Internet Economics

Joseph P. Bailey[1]

Working Paper

February 3, 1995 Draft

[Please do not cite or distribute without consent of the author]

---

[1]The author is affiliated with the Research Program on Communications Policy at the Massachusetts Institute of Technology and can be reached by email, bailey@far.mit.edu, or by phone, 617/253-6828.

The author wishes to thank NASA and NSF for funding and administrative support for this research. In particular, Scott Behnke, David Brown, Thomas Lee, Tony Villasenor, Walter Wiebe and Steve Wolff helped greatly with the formulation of this paper.

Abstract

The Internet, a growing network of networks, is an often touted and often misunderstood technology. It has its own infrastructure complete with service providers, private networks, user communities, international links, etc. similar to telephone or telegraph industries. However, the Internet is more decentralized and less application-specific than traditional telecommunication industries. These differences extend beyond the technology of the Internet to the economics and policies.

This paper explores Internet economics, a growing field which encompasses the technology, economics, and policy surrounding the Internet. It identifies different communities who have overlapping interests in this field and how their preconceived notions of how the Internet operates create contradictory views on how this technology will continue to evolve in a self-sustaining manner. Specifically, it explores the issue of usage sensitive pricing versus flat-fee pricing as applied to the Internet. It provides anecdotal evidence to support general views communities have towards pricing issues. It identifies roadblocks to implement usage sensitive pricing on the Internet. It describes short-term projects that will provide better data to enable a better understanding of the issues.

[T]he expected increase in video and audio applications suggests that pricing as a mechanism for rationing access to infrastructure will become more generally needed over time. -- "Realizing the Information Future", NRC, 1994, p. 8.

Introduction

Understanding the Internet as an economic system is difficult to do because of its almost anarchic form. Unlike the phone system which developed as a single network to provide one service (voice), the Internet is home to many applications and networks. Also unlike the phone system, the Internet has no billing infrastructure in place. Instead of users getting charged after they use the service on a usage-sensitive basis, they are charged a flat fee for some kind of connect charge. Also, because of the aggregation of users at a local site, the flat fee is usually charged at a much higher level than the user level.[2]

However, the flat fee model may not work as the Internet grows. While Internet traffic is growing around 20% per month, new networks are growing about 7% per month.[3] This translates into more traffic generated by the same number of users. The traditional way to accommodate this network growth is to increase capacity and charge it back to the users. However, since usage is increasing faster than the number of users, the current flat fee will have to increase even though the quality of service doesn't necessarily increase.[4] What is necessary is an Internet cost recovery system which is equitable for the users. While some have suggested

---

[2]This is largely due to the technology which doesn't distinguish network traffic to the different users. While Internet Protocol (IP) packets do have information that distinguishes where the come from (the IP source address) this information isn't used at the network routers which only look at the destination IP address to increase throughput and keep costs lower.

[3]A. Rutkowski, presentation to the Japanese Internet Society (ISOC), July 1994. See http://www.isoc.org/interop-tokyo.html.

[4]For example, a person who only uses the Internet for email will have to absorb some of the increase in price even though email usage doesn't lead to wide area networking investments. It is multimedia applications like CU-SeeMe and NCSA Mosaic that lead to these investments.

usage sensitive billing as such a mechanism it is unclear whether or not usage-sensitive billing is possible with the current technology.[5] Even if it is, usage sensitive billing will undoubtedly discourage use of the Internet which goes against the research and educational goals of the National Research and Educational Network (NREN).[6]

Is the Internet Broken?

The Internet, as a system, is viewed very differently by different user communities. Three main user communities are acknowledged here and they each answer the question "is the Internet broken?" differently. Figure 1 shows a Venn diagram of these user communities. Although these communities are mostly separate and some people/organizations operate where there is overlap, there is no common answer for all of these communities. It is this area (denoted by a "*" in Figure 1) where we need to find common ground and an understanding of Internet economics.

