Pricing in Computer Networks: Reshaping the Research Agenda

S. Shenker	D. Clark	D. Estrin	S. Herzog
Xerox PARC	MIT	USC/ISI	USC/ISI
shenker@parc.xerox.com	*ddc@lcs.mit.edu*	*estrin@usc.edu*	*herzog@isi.edu*

Abstract

As the Internet makes the transition from research testbed to commercial enterprise, the topic of pricing in computer networks has suddenly attracted great attention. Much of the discussion in the network design community and the popular press centers on the usage-based *vs.* flat pricing debate. The more academic literature has largely focused on devising optimal pricing policies; achieving optimal welfare requires charging marginal congestion costs for usage. In this paper we critique this *optimality* paradigm on three grounds: (1) marginal cost prices may not produce sufficient revenue to fully recover costs and so are perhaps of limited relevance, (2) congestion costs are inherently inaccessible to the network and so cannot reliably form the basis for pricing, and (3) there are other, more structural, goals besides optimality, and some of these goals are incompatible with the global uniformity required for optimal pricing schemes. For these reasons, we contend that the research agenda on pricing in computer network should shift away from the optimality paradigm and focus more on structural and architectural issues. Such issues include allowing local control of pricing policies, fostering interconnection, handling multicast appropriately, and allowing receivers to pay for transmission. To illustrate our point, we describe how these goals might be accomplished in the context of a different pricing paradigm: *edge* pricing. In addition, we argue that in the context of this edge pricing paradigm, usage-based pricing and flat pricing are not radically different but instead both reside along the single continuum of *usage-constraining* pricing policies.

This research was supported in part by the Advanced Research Projects Agency, monitored by Fort Huachuca under contracts DABT63-94-C-0073 (SS), DABT63-94-C-0072 (DC), and DABT63-91-C-0001 (DE,SH). The views expressed here do not reflect the position or policy of the U.S. government.

1 Introduction

In a few short years, the Internet has made a dramatic transformation from nerdy enigma to trendy hangout. With its millions of users and diverse application offerings, the Internet is now seen by many pundits as the archetype of the future global information infrastructure. Because of its heavily subsidized origins, commercialization has come late to the Internet. As the Internet confronts this belated and somewhat awkward transition from research testbed to commercial enterprise, there has been much recent discussion about the role of pricing in computer networks. Numerous workshops and conferences have been held on the topic in both the academic community and the network design community; the popular press has also seized upon the issue as one of broad interest (see, for instance, [44, 31]).

In the popular press and in the network design community, the agenda has been dominated by debates over whether to move from the present system of charges based on the speed of the access line (so-called "flat pricing") to basing charges on actual usage. Some contend that usage-based pricing is unnecessary, and would have disastrous consequences for the Internet. Others argue that moving away from flat pricing towards usage-based pricing is essential for the Internet's efficiency, and therefore is the key to its future economic viability. Unfortunately, little has been clarified by this heated debate except the depth of the participants' convictions. We hope to demonstrate in this paper that usage-based charging and flat-pricing are really two ends of a single continuum, so the difference between them is not one of fundamental principle but merely of degree, and that hybrids of the two approaches will likely be commonly used in the future.

The academic discussion of pricing in computer networks has concentrated on a rather different issue. This literature typically assumes the necessity of usage-based pricing and focuses on achieving optimal efficiency – maximal welfare – in certain simplified models using usage-based pricing schemes. The satisfaction a network user derives from her network access depends on the nature of the application being used and the quality of service received from the network (in terms of bandwidth, delay, packet drops, etc.); the network's resources are used most efficiently if they maximize the total user satisfaction of the user community. To achieve optimal efficiency, usage-based charges must equal the marginal cost of usage. Since the physical transmission of packets is essentially free, the marginal usage cost is almost exclusively a congestion cost; congestion costs are the performance penalties that one user's traffic imposes on other users. This *optimality* paradigm dominates the research agenda; much of the literature discusses pricing schemes based on computations of these marginal congestion costs.

The main purpose of this paper is to advocate shifting the research agenda away from the reigning optimality paradigm and towards a more architectural focus. We will use the phrase *pricing architecture* to refer to those components of the pricing scheme that are independent of the particular local pricing decisions and reflect nonlocal concerns, such as how receivers rather than senders can be charged for usage and how to appropriately charge multicast transmissions. These architectural issues, rather than the detailed calculation of marginal congestion costs, should form the core of the research agenda. To motivate this shift in research emphasis, we discuss both the economic issues and also the mechanistic design issues central to computer network pricing. Our treatment of these issues is designed to be accessible to both the network design community and the economic community, with the intention of providing some common context for these two communities and thereby increasing the opportunity for dialogue.

Our paper has three distinct parts. The first part, in Section 2, critiques the optimality paradigm.[1] We first contend that usage charges may, and perhaps should, exceed marginal congestion costs. Moreover, we argue that these marginal costs are inherently inaccessible and so the quixotic pursuit of their precise computation should not dominate the research agenda. The second part, in Section 3, presents a rather different paradigm for network pricing: *edge* pricing. This term refers to where the charges are assessed rather than their form (*e.g.*, usage-based or not) or their relationship to congestion (*e.g.*, marginal congestion costs or not). This emphasis reflects our belief that architectural issues are more important than the detailed nature of the charges themselves. The third portion of the note, in Section 4, describes two fundamental architectural issues and some preliminary design approaches. We conclude in Section 5 with a brief summary. Because much of our discussion requires some familiarity with network mechanisms, in Appendix A we present an extremely short overview of the relevant material.

2 A Critique of the Optimality Paradigm

The optimality paradigm may have particular relevance for isolated settings in which the network provider's goal is to maximize welfare, such as in a nonprofit research network or an internal corporate network. In this paper, however, we are addressing the role of pricing in a commercially competitive environment. The current Internet service provision market has multiple independent service providers (ISPs), and competition appears to be increasing rapidly. We claim that the optimality paradigm is not an adequate foundation for pricing in such a competitive setting.

The optimality paradigm places a special focus on marginal congestion costs. Our critique is posed in the form of three questions:

- Are marginal congestion costs relevant?
- Are marginal congestion costs accessible?
- Is optimality the only goal?

2.1 Are Marginal Congestion Costs Relevant?

It is a standard result that the overall welfare (the sum of provider profit and consumer surplus) is only maximized when prices are set equal to marginal cost, where these marginal costs take into account all externalities. In computer networks, these externalities include both congestion effects, where one user's use imposes a performance penalty on other users, and also connectivity effects, where a user benefits from other users being connected to the network.

[1]The authors include themselves in this critique, having adhered to this optimality paradigm in previous publications, such as [9, 8].

Competition between network service providers will typically drive prices to these marginal costs. If the marginal cost prices are sufficient to recover the *facility* costs of building and operating the network infrastructure, then these marginal cost prices are a stable competitive equilibrium[2], and so computing marginal congestion costs would be central to network pricing schemes. But it is doubtful that such marginal cost prices will recover the full facility costs of computer networks. Within the context of the congestible resource model studied in [29], marginal costs only cover the cost of the facilities priced at marginal expansion cost (*i.e.*, the total congestion costs are equal to the product of the total capacity times the marginal cost of capacity). If the facility costs are a sublinear function of capacity (*i.e.*, $f(x) > f'(x)x$ for $x > 0$) then facility costs will not be fully recovered by marginal cost pricing (see also [43, 33]). While the cost structure of networks is in flux as technologies rapidly evolve, it seems clear a large portion of the facility costs arises from the fixed (*i.e.*, not related to capacity) costs of deploying the physical infrastructure. Consequently, we assume in this paper that while marginal congestion costs may be nontrivial, they will be much less than the total facility cost of providing network service.

