Internet Economics: What Happens When Constituencies Collide

Joseph P. Bailey <bailey@rpcp.mit.edu>
Lee W. McKnight <mcknight@rpcp.mit.edu>

Abstract

This paper describes the emerging field of Internet Economics and some of the constituencies which are shaping it. It defines the motivating factors for looking at this area, reviews some recent research results, and explores areas of overlapping interest. The paper concludes with six recommendations for further work.

1 Introduction

Internet economics is a growing area of concern to members of the technical, business, academic, and user communities. While new applications, new users, and new connections have spurred on the Internet to become an important medium for communication, information dissemination, and electronic commerce, the economic policies currently undergirding the Internet may be less extensible. This paper examines how the different constituencies' often conflicting and overlapping interests are beginning to yield to a rough consensus.

In particular, the different constituencies are beginning to share the recognition that the growth of the Internet is explained by its economic feature of positive network externalities, its technical characteristic of statistical sharing, and its policy objective of interoperability. This definition of the critical features of the Internet as well as other findings discussed in this paper are partially a result of the authors' work in organizing an Internet Economics Workshop (hereinafter referred to as "the workshop") sponsored by the National Science Foundation and Advanced Research Projects Agency [29]. The workshop was held in Cambridge, MA at MIT on March 9 & 10, 1995.[1]

2 Motivation

The paper first identifies and characterizes the different academic disciplines which have done work related to Internet economics: technology, economics, and policy. Each of these areas has important contributions to make to the field but there has been little consensus across disciplines (or even within disciplines) on which direction to take. Figure 1 shows a Venn diagram which tries to characterize the constituencies' current motivation to look at Internet economics.

Figure 1: Constituencies' Objectives

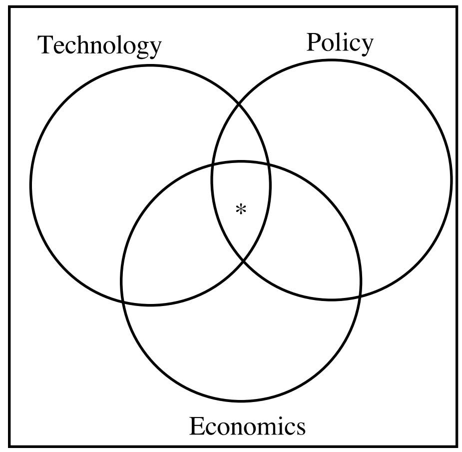

Figure 1 suggests that these communities do not at present agree in their views on Internet Economics. The workshop tried to find the common ground of these constituencies which is denoted by the "*" in Figure 1. First, we identify the "full circle" viewpoint of each of these constituencies.

2.1 Technical

Traditionally, the technical community hasn't concerned itself with the economic issues of the Internet. The current version of Internet Protocol version 4 (IPv4) offers users a best-effort service and it hasn't been necessary to ration a resource. Agreements between Internet Service Providers (ISPs) state that they intend to process traffic as quickly as possible usually on a first-come-first-serve basis. The Internet "cloud" is an inherently statistically shared resource of best-effort service.

When the shared resources of the Internet becomes congested, users experience additional latency or delay. Users have adapted their behavior to circumvent this latency by using the Internet during "off peak" hours for reduced latency. The possibility of not experiencing delays has traditionally been solved by technical solutions such as a priority queuing method

1. The workshop was partially funded by the National Science Foundation, grant #NCR-9509244, and the Advanced Research Projects Agency, contract #N00174-93-C-0036. Please consult the workshop home page (http://rpcp.mit.edu/Workshops/cfp.html) for workshop notes and other workshop information.

for certain interactive applications (e.g. telnet).[2]

Statistical sharing of common facilities has allowed the infrastructure to grow quickly and robustly. Congestion problems experienced in some links were solved by increasing the bandwidth between these links. This method of "over-provisioning" has worked well in the past but may not be sufficient to improve the quality of service with more recent technical developments – namely the introduction of a new version of the Internet protocol.

Internet Protocol version 6 (IPv6) promises to offer multiple qualities of service beyond the current best-effort service. If these services are not priced differently, it is possible there may not be enough of the shared resource to service each request. Methods for rationing or pricing this service are now being considered by engineers – a possibility that has been considered by economists for quite some time.

