in the transition from a government-funded to a commercial Internet. This movement has been welcomed by private providers of telecommunication services and businesses seeking access to the Internet.

No one is quite sure about how this privatization will work; in particular, it is far from clear how use of the privatized Internet will be priced. Currently, the several Internet backbone networks are public goods with exclusion: usage is essentially free to all authorized users. Most users are connected to a backbone through a “pipe” for which a fixed access fee is charged, but the user’s organization nearly always covers the access fee as overhead without any direct charge to the user.¹ None of the backbones charge fees that depend at the margin on the volume of data transmitted. The result is that the Internet is characterized by “the problem of the commons,” and without instituting new mechanisms for congestion control it is likely to soon suffer from server “overgrazing.” We shall propose an efficient pricing structure to manage congestion, encourage network growth, and guide resources to their most valuable uses.

We first describe the Internet’s technology and cost structure, since a feasible and efficient pricing scheme must reflect both technology and costs. We then describe congestion problems in the network, and some past proposals to control it. We turn to pricing by first describing in general terms the advantages and disadvantages of using pricing to control congestion, followed by the details of our proposed pricing structure. We devote particular attention to a novel feature of our proposal: the use of a “smart market” to price congestion in real time.

We wish to thank Guy Almes, Eric Aupperle, Hans-Werner Braun, Paul Green, Dave Katz, Mark Knopper, Ken Latta, Dave McQueeney, Jeff Ogden, Chris Parkin, Scott Shenker and Paul Southworth for helpful discussions, advice and data. We are also grateful to James Keller and Miriam Avins for extensive, helpful editorial advice. MacKie-Mason was visiting the Department of Economics, University of Oslo when this paper was completed.

¹ Most users of the NSFNET backbone do not pay a pipeline fee to ANS, the service provider, but instead pay for a connection to their “regional” or mid-level network, which then is granted a connection to the NSFNET.

1. Internet Technology and Costs

The Internet is a network of networks. We shall focus on backbone networks, although most of our pricing ideas apply equally well to mid-level and local area networks. There are essentially four competing backbones for the Internet: ANSnet, PSInet, Altnet, and SprintLINK.² ANS is a non-profit that was formed in 1990 to manage the publicly-funded NSFNET for research and educational users. ANSnet now provides the backbone service for NSFNET, as well as backbone service for commercial users through its subsidiary, ANS CO+RE, Inc. PSInet and Altnet are independent commercial providers of backbone Internet services to commercial and non-commercial users. Sprint, of course, is a major telecommunications provider as well as a provider of Internet transport services.

The Internet networks use packet-switching communications technology based on the TCP/IP protocols. While much of the traffic moves across lines leased from telephone common carriers, packet-switching technology is quite different from the circuit-switching used for voice telephony. When a telephone user dials a number, a dedicated path is set up between the caller and the called number. This path, with a fixed amount of network resources, is held open; no other caller can use those resources until the call is terminated.³ A packet-switching network, by contrast, uses “statistical multiplexing”: each circuit is shared by many users, and no open connection is maintained for a particular communications session. A data stream is broken up into small chunks called “packets.” When a packet is ready, the computer sends it onto the network. When one computer is not sending a packet, the network line is available for packets from other computers. The TCP (Transmission Control Protocol) specifies how to break up a datastream into packets and reassemble it; the IP (Internet Protocol) provides the necessary information for various computers on the Internet (the routers) to move the packet to the next link on the way to its final destination.

The data in a packet may be 1500 bytes or so. Recently the average packet on NSFNET carries about 200 bytes of data (packet size has been steadily increasing). On top of these 200 bytes the

² In addition, a new alliance called CoREN has been formed between eight regional networks and MCI. This represents a move away from the traditional backbone structure towards a mesh-structured set of overlapping interconnections.

³ Some telephone lines are multiplexed, but they are synchronous: $1/N$ th of the line is dedicated to each open circuit no matter how lightly used is that circuit.

TCP/IP headers add about 40; thus about 17% of the traffic carried on the Internet is simply header information.

Packetization allows for the efficient use of expensive communications lines. Consider a typical interactive terminal session to a remote computer: most of the time the user is thinking. The network is needed only after a key is struck or when a reply is returned.⁴ Holding an open connection would waste most of the capacity of the network link. Instead, the computer waits until after a key is struck, at which point it puts the keystroke information in a packet which is sent across the network. The rest of the time the network links are free to be used for transporting packets from other users.