In such cases, there is no stable competitive equilibrium (see [17, 41] for a more thorough discussion of this point); any stable situation must have some prices that exceed the associated marginal costs. What guides the setting of prices in such a situation? While there are few general results applicable here, one could argue that the resulting prices of each firm will satisfy the Ramsey condition of maximizing the consumer surplus while still fully recovering costs (because otherwise competing firms would enter and lure customers away by offering more surplus). Thus, in raising prices to increase additional revenue, network service providers will do so in a manner that retains, to the greatest extent possible, the maximality of welfare (since maximizing welfare at a fixed level of profit is equivalent to maximizing consumer surplus at a fixed level of profit).

It is useful, in the following discussion, to artificially break the pricing structure into two distinct pieces.[3] One component of network charges is the attachment fee; this is the fee charged for gaining access to the network and is independent of any actual or potential usage. The other component is what we will call a usage-constraining fee. There are marginal costs associated with both attachment and usage, and welfare is optimized when they are set equal to their respective marginal costs. If all users derived significant benefit from their network connection, the Ramsey pricing scheme would be to raise attachment fees but keep usage fees at the marginal cost levels, thereby retaining the optimal usage behavior and merely recouping additional revenue from attachment. This is the argument most commonly used to motivate the continued use of marginal congestion pricing in cases where marginal prices by themselves do not fully cover costs.[4]

---

[2] We use the term stable only to mean that the revenues cover costs; we do not use the term to refer to any other dynamical properties of the equilibrium.

[3] We ignore nonlinear pricing policies here in order to simplify the discussion.

[4] We should make clear that we are assuming that network providers are not also controlling, or directly profiting from, the content delivered over their networks. However, in *bundled* networks such as cable TV, where the application and the network transport are sold as a single unit, there are many more opportunities to recover costs. Profits on content and revenue from advertising [10] are important aspects of pricing in bundled networks. We do not consider such bundled networks in this paper, and restrict ourselves to the analysis of pricing pure Internet access without bundled services. The nature of the ISP market is still very much in flux, and there may be other sources of revenue in the future, such as renting space on provider-

Unfortunately, the assumption of uniformly large benefit from network access does not appear to apply to current computer networks. The low rates of penetration of Internet connectivity, and the high rate of churn in subscriptions to online services such as AOL and CompuServe, suggests that in addition to the many users that derive great value from their network connection, there are probably also many other users whose valuation of network connectivity is marginal, and who would disconnect if attachment fees were raised.[5] Thus, we expect that both usage and attachment prices will affect welfare, and there will be a unique price point that produces a positive optimal welfare. Assuming smoothness throughout, deviating from the optimal pricing point produces welfare changes that are second-order in the price deviations. The matrix of second derivatives will depend in detail on the individual utility functions and there is little reason to expect that, in general, consumer surplus is maximized when only attachment fees are raised.

In addition, the Ramsey pricing scheme could be different for different subpopulations of users. For instance, low-volume users who derive very little benefit from being connected to the network would more likely absorb an increase in usage charges without detaching from the network. This is consistent with what we observe; some commercial Internet providers charge based on volume to attract low-volume, marginal benefit users who might not otherwise purchase access. In contrast, most large institutions, which typically derive great value from their network connection, pay substantially for the attachment. The traditional and cellular telephony markets also display extensive second-degree price discrimination (i.e., nonlinear pricing schemes where the per-unit price depends on the quantity purchased); there are many different pricing plans, some with lower attachment charges and higher usage charges, and others with the reverse. We expect a similar use of second-degree price discrimination to increase revenue in computer networks.

There are other considerations that suggest that usage charges must remain at significant levels, even if congestion is extremely low (and so marginal congestion costs are extremely low). Assume that an entering network service provider can steal away a subpopulation of users from their current service provider if the entrant can supply this subpopulation with sufficient bandwidth to satisfy their needs at a cost less than the total fee being charged by the current provider.[6] Then we must impose a "core" condition on the pricing structure, mandating that no subset of users can be charged more than the cost of providing that subset service. If one believes that bandwidth is responsible for any significant portion of the cost of networks, then usage charges must be used to satisfy this core condition. Usage charges are needed to price discriminate between low-volume and high-volume users; otherwise, a competing network provider would steal all the low volume users away by offering a network provisioned at much lower levels with much lower prices. Thus, this core criterion requires

---

supplied web-servers, that may complicate the rather simplified case we are analyzing.

[5]It is possible that, in the grand and glorious future, the GII will have a single Internet-like network infrastructure and all households have a single network connection that carries their telephony, television, and data traffic. At that point, it may well be true that essentially all users have high valuation of their network connection and raising attachment fees is the appropriate way to raise revenue. However, we are a long way from this utopian vision, and we should design our current network pricing policies to fit the present situation.

[6]This assumes seamless interconnection, so switching providers does not affect connectivity. Otherwise, the decision to switch providers involves many other factors besides cost.

that users who regularly consume (or who plan to consume) significantly less bandwidth be charged less, with the difference reflecting the percentage of cost due to bandwidth. Of course, if bandwidth is relatively cheap (*i.e.*, is a very minor portion of the network cost) then this "core" argument has little bite.

2.2 Are Marginal Congestion Costs Accessible?

When prices are required to fully recover costs, we think there is little reason to expect usage prices will equal the marginal congestion costs. We now put that conclusion aside and ask: if we nevertheless attempted to set prices to these marginal congestion costs, could we actually do so? It turns out that computing these congestion costs is quite difficult.

The relationship between what happens to a packet traversing a network and the resulting change in a user's utility is extremely complicated. When we look at the fate of a single packet, congestion can cause it to be delayed or dropped. Some applications are very sensitive to this extra delay (or being dropped), and others are not. Pricing schemes seeking to achieve optimal efficiency must take these different delay and drop sensitivities into account. While in simple theoretical models it is convenient to use the abstraction that a user's utility is a function of, say, average bandwidth and delay (as in [38]), the real world is significantly more complicated; see the discussion in [7] about the properties of best-effort traffic. Unfortunately, we have little beyond these simple theoretical models to guide us.

Moreover, most applications involve a sequence of packets, and the effect on utility due to the dropping or delay on one individual packet depends on the treatment given the rest of the packets. For instance, the performance of a file transfer depends on the time the last packet is delivered; for large files this transfer time depends almost exclusively on the throughput rate and not on the individual packet delays (see [7] for a more thorough discussion of this point). It is extremely difficult, if not impossible, for the network to infer the effect on the transfer time arising from delaying any of the individual packets, especially since the transfer time is also a function of the user's congestion control algorithm. To make matters even worse, often an entire suite of applications is used simultaneously, and then the user's utility depends on the relationship between the delays of the various traffic streams (*e.g.*, a teleconference may involve an audio tool, a video tool, and a shared drawing tool).

Our understanding of this relationship between handling of individual packets and the overall utility is rather primitive, and the relationship changes rapidly with technology (*e.g.*, advances in congestion control could greatly decrease the sensitivity to randomly dropped packets). An important aspect of the problem is that the Internet architecture is based on the network layer not knowing the properties of the applications implemented above it. If we believe that network service providers will sell raw IP connectivity (*i.e.*, they just provide access at the IP level, and do not interpose any application-level gateways), then they have to price based solely on the information available at the IP level, and this greatly restricts the extent to which they can adjust prices to fit the particular applications being used. See [28] for a discussion of the implications of this layering for content provision.