2.2 Economics and the Internet

Perhaps two of the most commonly held misconceptions about the Internet involve its economics. The first is that it is paid for by the government when, in fact, in 1994 NSF paid for less than 10% of the Internet cloud. The second is that the Internet is free. MacKie-Mason and Varian detail the frequently asked questions concerning the economics of the Internet [20] and the economics for the infrastructure [23]. We note further that for most users, the cost of local area networks and computers far exceeds the costs which can be ascribed exclusively to Internet connectivity. Below, we will briefly analyze the economics of the Internet to provide context for the remainder of the paper as was explored during the workshop.

A common good, as defined by economists, is "[a] resource, such as air and water, to which anyone has free access."[3] Economic theory suggests that the Internet can also be included as a common good since the "cloud" is common to all Internet users. Therefore, the Internet, like the environment, may be prone to abusive behavior by one of its users leading to a degradation in that resource. This has been characterized as a "tragedy of the commons" by Garrett Hardin [16].

Economists have questioned the current practice of users experiencing latency when the Internet is congested suggesting that some users may pay more to circumvent congestion.

The economic community feels that the pricing structure of the Internet is sub-optimal. As the disparity between heavy users and lighter users of the Internet widen, more equitable and efficient pricing mechanisms will have to be created to enable self-sustaining growth of the Internet. 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. All of the cost associated with people, electronics, leased lines, etc. have been paid for already and are seen as fixed costs. The congestion cost means that when a user on the Internet sends or receives data, they are using bandwidth that could otherwise be used by other users of the network. It has been argued that usage based pricing that reflects this congestion cost would be a more equitable pricing policy [22].

Although this view is disputed across disciplines, some argue that it is only a matter of time before flat-fee pricing mechanisms fail as an expected Internet crunch approaches and congestion costs increase. 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, IPv6, etc.) and may require appropriate pricing mechanisms to emerge with these new services.

On the other hand, there is a powerful economic incentive to cooperate induced by the significant positive network externalities associated with being on the Internet. Cooperation through statistical sharing benefits each individual Internet user and has been sufficient motivation to date to avoid the worst case abuses which one could imagine.

The challenge then is not only to develop effective Internet pricing mechanisms, but to do so without losing the benefits of interoperability and positive network externalities currently gained through the Internet's reliance on statistical sharing. Achieving economic efficiency without inciting users to abandon the Internet's technical approach of statistical sharing, and hence without abandoning the Internet as such, is the challenge policy makers, businesses, and the public must now face.

2.3 Policy

The U.S. Government's role in the development of the Internet has changed with the NSFnet transition plan, but continues to be an important role. Even with the privatization of the federally-invested Internet backbone, the NSFnet, there are other federally-invested internets (ESnet, NSI, etc.).

2. The NSFnet experienced such congestion over its 56kb/s backbone in 1986 that it did just this.

3. Pyndyck and Rubinfeld, (1995) Microeconomics, Prentice-Hall, p.670.

As an investor and large user of the Internet, the U.S. Government is very concerned with the self-sustaining economic growth of the Internet. The mission agencies, such as NASA and the Department of Energy, operate their own wide area IP network which interconnects with the Internet. Therefore, the government continues to be conscientious of how it invests in them. One of the main documents that outlines this guidance is the Office of Management and Budgets (OMB) Circular A-130, Management of Federal Information Resources. This circular outlines how federal agencies should coordinate their efforts so that they may provide Internet services in a cost effective manner.

Specifically, the government is interested in how it allocates and recovers costs for Internet services. OMB Circular A-130 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 changed in July of 1994 from its original 1985 form in the area of cost accounting and recovery for federal networks [28]. While the old circular was very direct about how to account and recover for costs, the new circular is more general. This change was a result of the difficulty of implementing the recommendations from the original circular for shared data networks such as the IP networks federal agencies have. Other users of IP networks have had similar difficulties in developing accounting and cost recovery systems for IP networks which were considered fair by users.

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 [5], 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.

The Federal Networking Council and its Advisory Committee, which facilitate information exchange and action items for federal IP networks, have been involved with the issue of cost accounting and recovery.

3 Definitions

Throughout the aforementioned workshop, we found it difficult for the different constituencies to agree on the definitions surrounding Internet Economics. For example, participants couldn't agree whether or not the current Internet was characterized by a "usage-based" pricing system. It was argued that people buy leased lines based upon expected usage thereby leasing a higher bandwidth, more expensive trunk and, therefore, pay on a usage-based model. We do not expect the struggles over the definition of usage-based pricing to be resolved soon. Instead, we use the following definitions to characterize proposals to price the Internet.