Figure 1

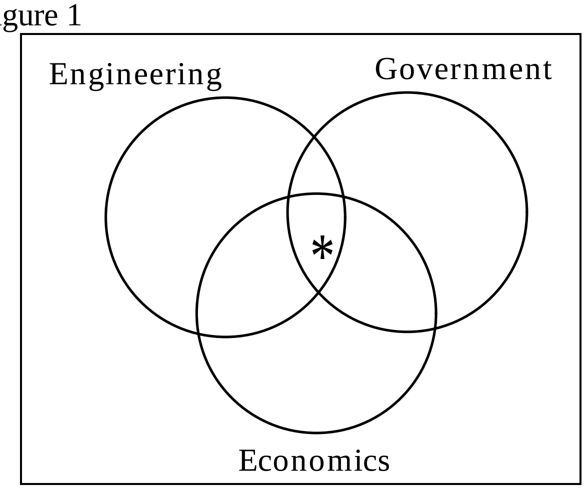

This paper hopes to outline the views of the three communities in Figure 1 and outlines the work done in this area.

---

[5]MacKie-Mason (1994) has suggested usage-sensitive billing. Bohn (1994) suggests a less sensitive billing scheme but one more feasible with today's technology.

[6]See Gupta (1994).

Government

The government community feels that the funding structure for the Internet may be broken. The Internet has many components which are funded by the government (NSFnet, ESnet, NSI, etc.) and this funding will not go away by privatizing and commercializing just the backbone, the NSFnet.[7] Therefore, the government stresses the need to be conscientious in the way different parts of the government deploy their networks. One of the main documents that outlines this guidance is the Office of Management and Budget's (OMB) Circular A-130, "Management of Federal Information Resources." This circular promotes the sharing of resources between the different federal agencies to save money. It also states that the costs for these information services should be accounted for so that there will be an equitable sharing of the costs to provide these network services. However, A-130 has changed recently in July of 1994 from its original 1985 form in the area of cost accounting and recovery for federal networks. The old circular states:

- "Price each services...to the users of that service on an equitable basis..."[8]
- "Directly distribute to the recipient of the services the full costs of...[software, hardware,] and telecommunications equipment..."[9]

The new and revised circular states:

- "For Information Processing Service Organizations (IPSOs) that have costs in excess of $5 million per year..."

---

[7]In fact, most of the 40,000+ networks considered to be part of the Internet are already commercial.
[8]OMB (1985) Appendix II, section 4.c.2.
[9]ibid. Appendix II, section 4.c.3.

•"Recover the costs incurred for providing IPSO services to all service recipients on an equitable basis commensurate with the costs required to provide those services..."[10]

While the old circular was very direct about how to account and recover for costs, the new circular is more general even though the information was moved from an appendix to the main body of the circular. This change was a result of the difficult implementation of the recommendations from the original circular to shared data networks like the IP networks the federal agencies have. Networks like the Defense Data Network (DDN) were hurt by following OMB A-130 too closely (as will be seen later in this paper).

Economics

The economic community feels that the pricing structure isn't optimal. For a better pricing structure, one that is more economically optimal, we could imagine a system where the price equals the marginal cost. However, costs in an IP network are totally fixed except for some kind of congestion cost. This means that when a user on the Internet sends out an e-mail message or initiates an ftp session, they are only using up bandwidth that could otherwise be used by other users of the network. All of the cost associated with people, electronics, leased lines, etc. have been paid for already and are seen as fixed costs.

Congestion, while a cost to other users, is also a deterrent for use. For example, many Internet users will experience high latency (delay) when they use the Internet during peak hours (i.e. business hours) and will choose to do their intensive computing some other time. These applications are fixed-delay in the sense that if the application doesn't respond to the user within an allotted amount of time, there is little or no value to the user. Video conferencing, telnet sessions, etc. are examples of these applications.

---

[10]OMB (1994) Section 8b.4.e

However, congestion, as a measure of marginal cost, has been the focus on the debate in the economic community on the subject on how to price the Internet equitably. Proposals for mapping congestion cost to a monetary cost to the user fall in two main categories, dynamic pricing and pseudo-dynamic pricing – they are defined in Appendix A.