There have been many pricing proposals in the recent literature and we do not attempt to review them all here; see [2, 3, 14, 25, 24, 23, 21, 27, 32, 34, 35, 37, 36, 42, 18, 45] for a few representative samples. The most ambitious pricing proposal for best-effort traffic is the "smart-market" proposal of MacKie-Mason and Varian described in [29] (see also

([27, 26]). In this scheme, each packet carries a "bid" in the packet header; packets are given service at each router if their bids exceed some threshold, and each served packet is charged this threshold price regardless of the packet's bid. This threshold is chosen to be a market clearing price, ensuring the network is fully utilized. The threshold price can be thought of as the highest rejected bid; having the packets pay this price is akin to having them pay the congestion cost of denying service to the rejected packet. The key to this proposal is incentive compatibility; users will put their true valuation in the packet since, as in standard second-price auctions, it only affects whether they get service but not how much they pay. By putting their true valuation of service in the packet header, users will get service if and only if it costs them less than their valuation of the service.[7]

This proposal has stimulated much discussion and has significantly increased the Internet community's understanding of economic mechanisms in networks. However, there are several problems with this proposal that prevent it from achieving true optimality. First, the most fundamental problem is that submitting a losing bid will typically lead to some unknown amount of delay (since the packet will be retransmitted at a later time), rather than truly not ever receiving service, so the "bid" must reflect how much utility loss this delay would produce rather than the valuation of service itself; thus, accurate bids cannot be submitted without precisely knowing the delay associated with each bid level, and neither the network nor the user knows this delay. Second, there are complications when the packet traverses several hops on its way to its destination. The valuation is an end-to-end quantity (the user only cares about the packet reaching its final destination and does not care about any partial progress), yet the valuation is used on a hop-by-hop manner to determine access at each hop; one would have to extend the bidding mechanism to evaluate the entire path at once, and this entails a distributed multiple good auction of daunting complexity.[8] Third, the bid is on a per-packet basis, yet many applications involve sequences of packets. It is impossible to independently set the valuation of a single packet in a file transfer, when the true valuation is for the set of packets.

Wang *et al.* [45] have proposed a pricing scheme for flows making network reservations (*i.e.*, asking for a quality of service that entails admission control and some assured service level) where prices optimize a given objective function. Gupta *et al.* [18] adopt a similar approach for a best-effort network with priorities. As in any conventional economic setting, the optimality of the pricing scheme depends on knowing the demand function. In settings where the supply and delivery are not time critical, such demand functions can be estimated over long periods of time. However, in computer networks, a user's utility depends on the delay in meeting her service request, and so one cannot merely consider the long term average demand but must also respond to instantaneous fluctuations when setting prices. In addition, the problem of denial of service leading to some delay, rather than an eternal denial of service, makes the valuations of the flows not directly related to congestion costs. Consequently, determining optimality in the presence of fluctuating demand is extremely difficult.

[7]As an aside, note that the pricing scheme is embedded within the architecture in this proposal. The bids are translated into the prices charged.

[8]If one believes that the major source of congestion is at the edge of the network, then one could only apply the smart market at the edge points. This removes the end-to-end versus per-hop problem, and could be used in our edge pricing scheme as the method of charging. See the discussion in Section 3.3.

We contend that the failure of these mechanisms to achieve true optimality is not a failure of imagination, but rather evidence that the task is beyond the scope of any practical algorithm. The keys to efficiency – knowing the service degradation that will result from a particular network action (*i.e.*, how much delay and/or loss), and knowing the user's utility loss as a result of this service degradation – are fundamentally unknowable.

This is not to imply that usage pricing schemes are of little utility. When compared to a situation with no usage-constraining charges, usage charges greatly increase the efficiency of the network. Simulations and calculations in [27, 19, 30] have clearly demonstrated the significant advantages usage pricing has over free entry. Our point is merely that such pricing schemes do not achieve true optimality, and that the significant efficiency gains demonstrated could probably also be achieved with explicitly suboptimal schemes.

2.3 Is Optimality the Only Goal?

We argued in the previous section that marginal congestion costs are inherently inaccessible. This critique applies equally to attempts to compute the optimal Ramsey prices. However, since price deviations away from these optimal points typically produce only second-order deviations in the total welfare, perhaps such deviations are not of much concern. Moreover, in the pursuit of optimality in simplified models, some more basic structural issues have been somewhat neglected. In this section we identify some of these structural issues and urge that they be given significant attention in the design of pricing policies.

Pricing policies should be compatible with the structure of modern networking applications. One of the recent developments in the Internet is the increasingly widespread use of multicast, in which a packet is delivered to a set of receivers, rather than just a single receiver. By sending packets down a distribution tree, and replicating packets only at the tree's branch points, multicast greatly reduces the load on the network. Therefore, it is crucial that pricing give the proper incentives to use multicast where appropriate.

Another important aspect of network applications is that the benefit of network usage sometimes lies with the sender of the traffic, and sometimes with the receiver(s). Pricing mechanisms should be flexible enough to allow the charges to be assessed to either, or some combination of both, endpoints. This is a very important goal in computer networks; the ability to charge receivers would facilitate the free and unfettered dissemination of information in the Internet, since the providers of such information would not have to pay the cost of transport. Note that this goal is not achieved by the flat pricing approach; currently the source's access charge is paid for exclusively by the host institution. This has not yet caused a problem on the Internet, since the elastic and adaptable data applications can easily adjust to overloaded conditions. However, when real-time applications, and other applications that adapt less well to congestion, are in widespread use the pinch at the source's access point will be felt more acutely.[9]

[9]The ability to assign charges to the receiving end could, in some cases, be handled by a higher level protocol that redistributes the basic charges determined by the network. However, there are several disadvantages to requiring such a higher level protocol: it requires the ability to transfer funds at a higher level, it cannot deal with capacity-based charging (described in Section 3.3), and in the multicast case the required information (such as the membership of the group and the network topology) may not be available at the higher layer. Thus, we think it preferable to build the flexibility of assignment into the basic charging

Pricing policies should also be compatible with the structure of the network service market. There are numerous independent service providers, and many of these are small providers who merely resell connections into bigger provider networks. The interconnection arrangements between providers are somewhat *ad hoc* (see [17, 41]) and changing rapidly. Interconnection among these networks is crucial for maximizing social welfare. Pricing schemes should not hinder interconnection by requiring detailed agreement on pricing policies and complicated per-flow transfers (i.e., a separate transfer for each flow) of money when carrying traffic from another interconnected network. In addition, these independent service providers should be able to make local decisions about the appropriate pricing policies. This implies that the pricing policy should not be embedded into the network architecture. Instead, the network architecture should provide a flexible accounting infrastructure that can support a wide variety of locally implemented pricing schemes. For instance, there are some contexts (such as managing an internal corporate or university network) where the goal of pricing is merely to encourage efficient use of the network resources. Often in these contexts there are incentives that can be used (*e.g.*, quotas) instead of money. While in this paper we have focused on monetary incentives, the underlying accounting structure and pricing architecture should allow the use of these other incentive forms if they are locally applicable.

Note that achieving optimality necessarily involves uniform implementation of a single pricing scheme across the network; optimality involves setting prices at exactly the marginal congestion costs and so the accounting scheme becomes a distributed computation of those congestion costs. Thus, the optimality paradigm is fundamentally inconsistent with the need for locality in pricing. Given that no pricing scheme has claim to being truly optimal, the need for local control should take precedence over the desire for absolute optimality.