1. Flat-rate pricing: Currently, most Internet users pay a fee to connect while they are not billed for each bit sent. For example, a user may pay for a T1 link *regardless* of how many bits they receive or send.

2. Usage-sensitive pricing: Users pay a portion of their Internet bill for a connection (this price could be zero, but rarely is) and a portion for each bit sent and/or received. The marginal monetary cost of sending or receiving another bit is non-zero during some time period. It is possible, for example, to have usage-sensitive pricing during peak hours and flat-rate during off-peak but we define the overall system as being usage-sensitive pricing.

3. Transaction-based pricing: Like usage-sensitive pricing, the marginal monetary cost of sending and/or receiving another bit is non-zero. However, the prices are determined by the characteristics of the transaction and not by the number of bits.

While we realize that these definitions may be different for different constituencies, these definitions are roughly agreed upon and necessary to discuss the research done in this area.

4 Internet Pricing Research

Pricing policies associated with Internet economics include pricing for congestion control and interconnection. Many of the assumptions underlying the application-specific pricing policies in telecommunications, such as for telephony, are difficult and costly to implement for the Internet. Furthermore, the Internet culture and economics are much different than that of other telecommunication networks. These differences lead to fundamental differences in the way we look at the Internet as an economic system.

While some work has been done to bring these constituencies' work in one place [3] there is need to develop a framework to think about Internet economics. David Clark opened the Internet Economics workshop by presenting a possible framework [10] and argued that the current telephony models do not make much sense for the Internet. Clark's framework emphasized that the Internet is complex and, unlike the telephone system, is a statistically shared resource.

We further argue that, unlike the telephone system, it is difficult for the Internet to support transaction cost pricing because of its distributed computing environment.

The paper critiques some more recent proposals for Internet-specific pricing and discuss why these proposals have not been implemented. We also provide some evidence where new pricing policies have been tried and provide analysis of their failures and limited successes.

For example, interconnection points with flat-fee settlements may provide improper incentives for use and may lead to inequitable cost recovery by the interconnected firms.

4.1 Statistics Collection

The growth in the number of Internet users and applications is manifested in the statistics characterizing network traffic. The term "Internet crunch" [6] coupled with the privatization of the National Science Foundation's network backbone (NSFnet) raise several economic and policy issues that were addressed at the workshop.

To better understand how the Internet is used and how potential pricing policies may affect this use, supporting research has been done to characterize Internet traffic. This work has focused on traffic flows and, as a by-product, developed new network monitoring tools [7] [9]. Currently this information is used by network managers to find bottlenecks in their networks for quick response to poor Internet service. They also use this data to understand long-term trends to form strategies to plan for network growth.

We encourage the development of statistics collection tools to better understand Internet economics. By characterizing user behavior it may be possible to model potential usage-sensitive pricing mechanisms. Also, the statistics collection tools will be able to support the billing of usage-sensitive pricing plans. Pricing is, perhaps, the most visible component of Internet economics.

4.2 Pricing and Resource Allocation

A critical aspect of this complex problem involves pricing policies and the Internet. While much work has been done in the area of telecommunications pricing [26] with regard to telephony, only recently has work on pricing data networks been undertaken. However there are motivations to pricing statistically shared networks [11] including the disparity of bandwidth used by different applications supported by the Internet. The considerable technical differences between Internet and telephony, as previously argued in this paper, prohibit the possibility of using telephony pricing mechanisms to the Internet. Kelly details the economics behind this problem and introduces the concept of "effective bandwidth" which may eventually shape Internet pricing [17] [18].

New usage-sensitive pricing systems for *some* Internet traffic have been suggested for the Internet [6] [22], but have not yet been implemented. These proposals allow the user to send and receive traffic with the current best effort service but allow some priority traffic to avoid congestion. This priority traffic is billed on a usage-sensitive basis. While, many reasons exist for and against new Internet pricing policies, there seems to be little consensus on this subject. Others critique these new systems with regard to their underlying assumptions [30] and the current practice of overprovisioning leading to expanded capacity [13]. We feel there is good reason why these new pricing policies have not taken hold in the Internet community that go beyond the economic criticisms. The other constituencies are not motivated to implement a "technical fix" to the perceived economic problems when methods such as overprovisioning are more consistent with the Internet culture. More recent models have been developed which expand the thinking of Internet pricing as a "congestible resource" [21] but many of the implementation problems still persist.