MacKie-Mason (1993) has proposed a pricing system for the Internet that falls under the category of dynamic pricing. Each packet sent includes a willingness-to-pay value data along with the data you send. At the network routers, there is a "bidding" process to determine market price and queuing order. The queue is set up such that the people who pay more for the service get higher priority and, hence, get processed first. The people who are willing to pay nothing for their packets wait until there is no congestion at that router. The host who sent the packet is then charged by packet or byte at the market price. Architecturally, this means that there needs to be greater processing at the network routers than currently exist.[11] Accounting for packets sent and/or received by particular users along with their dynamic price would have to be done at the router as well (which involves electronic storage of this data). Finally, and perhaps most difficult, the infrastructure to bill particular users for their priority traffic by collecting money from the "owners" of IP address.[12]

Another proposal bye Bohn (1994) et al. falls into the category of pseudo-dynamic pricing. Their proposal recommends that each network user pays for "chits" that allow you to send priority packets. You include chits in your packets by inserting them in your precedence bits in the IP

---

[11]Currently, routers try to minimize the amount of time necessary to process IP packets by only reading the destination into random access memory (RAM). If tailored queuing mechanisms were put in place, the throughput of the router decreases or the cost would go up to outfit the router with more processing power.

[12]This proposal doesn't address the overhead costs to implement this system or the possibility that someone may "cheat" the system by including a fake IP source address in their packets.

header (3 bits). More than one chit may be used for one packet, thereby giving it a higher priority level than packets with fewer chits. Therefore, eight different queues are established depending on the precedence you pay for. All of the accounting goes on at the host computer who agrees to purchase these chits at some market price ahead of time. The billing is therefore done beforehand and not after like the MacKie-Mason proposal. This proposal has the advantage that it could be implemented incrementally and it is one that has been accepted before by the Internet community.[13]

Engineering

The engineering community generally believes that the Internet is not broken since it provides a network for reliable applications. If there are problem with service (e.g. I experience large delays when trying to run some applications) there is no short term solution to the problem. Although people are currently talking about bandwidth-on-demand this is very different than the method of operation the IP community has existed in. The asynchronous transfer mode (ATM) community feels that it has solved this problem, but the technology isn't as mature and deployed as IP today.

In the long run, congestion can be mitigated by reconfiguration of the network infrastructure. For example, growth of the NSFnet backbone has led to a few different network architectures over the years. The NSFnet progressed from a 56 Kbps network in 1986 to a 45 Mbps network in 1994. Almost a thousand-fold increase in capacity in less than ten years. If higher capacity lines were not added to the system, it is unclear if the Internet would be such a popular medium today. As capacity got larger with the NSFnet, a priority queuing system that had previously been in place was abandoned.[14]

---

[13] In 1986, when the NSFnet was operating at 56 Kbps, there was very high latency. To mitigate this problem, the engineers of the NSFnet decided to give priority to telnet traffic (the largest fixed-delay application at the time) by queuing by port number. This didn't use the precedence bits suggested by the Bohn proposal.

[14] Operating at T1 capacity (1.544 Mbps) it was no longer necessary to queue by port number.

As Pool (1983) argues, bandwidth is cheap, so this method of over-provisioning may continue to exist in the IP community. Certainly if the latest architecture that will connect up the supercomputer centers (vBNS) ever becomes incorporated into the Internet, there may not be a need to do priority queuing. What is most significant to the resistance of priority queuing is the question of what IP really is – a best-effort-service. This is being questioned by groups like the Internet Engineering Task Force (IETF) even today. Figure 2 outlines a matrix of options similar to that recommended by Schenker (1993) to the IETF.