While true optimality is not an appropriate goal, pricing should still be used to achieve reasonable levels of efficiency. It is important that the underlying accounting infrastructure allow prices to be based on some approximation of congestion costs. There is an important distinction lurking here. It is important to allow prices to be *based on* some approximation of congestion costs, but it is important to not force them to be *equal* to these congestion costs. As we argued, the need for full cost recovery militates against such an assumption of equality. Meeting any reasonable efficiency goal, however, would likely require that prices depend on such congestion costs.

Rather than start with mechanisms designed to precisely calculate marginal congestion costs, we might first ask: what are the absolutely minimal requirements for providing some estimate of congestion costs? One minimal requirement is that pricing should encourage the appropriate use of quality of service (QoS) signals (by this we mean the signals sent by applications to the network requesting a particular quality of service; see Appendix A). This is crucial for making the new QoS-rich network designs effective, and would enable them to achieve significant increases in network efficiency. An additional requirement is that pricing should discourage network usage during times of congestion, but not discourage it during relatively uncongested times. Our basic point is that perhaps these minimal requirements are sufficient to achieve reasonable approximations, and that attempts to more accurately calculate Ramsey prices are of little (indeed second-order) value and distract us from the more important but often overlooked structural concerns.

---

mechanism itself.

3 A New Pricing Paradigm: Edge Pricing

After having critiqued the reigning optimality paradigm, we now present a very different pricing paradigm: edge pricing. We motivate the edge pricing paradigm by describing a series of approximations to true congestion costs.

3.1 Approximating Congestion Costs

Computing the true congestion costs requires that you can compute other users' loss in utility due to one user's use. This requires knowledge not only of the utility of users, which in the Internet architecture is fundamentally unknowable, but also knowledge of the current congestion conditions along the entire path. Such detailed knowledge entails a sophisticated accounting scheme that transcends administrative boundaries by following the entire path. Having already concluded that our estimates of utility loss are extremely rough estimates, can we also replace the knowledge of current congestion conditions along the entire path with a reasonable, but more easily accessible, estimate? Consider the following two approximations.

The first approximation is to replace the current congestion conditions by the *expected* congestion conditions. This is essentially QoS-sensitive time-of-day pricing. The time-of-day dependence builds in expectations about the current congestion conditions. The QoS dependence reflects the fact that the effect one flow's packets have on another flow's packets depends on the respective service classes of the flows; packets in higher quality service classes impose more delay on other packets than do packets in lower quality classes. This approximation of QoS-sensitive time-of-day pricing has the problem that it does not reflect any instantaneous fluctuations in traffic levels; packets sent during a lull in the network would still be charged full price even though the actual congestion costs were quite small. Such insensitivity to instantaneous conditions would seem to remove any incentive for users to redistribute their load dynamically; just as in the telephone network, time-of-day pricing encourages users to time-shift their calls to later (or earlier) hours when rates are lower, but does not encourage them to adjust to the instantaneous conditions. (Of course, in the telephone network there is no way for users to detect the current load.)

We claim that the inability to charge less during periods of low congestion is not a serious problem because, in many cases, one can substitute the congestion-sensitivity of service for the congestion-sensitivity of prices. During a lull in the network, lower quality classes give as good service as high quality classes do during congested periods. Users who monitor the service they are getting from the network and adjust their service request accordingly can take advantage of this variability. The way user costs are lowered during times of reduced network load is not that the network lowers the price of service classes but that users request lower service classes and are charged the lower price of that class.

We refer to "users" as the entities adapting to current conditions, to distinguish this from the network adapting, but we should note that in reality adaptation does not require significant effort from the human user (see [26] for a similar discussion of the role of adaptation). Instead, adaptation routines will be highly automated and embedded within applications or the end system's operating system. Many current network applications are already designed to adapt to network conditions, and so relying on users to adapt to current conditions, rather than the network, is quite consistent with current practice. In fact, this reflects a basic Internet design philosophy; to the extent possible (and routing is the one place where it is frequently less possible), the intelligence and responsibility to adapt to current conditions should be placed on the outside of the network; the fundamental infrastructure inside the network should remain fairly simple, intentionally ignorant of the applications it is supporting, and should not try to adapt on behalf of these application. Applied to this case, this philosophy argues for relatively static pricing policies with end users varying their service requests in response to current congestion conditions. This removes from the network the responsibility of accurately assessing current conditions and their likely impact on users' utilities, and puts the onus on individual applications/users to make that assessment for themselves; given that applications have very different sensitivities to service quality, it seems preferable to place the bulk of the variability where it can be done in the most informed way.

If expected congestion were the only approximation, then we would essentially have a pricing scheme where prices were computed per-link based on the time-of-day and quality of service requested. The second approximation is to replace the cost of the actual path with the cost of the *expected* path, where the charge depends only on the source and destination(s) of the flow and not on the particular route taken by the flow. From a user's perspective, they have requested service from one point to another (at least in the unicast case); the actual path the data takes is typically determined by the network routing algorithms (except in the case of source routing). Having the price of the service depend on the network's decision about routing seems an unnecessary source of price variation that makes it harder for the user to make informed plans about network use. Moreover, when alternate paths are taken by the network in response to congestion, the extra cost due to the congestion should not necessarily fall only on those flows that have been redirected. Certainly in the telephone network, the price of a telephone call does not depend on the network's choice of route.

3.2 Edge Pricing

When we combine these two approximations, the price is based on the expected congestion along the expected path appropriate for the packet's source and destination. Therefore, the resulting prices can determined and charges assessed locally at the access point (*i.e.*, the edge of the provider's network where the user's packet enter), rather than computed in a distributed fashion along the entire path. We will call this local scheme *edge* pricing. A similar approach to pricing in computer networks has been suggested by Jacobson [22]. The prices charged at the edge, or access, point may depend on information obtained from other parts of the network, but the entire computation of charges is performed at the access point. In Section 4 we discuss the multicast case where the relevant information is difficult to obtain.

As discussed in [7], edge pricing has the attractive property that all pricing is done locally. Interconnection here involves the network providers purchasing service from each other in the same manner that regular users purchase service. When a user connected to provider A's network sends a packet, it is applied to that user's bill according to whatever pricing policy provider A has.[10] If the destination of the packet is on provider B's network, then

[10] We use the term "bill" here only to connote that the packet is applied to the contract the user has with provider A; as we mention below, the contracted pricing policy may very will be a flat price with a limit on

when the packet enters provider B's network the packet is charged against provider's A bill with provider B. There are no per-flow settlement payments, in the sense that the various providers do not redistribute the charge levied to the end user among themselves. Instead, each provider takes full responsibility for every packet they forward; a sequence of bilateral agreements between the adjacent service providers along the path performs the necessary function of cost shifting. These bilateral agreements apply only to the aggregate usage by these providers, and so greatly simplify the transfer of payments between providers.

The beauty of this is that billing structures are completely local. The exact nature of the pricing scheme is simply a matter between the user and the service provider. Because the decisions are local, service providers can invent ever more attractive (and complicated) pricing schemes and can respond to user requirements in a completely flexible fashion. No uniform pricing standards need be developed since interconnection involves only bilateral agreements that allow each provider to use their own pricing policy. Locality allows providers to experiment with new pricing policies and gradually evolve them over time; in fact, pricing policies will likely be one of the important competitive advantages available to providers when competing with each other. For instance, locality allows providers to offer specialized pricing deals such as bulk discounts. It is hard to imagine implementing a meaningful bulk discount when charging is done in a nonlocal per-link basis; a user's usage of any particular link, or of any particular service provider outside of the local one, is probably quite limited, and so such discounts are much less meaningful.