A usage-sensitive pricing mechanism for the Internet isn't entirely counter cultural for countries with high cost Internet access. In fact, usage-sensitive Internet pricing has worked in New Zealand [8]. We propose that while this is an interesting data point, their solution may not be practical in the United States where there are multiple ways to access the Internet, not one high cost link.

In a market similar to the U.S. we have already seen this happen. 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.

Newly developed mechanisms have been developed to price at the TCP layer [15], not at the IP layer like the previous proposals. While there have been no instances of this pricing mechanism actually implemented to bill users, it is a step towards transaction-based pricing. There is a security concern with collecting this data since it could be used for traffic analysis. McKeown argued at the workshop that this proposal could be used for UDP traffic as well which we feel is important given the growing use of the UDP for applications such as INETPhone and MBone.

The flat-rate pricing model does have tremendous advantages. In fact, the minimalist view of pricing B-ISDN services flat-rate is very similar to people's fondness of flat-rate pricing for the Internet [1]. The administrative overhead for the billing systems could be avoided with a flat-rate pricing system. The percent of the overhead is unknown, but perhaps we can learn from the telephony case where the billing overhead accounts for approximately 50% of the phone bill [4]. The flat-rate system is also closer to the cost of the Internet. The marginal cost of sending an additional packet on the Internet is zero in the short-run [23]. Flat-rate pricing does not accomplish the goal of providing a congestion control mechanism for better resource allocation.

It was evident during the workshop that not all traffic on the Internet needed to be treated the same and, therefore, that service qualities can affect resource allocation and pricing [31]. If, for example, a user wanted to send a guaranteed service quality packet (fixed delay and fixed bandwidth) it would have to avoid congestion in the network experienced by best effort service quality packets. At the same time, pricing based on service quality may be the only way to limit the demand on the guaranteed service quality.

Users concerns with implementation of usage-sensitive pricing may be solved with expenditure controllers to ensure that they do not incur usage-sensitive pricing greater than expected [14]. While this solution need not be implemented in today's flat-fee Internet, it may prove to be a mechanism to convince users that they cannot be incur larger-than-expected bills.

4.3 Interconnection Agreements

The driving economic forces behind interconnection are positive network externalities. The user's benefit of a network increases when they can interoperate with more users. This is clear with an application such as e-mail (since a user can correspond with more people) but is less clear when potential users may congest our network. While this problem was identified at the Internet Economics Workshop, no formal model was presented which described the differences in network externalities.

It was realized that there are many models of interconnecting two users together but they fit into a few major categories. These categories are [2]:

- Bilateral agreements – ISPs agree to interconnect for economic reasons
- Third party administrator - provides interconnection between a number of hosts
- Cooperative agreement - e.g., certain governmental agencies who interconnect their IP networks but do not act opportunistically (i.e. try and extract profits) if they operate an interconnection point.

These three models of interconnection characterize the existing technical architectures actually connecting Internet networks [33]. There are many methods of interconnection which is a reflection of the heterogeneous nature of Internet technology. Unfortunately, heterogeneity (of users, computers, networking environments) sometimes leads to non-interoperable systems.

Economic payments for interconnection, while different for the three models characterized, can either encourage or discourage resale. For example, with a flat-fee interconnection agreement, there is the possibility of two firms to aggregate their traffic before the interconnection point and split the fee of the interconnection. The Commercial Internet eXchange (CIX) used a policing policy to ensure that this did not happen. It would be possible, however, to design an economic model which was usage-sensitive so there would be no reason to police for resale. No usage-sensitive interconnection pricing examples were brought forward during the workshop.

Regulation may be likely if a third party administrator, for example, excluded some firms who wanted service. This, however, does not seem to be the case when actually regulation of the regional Bell operating companies seems to be hindering interconnection agreements and not helping them [12].

4.4 Commerce and Information Security

The last area of Internet economics explored by the Internet Economics Workshop was commerce and information security on the Internet. While there is much research work in designing architectures for billing [32] and electronic payment mechanisms [25] [27], there is less disagreement among the constituencies as compared to pricing issues. The possibility of using billing and electronic payment currently is developed for people who want to buy and sell information or good – not bandwidth. Therefore, the discussion from this section was not as controversial.