The other distinguishing feature of Internet technology is that it is “connectionless.”⁵ This means that there is no end-to-end setup for a session; each packet is independently routed to its destination. When a packet is ready, the host computer sends it on to another computer, known as a router (or switch). The router examines the destination address in the header and passes the packet along to another router, chosen by a route-finding algorithm. A packet may go through 30 or more routers in its travels from one host computer to another. Because routing is dynamically calculated, it is entirely possible for different packets from a single session to take different routes to the destination.⁶

The postal service is a good metaphor for the technology of the Internet (Krol (1992), pp. 20–23). A sender puts a message into an envelope (packet), and that envelope is routed through a series of postal stations, each determining where to send the envelope on its next hop. No dedicated pipeline is opened end-to-end, and thus there is no guarantee that envelopes will arrive in the

⁴ Some interactive terminal programs collect keystrokes until an `Enter` or `Transmit` key is struck, then sends the entire “line” off in a packet. However, most Internet terminal sessions use the `telnet` program, which sends each keystroke immediately in a separate packet.

⁵ Some packet-switching networks are “connection-oriented” (notably, X.25 networks, such as Tymnet and frame-relay networks). In such a network a connection is set up before transmission begins, just as in a circuit-switched network. A fixed route is defined, and information necessary to match packets to their session and defined route is stored in memory tables in the routers. Thus, connectionless networks economize on router memory and connection set-up time, while connection-oriented networks economize on routing calculations (which have to be redone for every packet in a connectionless network).

⁶ Dynamic routing contributes to the efficient use of the communications lines, because routing can be adjusted to balance load across the network. The other main justification for dynamic routing is network reliability, since it gives each packet alternative routes to their destination should some links fail. This was especially important to the military, which funded most of the early TCP/IP research to improve the ARPANET.

sequence they were sent, or follow exactly the same route.

The TCP protocol enables packets to be identified and reassembled in the correct order. TCP prefaces the data in a packet with a header containing the source and destination ports, the sequence number of the packet, an acknowledgment flag, and so on. The header takes up 20 or more bytes. TCP sends the packet to a router, a computer that is in charge of forwarding packets to their next destination. At the routers, IP adds another header (another 20 or more bytes) containing source and destination addresses and other information needed for routing the packet. The router then calculates the best next link for the packet to traverse, and sends it on. The best link may change minute by minute, as the network configuration changes.⁷ Routes can be recalculated immediately from the routing table if a route fails. The routing table in a switch is updated nearly continuously.

Over the past five years, the speed of the NSFNET backbone has increased from 56 Kbps to 45 Mbps (“T3” service).⁸ The newer backbones have also upgraded to 45 Mbps. These lines can move about 1,400 pages of text per second; a 20-volume encyclopedia can be sent across the Internet in half a minute. Many regional networks still provide T1 (1.5Mbps) service, but these too are being upgraded.

The transmission speed of the Internet is remarkably high. We recently tested the transmission delay at various times of day and night for sending a packet to Norway from Ann Arbor, Michigan. Each packet traversed 16 links: the IP header was read and modified 16 times, and 16 different routers calculated the best next link. Despite the many hops and substantial packetization and routing, the longest delay on one representative weekday was only 0.333 seconds (at 1:10 PM EST); the shortest delay was 0.174 seconds (at 5:13 PM EST).⁹

Current backbone network costs

The postal service is a good metaphor for packet-switching technology, not for the cost structure of Internet services. Most of the costs of providing the Internet are more-or-less independent of the

⁷ Routing is based on a dynamic knowledge of which links are up and a static “cost” assigned to each link. Currently routing does not take congestion into account. Routes can change when hosts are added or deleted from the network (including failures), which happens often with about 2 million hosts and over 21,000 subnetworks.

⁸ “Kbps” is thousand (kilo) bits per second; “Mbps” is million (mega) bits per second.

⁹ While preparing the final manuscript we repeated our delay experiment for 20 days in October–November, 1993. The range in delay times between Ann Arbor and Norway was then 0.153 seconds and 0.303 seconds.

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts



Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research



With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips



Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS

Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS

Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS

Sync your system to PACER to automate legal marketing.