3.3 Forms of Pricing

Edge pricing describes the place at which charges are assessed, but is completely neutral about the nature of these charges. In most of the literature, there is a sharp distinction between usage and attachment charges; this differentiates the fixed (or flat) portion of the price and the variable usage-dependent portion of the price. Thus, the cost of upgrading the speed of a user's access line is considered an attachment charge. We think this division is somewhat misleading, since there is a natural continuum between the two.[11] We instead choose to refer to them all as usage-constraining prices.[12] Per-packet charges are clearly designed to constrain usage, but so are limits on a user's peak sending rate.

The continuum of usage-constraining charges can perhaps best be explored by defining its two endpoints. At one end of the continuum, prices can be based on actual usage, in the form of per-packet and/or per-reservation charges; this is the traditional form of usage-based pricing. At the other end of the spectrum, users could purchase a capacity from the network and then be allowed to use, without any additional charge, up to that capacity. One form

peak rate, in which case there is no additional charge per packet.

[11]The distinction between fixed and variable prices may be extremely important to individual users; users on fixed budgets may need fixed prices, whereas users with extremely variable demand may need the ability to only pay for usage. Our point is that this distinction, while important to individual users, is not fundamentally important from an architectural or economic perspective. Both forms of pricing can be assessed locally, and both constrain usage.

[12]In our taxonomy, attachment prices would refer only to the price of attaching to the network and not refer at all to the speed of the access line. All other charges would be considered usage-constraining. Of course, in nonlinear pricing schemes (or when there is a spectrum of pricing menus offered, as in the current cellular telephony market) the distinction between the two is completely blurred.

of capacity could be defined in terms of just a peak rate, as in the current form of flat pricing. More generally, however, this capacity is defined in terms of a *filter* that is applied to the traffic. A usage filter characterizes flows as either conforming or not-conforming to the agreed upon capacity. Such filters can measure the usage over differing time horizons, such as controlling the long-term average rate, the short-term peak rate, and intermediate burst durations. This capacity framework is merely a generalized version of the current flat-rate pricing schemes; the extra flexibility allows pricing schemes to be more closely attuned to user requirements. See Appendix C for a more complete explanation of such filters. While not essential to our discussion here, we should note that there can be several possible actions that the service provider could take when a user exceeds her capacity; for instance, all such packets could be mapped into the lowest service class, or dropped, or queued until the flow is in compliance with the filter, or merely assessed an additional per-packet fee.

The units of usage that are applied against the capacity constraints, just like per-packet charges, can depend on many things such as time-of-day, destination, and QoS. High quality service classes might consume twice as many units as lower quality service classes, with similar increments for packets traveling further or over particularly congested links. Of course, to realize the goal of allowing users to send an unlimited amount of traffic when the network is empty, there should be a category of absolutely lowest quality of service that is essentially free. In fact, one could even use a "smart market" auction approach to pricing at the access point.[13]

These capacity constraints allow network providers to make informed provisioning decisions. Of course, provisioning decisions will also be heavily based on measurements of actual aggregate usage, but the capacity filter parameters give some additional input for estimates. If there is an infinite amount of multiplexing (*i.e.*, each user constitutes an infinitesimal share of the aggregate usage) and users are uncorrelated, then provisioning need only be based on the long-term average rates. The other capacity filter parameters are needed to make estimates of the magnitude of usage fluctuations away from this average value.

Because the overlimit behavior (when usage exceeds the capacity) can merely be an additional per-packet charge, there can be a continuum of pricing policies that stretch between purely usage-based charging and purely capacity-based charging. Within the spectrum of edge charging, the difference between capacity-based prices and usage-based prices is not a fundamental architectural issue. We expect that the market will invent, over time, increasingly attractive and flexible hybrids of these approaches. Telephony may provide an instructive example. Telephone companies offer a menu of local calling plans, some usage-based (*e.g.*, metered service), some capacity-based (*e.g.*, unlimited service), and some a combination of both (*e.g.*, a certain number of free minutes per month, plus a metered rate for calls in excess of this number). It is likely that the same will happen in computer networks, with some users choosing usage-based and others choosing capacity-based charges, and many being somewhere in between. Thus, the heated debate between advocates of usage-based and capacity-based pricing schemes will become completely irrelevant as users vote with their feet. Because in the edge pricing paradigm the decision between usage-based

[13]There may be some disadvantages with using the smart-market as the local pricing scheme (*e.g.*, it embeds the pricing policy in the architecture), but our point here is that it is not architecturally precluded by the edge pricing paradigm, and so firms are free to experiment with it.

and capacity-based, or anywhere in between, is completely local, and we expect that network provision will be competitive, the offered plans will likely reflect the true needs of consumers (and thus the architecture need not preclude one choice or the other to prevent providers from exploiting users).

The rest of this paper is devoted to exploring the infrastructure needed to support this edge pricing approach.

4 Architectural Issues

Edge pricing localizes the whole charging process; everything occurs at the access point. Yet, there are two inherently nonlocal aspects of pricing: (1) charging appropriately for multicast, and (2) the ability to charge receivers for the service.[14] These nonlocal aspects pose some fundamental architectural challenges to the edge pricing approach. We see these issues as forming the basis of a fertile research agenda in pricing in computer networks. In this section we discuss how the infrastructure might be designed to handle these nonlocal aspects. We describe the problems of multicast and charging receivers separately, and then we review some remaining open problems.

It is important to note that this design discussion is extremely preliminary, and is intended to be illustrative rather than definitive. That is, our purpose is to illustrate some of the issues involved by engaging in a design discussion, but we freely admit that the design directions advocated here may not, in the end, be the appropriate choices.[15]

4.1 Multicast

When unicast packets enter a provider's access point, the destination field is enough to determine the typical path of the packet. Unicast routes fluctuate occasionally, but the normal case is that unicast routes change on rather slow time scales. Thus, fairly static tables at the entry points can provide adequate information for pricing decisions, and it would be relatively trivial to design the distributed algorithms needed to construct and maintain these tables.[16] If addresses encode geographic information (as in Deering's recent proposal [11]) or

[14] Charging receivers for service is a nonlocal problem because, in approaches with explicit willingness-to-pay signaling, when both the source and receiver are serviced by the same provider the source's access point must be informed that the receiver is willing to assume responsibility for the transmission. Similarly, when the path from source to receiver traverses several different provider networks, the notification of receiver-paying must be communicated to both the exit access point and the entrance access point in each network. Other approaches can avoid this explicit nonlocal signaling by adopting some uniform standards, such as a certain portion of the multicast address space being set aside for receiver-pay groups, but these standards themselves are nonlocal in that they represent agreements between providers about a billing policy.

[15] In particular, there is a spectrum of design choices providing different levels of functionality and requiring more or less additional mechanism; in the following discussion we are not attempting to make a detailed evaluation of the functionality *vs.* mechanism tradeoff, but are merely illustrating some possible ways of achieving the aforementioned goals. There are more minimalist approaches to these problems that require less additional mechanism, and they should be considered when making design decisions for the Internet, but for this pedagogical discussion we have presented more straightforward, if more mechanistic, approaches.

[16] These tables would contain information describing how many usage units (for a capacity filter) each packet represents, or a monetary per-packet charge, or whatever other information is needed for the provider

provider information (as in the current IPv6 proposal [12]) then these tables are especially simple (see [15] for more information). Moreover, if the provider networks are small enough, one set fee for all intraprovider packets and another fixed fee for all interprovider packets might be sufficient.