There were threads in this panel which related to the rest of the workshop. The problems of trusting the firm that provides billing or interconnection has already been addressed by the electronic commerce researchers [19]. The Internet Economics Workshop was used as more of an informational session to disseminate the findings of researchers and practitioners.

We do believe, however, that the research on electronic commerce and information security will play a major role in the future. Internet economics is only recently concerned with pricing and billing mechanisms for access. The true economic benefit of network externalities will be realized when firms may interoperate with more users and reduce transaction costs for businesses and users. The effect this will have is an area for further research.

5 Consensus and Future Directions

There is a need to continue to develop a framework for Internet economics, building upon current research in areas such as statistical analysis, pricing, interconnection agreements, and electronic commerce. These focal areas of research are critical to understanding Internet economics. It was the intent of the Internet Economics Workshop to accelerate the process of finding common ground among the diverse interests and needs of government, industry, and academic groups to probe these areas of Internet economics.

The workshop brought these communities together with the hope they could find the common ground among them (as denoted by the "*" in Figure 1). The workshop allowed these communities to establish a shared framework for the Internet as an economic system that will provide a basis for long- and short-term actions by government and industry. This framework and the areas for future research are outlined in this section.

We identify six areas of opportunity where new solutions to Internet economic problems may be explored. They are:

1. Further discussion and analysis of Internet pricing models and policies, perhaps facilitated by electronic mailing lists and directed workshops, which would include all necessary parties so that integrated approaches may be implemented.

2. The Internet Engineering Task Force (IETF) should be a forum for standardizing practices of accounting methods and security for electronic commerce on the Internet with increased emphasis on the underlying problems of Internet economics.

3. New open, interconnection agreements should be sought out, drawing from the work done in related fields.

4. New services promised by a new version of the Internet Protocol, IPv6, must have incentive policies to encourage appropriate use of these services.

5. Increased data collection between Internet service providers will provide an understanding about the growth of the Internet and provide a rich data set to develop economic models and increase our understanding of the subtle interaction between Internet engineering, economics, and cultural practices.

6. Research on the relationship between the Internet culture and the technical, economic, and policy characteristics of the Internet, including empirical observation of user behavior and attitudes under different pricing models is also needed. Perhaps an "Internet Cultural Impact Assessment" is needed before any wholesale changes in Internet pricing are widely implemented.

These recommendations, many of which were discussed at the Workshop on Internet Economics, are being further developed by workshop participants for publication [24].

6 Conclusion

The effect of the Internet culture on Internet economics is very apparent and very important to members of the technical, economic, government, and user communities.

While work has been done separately on the economic development of the Internet, the Internet Economics Workshop was the first attempt at bringing these constituencies together. While we are biased in our assessment, we have received positive feedback from the participants and observers. Publications and public discussion via distribution lists may enable these constituencies to continue to work together towards the future directions of Internet Economics.

These research and implementation endeavors include efforts in getting a better understanding of the nature of Internet usage today as well as an understanding of future pricing policies. We encourage economists and engineers to collaborate in developing integrated models to predict the effect of possible pricing policies for congestion control and interconnection.

Better understanding would lead presumably to better economic performance while safeguarding the Internet culture and further stimulating the powerful innovation engine the Internet has become for the emerging Global Information Infrastructure.

References

[1] L. Anania, and R. Solomon, "Flat – The Minimalist B-ISDN Rate," Telecommunications Policy Research Conference, Maryland, October 31, 1988.

[2] J. Bailey, "Interconnection Agreements and Internet Economics," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[3] J. Bailey, "Internet Economics," unpublished document, NASA, October 1994.

[4] J. Bailey, S. Gillett, D. Gingold, B. Leida, D. Melcher, J. Reagle, J. Roh, R. Rothstein, and G. Seale, "Internet Economics Workshop Notes," Research Program on Communications Policy, MIT, March, 30, 1995.

[5] W. Barns, "Defense Data Network Usage Accounting Enhancement Approaches," The MITRE Corporation, 1989.

[6] R. Bohn, H.-W. Braun, K. Claffy, and S. Wolff, "Mitigating the Coming Internet Crunch: Multiple Service Levels via Precedence," Applied Network Research Technical Report, San Diego Supercomputer Center, 1994.