Figure 2

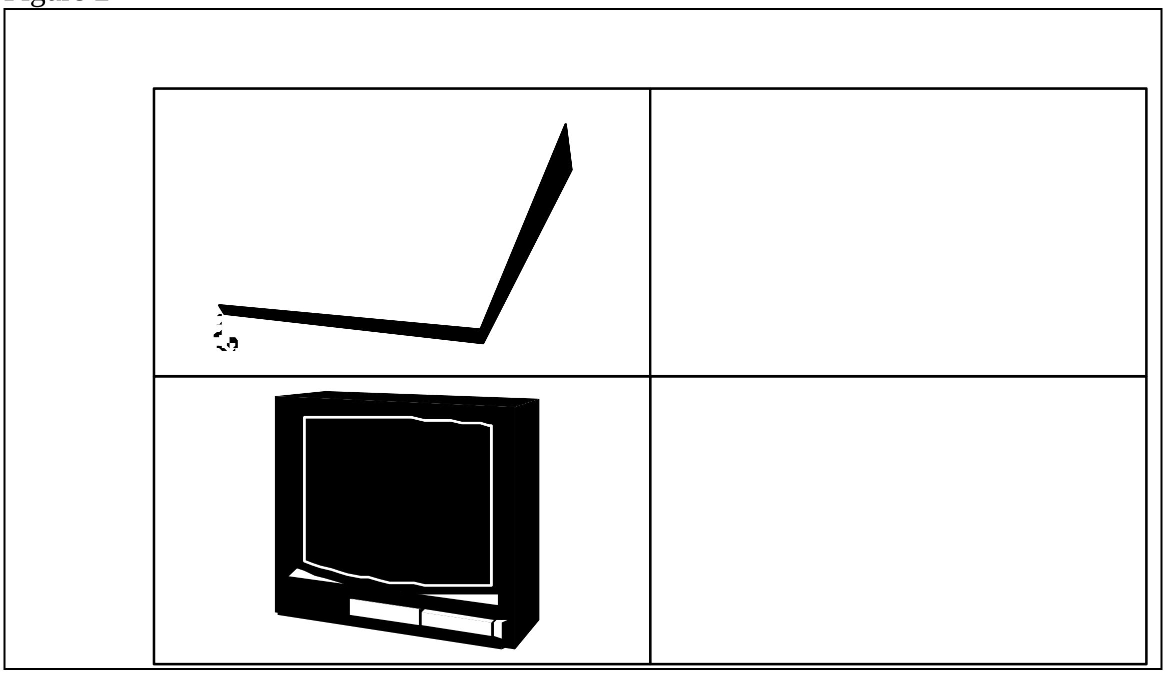

The Department of Defense (DoD) decided to implement a usage based cost recovery and pricing scheme for its inter-agency Internet called the Defense Data Network, the DDN (Barns, 1989), to implement the directive of OMB Circular A-130. The idea was to charge each of the branches of the military for their usage of the network to recover the costs. However, it became apparent to the different DDN users that it was difficult to budget money for networks since usage varied greatly from month-to-month. Instead of encouraging people to use DDN for their networking needs, the different branches of the US military developed their own IP networks which they paid for in the more traditional flat-fee model.

A usage based pricing scheme may not be preferred if consumers are given a choice between usage based pricing and flat-fee. For example, in Italy, an Internet service provider (ISP) had a usage based pricing policy which was replaced by a flat-fee model to meet consumer demands. They charged per email message sent, per megabyte of data received, and hourly fees for other services under their usage based model. Another service provider entered this market, offering the same services, but with a flat-rate annual fee. The migration to the later ISP was so overwhelming that the usage based pricing policy of the former was replaced by a flat-fee model.

These are two instances where the usage-sensitive cost recover model has failed. The reasons for their failure is important to remember when trying to implement new cost-recovery schemes for IP networks:

- It is difficult to budget money for a network expense that is uncertain.
- Users will build their own networks to avoid usage of the usage-based charging networks.
- There will be large administrative costs which may offset any perceived benefits.

Certainly these may not be the only reasons usage-based cost recovery schemes work out, but they outline the major ones learned from the DDN and Italian networks.

This is not to say that usage-cost recovery doesn't work. In fact, it works quite well in New Zealand where the international link to the United States is paid for on a usage-sensitive basis. A major reason that this was successful is that is was a policy implemented before the network was and it was, arguably, the only way to pay for such an expensive link. Australia is in a similar situation but they have not taken action on doing usage-sensitive cost recovery yet.

The NASA Science Internet Solution

NASA Science Internet (NSI) is faced with many of the problems of cost accounting and recovery discussed in this paper. NSI is a growing network where increased demand has been met with increased capacity. When faced with suggestions from the MacKie-Mason and Bohn proposal, NASA is unclear whether or not these solutions will actually help their users. One reason why they don't have an idea how they may be affected is because of the lack of accounting information currently collected.