Multicast packets pose more of a challenge. A multicast address is merely a logical name, and by itself conveys no geographic or provider information. While multicast routing identifies the next hop along the path for packets arriving at an interface, multicast routing does not identify the rest of the tree. Thus, estimating costs in the multicast case requires an additional piece of accounting infrastructure. Moreover, the set of receivers – the members of the multicast group – can change quite rapidly and so the mechanisms for providing the appropriate accounting information must be designed with care.

One can imagine several different approaches. The simplest would be to merely collect the location (*i.e.*, subnet numbers) of all receivers (with receivers outside of the provider's network being recorded as residing at the appropriate exit point of the network). From these locations one could compute the approximate costs of the appropriate tree.

Another approach would be to compute these costs on-the-fly by introducing a new form of control message – an accounting message – that would be initiated when the receiver sent its multicast join message (multicast *join* messages are the control messages sent by a receiver to join the multicast group; see Appendix A). These accounting messages would be forwarded along the reverse trees towards each source, recording the "cost" of each link it traversed and summing costs when branches merged. When these accounting messages reached a source's access point, the cumulative cost of reaching all receivers from that source would be available. Each provider would only need to record the cost information local to their network; that is, the costs would start accumulating when the accounting message entered the provider's network and would stop when the accounting message exited the network. No cost information crosses the provider boundaries; instead, this cost information is only used locally to compute the charges to apply on the edge of the network. This on-the-fly approach makes the charge for multicast depend on the true path rather than the typical path, which may cause unnecessary variability.[17] We should note that there might be groups (*e.g.*, cable TV channels) where the typical tree might be well enough known in advance so that such additional mechanisms are not needed; we discuss this briefly in Section 4.3.

Note that the additional piece of accounting infrastructure needed to compute these costs is local to the provider; that is, each provider can use its own algorithm. No standards need be established, no agreements with other providers need be made. Thus this protocol can incrementally evolve over time as we understand better the cost structure and traffic patterns of future networks. Independent evolvability is one of the biggest advantages of the edge pricing paradigm; while the total amount of mechanism needed to perform the necessary accounting may not be less than in other paradigms, the degree of independence of these accounting mechanisms is substantially higher. The ability for providers to act independently to upgrade their accounting will lead to rapid development of the required implementations; proposals that require a single uniform and standardized accounting infrastructure are much

to assess the appropriate charges.

[17] One could apply such a scheme to a logically overlaid network so the prices would be less dependent on the details of the path. For instance, the network could be divided up into area codes, with logical link costs recorded whenever the accounting message left one area code and entered another.

less likely to ever be implemented.

The above discussion applies to the complete spectrum of usage-constraining pricing schemes, from usage-based to capacity-based charging. However, much of the above discussion implicitly applied to best-effort service. The basic principles remain the same when pricing for reserved or assured levels of service, but the mechanistic details are quite different because of the presence of a set-up protocol like RSVP [46].

4.2 Charging Receivers

The second nonlocal problem we consider is assigning charges to receivers. This involves addressing the following three issues:

#1: *How does a receiver indicate to the network provider that it is willing to take responsibility for the source's traffic?* Here there are several alternatives, and we merely mention a few to illustrate some of the possibilities. In the best-effort multicast case, the join message might be extended to include a willingness-to-pay field. In the case of reservations, either unicast or multicast, the RSVP reservation message could carry similar information. The only case that does not already have a preexisting control message that could be used for this purpose is the unicast best-effort case. Here, we may require a new willingness-to-pay control message to be generated by the receiver, but there may also be other approaches. In addition, we may want to allow the source to indicate that it is not willing to pay, so that if the source's access point has not received a notification that the receiver(s) is(are) willing to pay then the packets are immediately dropped; such an indication that the source is not willing to pay could be contained in the packet header. Another approach – one that requires no additional signaling – is to divide the multicast addresses into sender-pays and receiver-pays categories, so that the assignment option is indicated by the choice of multicast address. Here the very act of joining the group communicates a willingness-to-pay.

#2 *How does the network "bill" the receiver?* One general approach here is to apply pricing when the packets traverse the receiver's access point. Thus, packets are "charged" according to the receiver's contract with its provider, not according to the sender's contract. If the capacity is exceeded, then the overlimit behavior (delaying, dropping, etc.) is applied at this exit point. If the packet traverses several providers, then this reverse charging is applied whenever a boundary is crossed; the packet is charged to the provider whose network the packet is entering, not the provider whose network the packet is exiting.

#3 *How does the network split the responsibility for the bill among the members of a multicast group?* If there are multiple receivers, the network not only needs to transfer the charges to the receivers, but also must apportion the cost among them in a reasonable manner. One way to do this is to assign fractional responsibilities to each of the receivers. Then, when the packet arrives at each receivers access point, the receiver is "charged" only the fraction of the normal amount. The variety of policies for assigning these fractions, as well as mechanisms for computing them, have been addressed in [20]. One could also use cruder approximations to compute these fractions, basing multicast prices on *ad hoc* discounts from the unicast cost.

4.3 Open Issues

The preceding general discussion merely presented some possible approaches. There are many other possibilities, and the ones mentioned above should be considered sketchy illustrations of the issues involved, rather than serious and complete design proposals. This initial design discussion leaves several fundamental issues unresolved; we mention a few of there here.

Our discussion of the need for charging receivers has focused on a narrow binary choice; either sources pay or receivers pay. One may want to consider a much broader spectrum of policies in which the costs are shared in a more flexible manner. This might be a fractional splitting (*e.g.*, the source pays 30% and the receiver(s) pay 70%), or perhaps the source pays for a certain portion of the path (*e.g.*, the source pays for the portion of the path within its local provider's network) and the receiver(s) pay the rest. We have not addressed the requirements of such source/receiver cost sharing.

We also have not considered the case where some receivers are willing to pay and others are not. Aside from the mechanistic questions, there are important unresolved policy questions about how to handle such a situation. Related to this is the fact that receivers may want to limit their exposure. The willingness-to-pay field may, in addition to indicating that the receiver is willing to pay, also indicate a cap on how much (in some arbitrary units) cost the receiver is willing to absorb. Such a cap may be necessary when joining what one expects to be large multicast groups. For instance, when a receiver in California joins a group for a virtual rock concert sourced from London with an expected audience of millions, the receiver may be willing to pay her share of a few dollars (or equivalent capacity units) but would certainly not be willing to absorb the bill for the entire 50 mbps video feed from London. However, such limits open up thorny strategic issues as receivers would be tempted to *free ride* on other receivers.

There may be other approaches to deal with the startup phase of multicast groups that will eventually become large. There may be some way that the organizer of the session, whether it be a rock concert or an IETF broadcast or a cable TV channel, could describe to the network beforehand an approximation to the likely distribution tree. This would enable the network to estimate the likely cost shares beforehand and thus greatly reduce the exposure of the first few group members. Other schemes to reduce such exposure have been discussed in [16].

The accounting mechanisms discussed in [20], which determine the appropriate multicast cost shares, are implemented on a link-by-link basis. Such methods must be extended to a more abstract set of logical links so that the cost shares can reflect a coarser level of granularity. Also, as discussed in [20], some cost-sharing approaches depend on the number of receivers downstream of each link. Such numbers are relatively easy to obtain within a provider's network (*e.g.*, by extending the multicast join mechanism). To do this accurately across providers would require each provider to reveal this number to other providers, and this raises an incentive question, since the cost share increases with the number of receivers (and so each provider would reveal only the existence of one receiver in their network).