[7] H.-W. Braun, K. Claffy, and G. Polyzos, "A Framework for Flow-Based Accounting on the Internet," Proceedings of SICON '93, 1993.

[8] J. Brownlee, "New Zealand Experiences with Network Traffic Charging," ConneXions, December, 1994.

[9] K. Claffy, H.-W. Braun, and G. Polyzos, "Internet Traffic Flow Profiling," Applied Network Research Technical Report, San Diego Supercomputer Center, 1994.

[10] D. Clark, "A Model for Pricing the Internet," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[11] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, "Pricing in Computer Networks: Motivation, Formulation, and Example," University of Southern California, 1993.

[12] B. Collet, "Telecommunications Policy: A Framework for Internet Service Providers," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[13] D. Crawford, "Pricing Network Usage: A Market for Information or a Market for Communication?" forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[14] K. Danielsen, and M. Weiss, "User Control Modes and IP Allocation," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[15] R. Edell, N. McKeown, and P. Varaiya, "Billing Users and Pricing for TCP," Department of Electrical Engineering and Computer Science, University of California, Berkeley, October 31, 1994.

[16] G. Hardin, "Tragedy of the Commons," Science, 1968.

[17] F. Kelly, "Charging and Accounting for Bursty Connections," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[18] F. Kelly, "Tariffs and Effective Bandwidths in Multiservice Networks," Proceedings of the 14th International Teletraffic Congress, The Fundamental Role of Teletraffic in the Evaluation of Telecommunications Networks, June 1994.

[19] C. Lai, G. Medvinsky, B. Neuman, "Endorsements, Licensing, and Insurance for Distributed System Services," Proceedings of the Second ACM Conference on Computer and Communications Security, November 1994.

[20] J. MacKie-Mason, and H. Varian, "Economic FAQs about the Internet," Journal of Economic Perspectives, vol. 8, no. 3, pp. 75-96, Summer 1994.

[21] J. MacKie-Mason, and H. Varian, "Pricing Congestible Network Resources," University of Michigan, 1994.

[22] J. MacKie-Mason, and H. Varian, "Pricing the Internet," University of Michigan, 1994.

[23] J. MacKie-Mason, and H. Varian, "Some Economics of the Internet," University of Michigan, 1993.

[24] L. McKnight, and J. Bailey, ed., "Internet Economics," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[25] G. Medvinsky, and B. Neuman, "NetCash: A Design for Practical Electronic Currency on the Internet," ACM Conference on Computer and Communications Security, 1993.

[26] B. Mitchell, and I. Vogelsang, Telecommunications Pricing: Theory and Practice, Cambridge University Press, New York, 1991.

[27] B. Neuman, and G. Medvinsky, "Requirements for Network Payment: The NetCheque Perspective," Proceedings of IEEE Compcon '95, March 1995.

[28] OMB, "Management of Federal Information Resources," Office of Management and Budget, Executive Office of the President, Circular A-130, July 1995.

[29] RPCP, "Workshop on Internet Economics,"

proposal to the National Science Foundation, Research Program on Communications Policy, MIT, December 1994.

[30] M. Sarkar, "An Assessment of Pricing Mechanisms for the Internet: A Regulatory Imperative," forthcoming, Journal of Electronic Publication, University of Michigan Press, http://www.press.umich.edu/jep, 1995.

[31] S. Shenker, "Service Models and Pricing Policies for an Integrated Services Internet," forthcoming, B. Kahin, and J. Keller, Public Access to the Internet, MIT Press, Cambridge, MA, 1995.

[32] M. Sirbu, and J. Tygar, "NetBill: An Internet Commerce System Optimized for Network Delivered Services," Carnegie Mellon University, 1995.

[33] P. Srinagesh, "Internet Cost Structures and Interconnection Agreements," 22nd Annual Telecommunications Policy Research Conference, Maryland, October 1994.

Author Information

Joseph P. Bailey is a doctoral candidate in the Technology, Management and Policy Program at MIT and a research assistant with the MIT Research Program on Communications Policy.

Lee W. McKnight is an associate director of the MIT Research Program on Communications Policy, a Lecturer in the MIT Technology and Policy Program, and a Principal Research Associate at the MIT Center for Technology, Policy, and Industrial Development.