While it is a difficult task to account for all the data over NSI, it is at least possible to capture a portion of it. Instead of deploying over 200 accounting devices (which is the number or routers NSI has), it would be possible to capture a portion of the data through the five main "hubs" of NSI. Figure 3 shows the location of each of the hubs along with the leased lines that run between them.

Figure 3

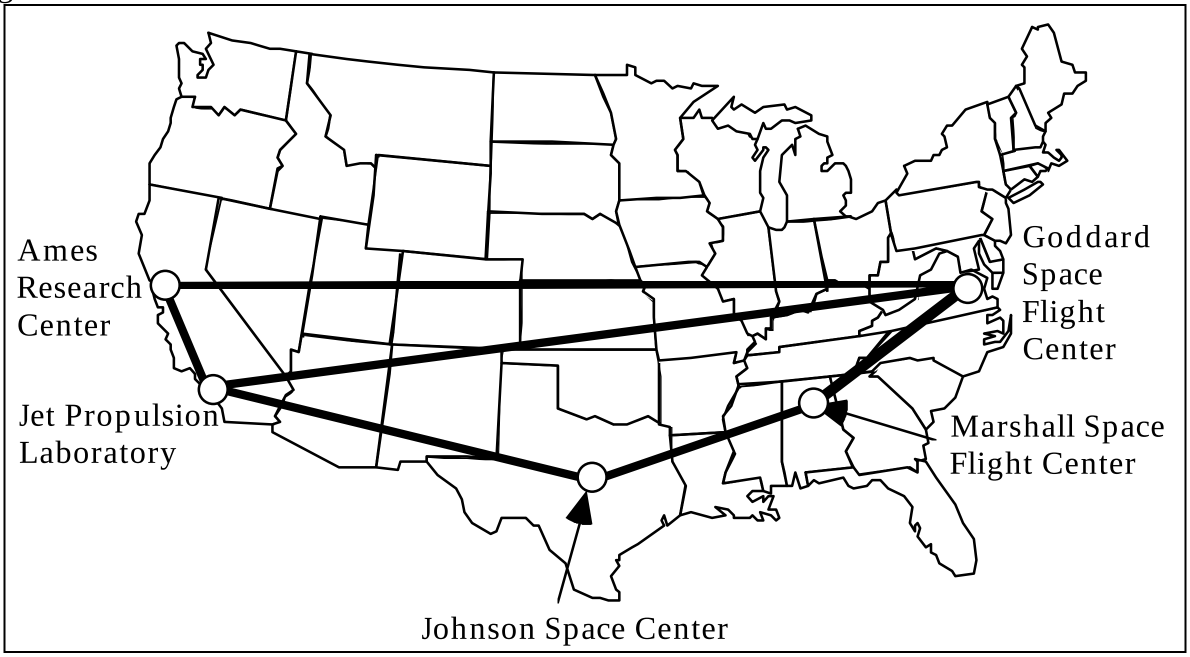

NASA plans to deploy workstations at each of these locations to account for packet flows consistent with the flow counting research done by Braun (1993) and Claffy (1994).

The hope is to characterize the way people use the network to better plan for future growth. Also important is the collection of data to ensure equitable sharing of costs within data. While it is unlikely that these accounting platforms may me the primary purpose of cost allocation, it could be used to corroborate some other policy-oriented cost allocation method.[16]

Once the hardware is in place, NASA foresees potential other uses such as separating out IP traffic flows for a particular hosts or set of hosts. This deployment, while broad in scope, will give the federal IP community a better idea how proposals like MacKie-Mason's and Bohn's may mitigate congestion problems by collecting actual network data.

---

[16] In a conversation with David Brown of Sterling Software, the author and he came up with many different policies such as cost per IP address allocated, costs allocated for each router (including shared routers), etc.

Conclusion

The Internet is still a young technology and the Internet infrastructure hasn't solved all pricing and billing issues even though this will be of greater concern as high bandwidth applications become more pervasive. Probably, like the telephone system, there will be multiple methods of paying for Internet connectivity in the future each giving different qualities of service. This change in technology is unlikely to be accepted when the engineering, government, and economic communities disagree on the future of the Internet. Most likely, the users will determine the best pricing strategy and flat fee may be the near-term solution.