Another issue arises in the case of receiver-pays with capacity-based charging, where the overlimit behavior is packet dropping. If the incoming traffic greatly exceeds the available capacity then the network has transported packets across the network only to consistently drop them at the exiting access point. We may need some slow, out-of-band, signaling in this case to *unjoin* the receiver from the group. Such signaling is not needed if the overlimit behavior is an additional per-packet charge.

5 Conclusion

Current discussions about pricing in computer networks are dominated by two main topics. The first topic is the debate between usage-based pricing and flat pricing which has embroiled the network design community and caught the attention of the popular press. Rather than being radically different, we think these two schemes reside along the single continuum of *usage-constraining* pricing policies. As in telephony, both pricing options, along with various intermediate hybrids, will likely be offered to users by their local provider. The detailed design of such schemes is perhaps best left to the marketing departments of the various network service providers. Thus, no particular pricing policy should be embedded into the network architecture. The challenge for the network design community is to provide a coherent network pricing architecture that allows individual providers to make their own choices about how to price service. This paper presents one such pricing architecture that achieves this goal: *edge* pricing.

The second topic, emphasized in the more academic literature, is the design of marginal cost pricing schemes that produce the optimally efficient use of network resources. We have critiqued this optimality paradigm on three grounds: (1) marginal cost prices may not produce sufficient revenue to fully recover costs and so are perhaps of limited relevance, (2) congestion costs are inherently inaccessible to the network and so cannot reliably form the basis for pricing, and (3) there are other, more structural, goals besides optimality, and some of these goals are incompatible with the globally uniformity required for optimal pricing schemes. For these reasons, we contend that the research agenda on pricing in computer network should shift away from the optimality paradigm and focus more on structural and architectural issues. Such issues include allowing local control of pricing policies, fostering interconnection, handling multicast appropriately, and allowing receivers to pay for transmission. To illustrate our point, we described how these goals might be accomplished in the context of the edge pricing paradigm.

Even though many of our detailed comments concern the particular pricing architecture of edge pricing, our intent in writing this paper is not to advocate that this scheme be adopted by the Internet; the proposal is extremely preliminary and there may be other schemes with similar properties. Rather, our intent is to initiate a dialogue about such pricing schemes and hopefully stimulate the creation of other pricing paradigms that meet our design goals.

Acknowledgements

We are grateful to Jeffrey MacKie-Mason, Hal Varian, Marjory Blumenthal, Padmanabhan Srinagesh, and Steve Deering for insightful comments on an earlier draft. We would also like to thank Van Jacobson for early conversations on edge pricing.

References

[1] R. Braden, D. Clark, and S. Shenker. Integrated services in the Internet architecture: an overview. Technical report, IETF, June 1994. RFC 1633.

[2] R. Bradford. Incentive compatible pricing and routing policies in multiserver queues. Technical Report MS L-195, Systems Research Group, Lawrence Livermore National Laboratory, 1994.

[3] Nevil Brownlee. New zealand experiences with network traffic charging. *ConneXions*, 8(12), December 1994. Available from URL: http://www.auckland.ac.nz/net/Accounting/nze.html.

[4] Stephen Casner and Stephen Deering. First IETF Internet audiocast. *Computer Communications Review*, 22:92–97, 1992.

[5] D. Clark, L. Zhang, and S. Shenker. Supporting real-time applications in an integrated services packet network: Architecture and mechanism. In *Proceedings of Sigcomm 92*. ACM Press, New York, 1992.

[6] David D. Clark. Adding service discrimination to the Internet. Technical report, MIT, August 1995.

[7] David D. Clark. A model for cost allocation and pricing in the Internet. Technical report, MIT, August 1995.

[8] Ron Cocchi, Deborah Estin, Scott Shenker, and Lixia Zhang. A study of priority pricing in multiple service class networks. In *Proceedings of Sigcomm '91*, September 1991. Available at ftp: //ftp.parc.xerox.com/pub/net-research/pricing-sc.ps.

[9] Ron Cocchi, Deborah Estrin, Scott Shenker, and Lixia Zhang. Pricing in computer networks: Motivation, formulation, and example. *ACM/IEEE Transactions on Networking*, 1:614–627, 1993.

[10] Robert Crandall. Funding the national information infrastructure: Advertising, subscription, and usage charges. Technical report, Brookings Institution, 1995.

[11] Stephen Deering. Personal communication, 1994.

[12] Stephen Deering and Robert Hinden. Internet protocol, version 6 (ipv6) specification. Technical report, IETF, June 1995. Internet Draft.

[13] Stephen E. Deering and David R. Cheriton. Multicast routing in datagram internetworks and extended LANS. *ACM Transactions on Computer Systems*, 8(2):85–110, May 1990.

[14] Richard J. Edell, Nick McKeown, and Pravin P. Varaiya. Billing users and pricing for TCP. *IEEE Journal on Selected Areas in Communications*, 13:1162–1175, 1995. Available from URL: http: //www-path. eecs.berkeley.edu/~edell/papers/Billing/article.ps.

[15] Paul Francis. Comparison of geographical and provider-rooted Internet addressing. In *Proceedings of INET '94*, June 1994.

[16] Rainer Gawlick. Admission control and routing: Theory and practice. Technical report, MIT, August 1995.

[17] Jion Gon and Padmanabhan Srinagesh. The economics of layered networks. Technical report, Bell Communications Research, Inc., 1995.

[18] Alok Gupta, Dale O. Stahl, and Andrew B. Whinston. Managing the Internet as an economic system. Technical report, University of Texas at Austin, July 1994. Available from URL: http://cism.bus.utexas.edu/ravi/pricing.ps.Z.

[19] Alok Gupta, Dale O. Stahl, and Andrew B. Whinston. The Internet: A future tragedy of the commons. Technical report, University of Texas at Austin, 1995.

[20] Shai Herzog, Scott Shenker, and Deborah Estrin. Sharing the 'cost' of multicast trees: An axiomatic analysis. In *Proceedings of Sigcomm '95*, August 1995.

[21] Michael Honig and Ken Steiglitz. Usage-based pricing and quality of service in data networks. *IEEE Journal on Selected Areas in Communications*, To appear, 1995.

[22] V. Jacobson. private communication, 1995.

[23] Hong Jiang and Scott Jordan. The role of price in the connection establishment process. *European Transactions on Telecommunications and Related Technologies*, 1995.

[24] Scott Jordan and Hong Jiang. Connection establishment in high speed networks. *IEEE Journal on Selected Areas in Communications*, 13:1150–1161, 1995.

[25] Yan-Jun Li and Joseph Hui. Service shadow price and link shadow price. In *Proceedings of the 2nd International Conference on Telecommunications Systems Modeling and Analysis*, 1994.

[26] Jeffrey K. MacKie-Mason, John Murphy, and Liam Murphy. ATM efficiency under various pricing schemes. Technical report, University of Michigan, Dublin City University and University of Auburn, March 1995.

[27] Jeffrey K. MacKie-Mason, John Murphy, and Liam Murphy. The role of responsive pricing in the Internet. Technical report, University of Michigan, Dublin City University and University of Auburn, June 1995. Available from URL: http: //www. ipps. lsa. umich. edu/ipps/papers/info-nets/Feedback_Role/feedback_role. ps.Z.

[28] Jeffrey K. MacKie-Mason, Scott Shenker, and Hal Varian. Service architecture and content provision. Technical report, University of Michigan and Xerox PARC, 1995.