It is unclear how long flat-fee pricing will persist as an expected Internet crunch approaches. Since not all applications demand the same kind of service out of the Internet, it may be necessary to prioritize data traffic for lower levels of latency. In particular, video and voice traffic may need a higher level of service when Internet congestion is high. New services are being addressed by the engineering community (through ATM, IPng, etc.) and appropriate pricing mechanisms will emerge with these new services.

NASA has taken a conservative, interdisciplinary approach to this situation by deploying network statistic accounting equipment at a number of their major Internet hubs. They don't plan on solving all of the problems surrounding Internet economics, but they hope to achieve an equitable cost recovery system for their users. This goal is certainly attainable and will hopefully provide a model for other Internet service providers.

Appendix A: Definitions

BILLING: The process of recovering costs from a customer on a full or partially based usage charge. This definition does not include fixed cost, contract-based services or publicly supported (tax dollar) services. This definition is akin to bills received for phone service. While some portion of the bill may be a fixed cost, there will be some usage-based charges as well.

COST ACCOUNTING: A method by which the money allocated for a service is itemized so that a person knows what services they receive for the money they spend.

COST RECOVERY: The process by which the money for a service is received. For example, in the case of the private sector, the revenue stream is received from the customers who pay money for a service.

DYNAMIC PRICING: Each packet would contain some kind of price bid and the sender of the packet would be charged the "market price" for sending a packet. This market price is dynamic and increases as congestion over the net increases. Low priority traffic, like email, would not get charged since those packets would only be sent when the congestion is below some certain threshold level.

PRICING: A market value associated with a customer's cost for receiving a service.

PSEUDO-DYNAMIC PRICING: A pre-determined price is set for network service and even though it is time sensitive. It doesn’t vary instantaneously like the dynamic pricing model.

References

Barns, W. W. (1989). “Defense Data Network Usage Accounting Enhancement Approaches.” The MITRE Corporation.

Barns, W. W. (1991). “Precedence and Priority Access Implementation for Department of Defense Data Networks.” The MITRE Corporation.

Bohn, R., H.-W. Braun, K. C. Claffy and S. Wolff (1994). “Mitigating the Coming Internet Crunch: Multiple Service Levels via Precedence.” Applied Network Research Technical Report.

Braun, H.-W., K. C. Claffy and G. C. Polyzos (1993). “A Framework for Flow-Based Accounting on the Internet.” Proceedings of SICON ‘93.

Claffy, K. C., Braun, Hans-Werner, George, C. Polyzos (1994). “Internet Traffic Flow Profiling.” San Diego Supercomputer Center.

Cocchi, R., S. Shenker, D. Estrin and L. Zhang (1993). “Pricing in Computer Networks: Motivation, Formulation, and Example.” University of Southern California.

Gupta, A., D. O. Stahl and A. B. Whinston (1994). “Managing the Internet as an Economic System.” University of Texas at Austin.

INI (1993). “Internet Billing Server Prototype Scope Document.” Carnegie Mellon University, Information Networking Institute.

MacKie-Mason, J. K. and H. R. Varian (1993). “Pricing the Internet.” University of Michigan.

MacKie-Mason, J. K. and H. R. Varian (1993). “Some Economics of the Internet.” University of Michigan.

Mills, C., D. Hirsh and G. Ruth (1991). “Internet Accounting: Background.” Network Working Group, Internet Engineering Task Force.

Mitchell, B. and I. Vogelsang (1991). Telecommunications Pricing: Theory and Practice. New York, Cambridge University Press.

NRC (1994). Realizing the Information Future: the Internet and Beyond. Washington, D.C., National Research Council, National Academy Press.

OMB (1985). “Management of Federal Information Resources.” Washington, D.C., Executive Office of the President, Office of Management and Budget.

OMB (1994). “Management of Federal Information Resources.” Washington, D.C., Executive Office of the President, Office of Management and Budget.

Pool, I. d. S. (1983). Technologies of Freedom. Cambridge, Harvard University Press.

Schickele, S. (1993). “The Internet and the Market System: Externalities, Marginal Cost, and the Public interest.” Proceedings of INET ‘93.

Shenker, S. (1993). “Service models and Pricing Policies for Integrated Services Internet.” Xerox PARC.