[29] Jeffrey K. MacKie-Mason and Hal Varian. Pricing the Internet. In Brian Kahin and James Keller, editors, *Public Access to the Internet*. Prentice-Hall, Englewood Cliffs, New Jersey, 1995. Available from URL: ftp: //gopher. econ. lsa.umich.edu/pub/Papers/Pricing_the_Internet.ps.Z.

[30] Jeffrey K. MacKie-Mason and Hal R. Varian. Pricing congestible network resources. *IEEE Journal on Selected Areas in Communications*, 13:1141–1149, 1995. Available at URL: ftp: //gopher. econ. lsa.umich.edu/pub/Papers/pricing-congestible.ps.Z.

[31] John Markoff. Traffic jams already on the information highway. *New York Times*, November 3:A1, November 3 1993.

[32] Richard P. McLean and William W. Sharkey. An approach to the pricing of broadband telecommunications services. Technical report, Bellcore, March 1993.

[33] H. Mohring and M. Hartwiz. *Highway Benefits: An Analytical Approach.* Northwestern University Press, Evanston, 1962.

[34] John Murphy and Liam Murphy. Bandwidth allocation by pricing in ATM networks. In *Proc. of IFIP Broadband Communications '94*, Paris, France, March 1994. Available from URL: ftp: //gopher. econ. lsa. umich. edu/pub/Archive/Murphy_Bandwidth_Pricing.ps.Z.

[35] Liam Murphy and John Murphy. Pricing for ATM network efficiency. In *Proc. 3rd International Conference on Telecommunication Systems Modelling and Analysis*, Nashville, TN, March 1995.

[36] Colin Parris and Domenico Ferrari. A resource based pricing policy for real-time channels in a packet-switching network. Technical report, International Computer Science Institute, Berkeley, CA, 1992. Available at ftp: //tenet. icsi.berkeley.edu/pub/tenet/Papers/ParFer92.ps.

[37] Colin Parris, S Keshav, and Domenico Ferrari. A framework for the study of pricing in integrated networks. Technical Report Tech. Rept. TR-92-016, International Computer Science Institute, Berkeley, CA, 1992. Available at ftp: //tenet. icsi.berkeley.edu/pub/tenet/Papers/PaKeFe92.ps.

[38] Scott Shenker. Making greed work in networks: A game-theoretic analysis of switch service disciplines. In *Proceedings of Sigcomm '94*, August 1994. Available at ftp: //ftp.parc.xerox.com/pub/net-research/greed-cr.ps.

[39] Scott Shenker. Service models and pricing policies for an integrated services Internet. In Brian Kahin and James Keller, editors, *Public Access to the Internet.* Prentice-Hall, Englewood Cliffs, New Jersey, 1995.

[40] Scott Shenker. Some fundamental design decisions for the future Internet. *IEEE Journal on Selected Areas in Communications*, 13:1176–1188, 1995.

[41] Padmanabhan Srinagesh. Internet cost structure and interconnection agreements. Technical report, Bellcore, 1995.

[42] Dale O. Stahl and Andrew B. Whinston. An economic approach to client-server computing with priority classes. Technical report, University of Texas at Austin, 1992.

[43] Robert H. Strotz. Urban transportation parables. In J. Margolis, editor, *The Public Economy of Urban Communities*, pages 127–169. Resources for the Future, Washington, D.C., 1965.

[44] John Verity. Planet internet. *Business Week*, April 3 1995.

[45] Qiong Wang, Marvin A. Sirbu, and Jon M. Peha. An optimal pricing model for cell-switching integrated services networks. Technical report, Carnegie Mellon University, May 1995.

[46] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala. RSVP: A Resource ReSerVation Protocol. *IEEE Network Magazine*, 7:8–18, 1993. Available from URL: ftp://ftp.parc.xerox.com/pub/net-research/rsvp.ps.Z.

A Internet Architecture and Mechanisms

This appendix describes a few relevant features of the Internet architecture. It is a very selective and sketchy overview, intended merely to provide a minimal background for reading this paper. The current Internet architecture is designed for point-to-point (or unicast) *best-effort* communication. Every packet header contains a source address and a destination address. Upon receiving a packet, a switch (or, equivalently, a router) consults its routing table to find, based on the packet's destination address, the appropriate outgoing link for the packet. The network makes no commitments about when, or even if, packets will be delivered. Sometimes the incoming rate of packets at a switch is greater than the outgoing rate, and so queues build up in the switch. These queues cause packet delays and, if the switch runs out of buffer space, packet discards. The network does not attempt to schedule use; sources can send packets at any time, and the network switches merely exert their *best effort* to handle the load.

This simple network architecture has been amazingly successful. However, there are efforts currently underway to extend this architecture in two ways. The first is to offer better support for multipoint-to-multipoint communications through the use of *multicast* (see [13] for the seminal paper on the topic). In the current Internet architecture, when a source sends a packet to multiple receivers, the source must replicate the packet and send one to each receiver individually. This results in the several copies of the same packet traversing those links common to the delivery paths (*i.e.*, those links that lie on more than one delivery path, where the delivery path is the route taken by the packet from source to receiver). In multicast, the source merely sends the packet once, and the packet is replicated by the network only when necessary (*i.e.*, the packet is transmitted only once on each link, and then is replicated at the split points where the delivery paths diverge and one copy is sent along each outgoing branch). Sources sending to a multicast group use the multicast group address as the packet's destination address; however, multicast addresses are merely logical names and do not convey any information about the location of the receivers (unlike a unicast address). Computers on a network wishing to receive packets sent to a particular multicast address send a *join* message to the nearest router. The routing algorithm then distributes this information to create the appropriate distribution trees (*i.e.*, trees from every source to every receiver) so that packets sent to the group reach each receiver. There are variety of routing algorithms that can accomplish this task, and we do not review them here. Note that senders are not aware of who is receiving the packets, since the multicast paradigm is receiver-driven. Efforts to standardize and deploy multicast are well advanced; the vitality of the current MBone [4] attests to the benefits of this technology.

The second extension to the Internet architecture is much more preliminary, and rather controversial. Efforts are underway to extend the Internet's current service offerings to include a wider variety of qualities of service (QoS). The current single class of best-effort service may not be sufficient to adequately support the requirements of some future video and voice applications (although this is a highly debatable point; see [40]). Moreover, offering all applications the same service is not an efficient use of bandwidth; providing a wider variety of qualities of service allows the network's scarce resources to be devoted to those applications that are most performance-sensitive. There are many ways in which these services could be extended, some as simple as merely providing several service priority levels and/or drop priority levels. See also the discussion in [7, 6] for other approaches to such extensions to best-effort service. Offering multiple qualities of service requires some form of incentives, such as pricing, to encourage the appropriate use of the service classes; see [5, 9, 8, 39] for a discussion of these issues.

More radical extensions to the service offerings are also contemplated. A working group of the Internet Engineering Task Force is preparing a proposal to offer several *real-time* services; a bounded-delay service, in which the network commits to deliver all packets within a certain delay, is an example of such a real-time service. These services are fundamentally different than best-effort in that the network is making an explicit and quantitative service commitment and therefore must reserve the appropriate resources. Such services require admission control procedures, whereby receivers request service (*i.e.*, issue a reservation request) and the network then either commits to the requested level of service (if it can meet the requirements), or denies the reservation request (if the current load level is to high to meet the requirements of the incoming request). In the proposed resource reservation protocol RSVP [46], receivers send their request for service to the network, and this request follows the reverse delivery tree towards all relevant sources (a single source if the application is unicast, or to all senders to the group if the application is multicast). [1] presents a slightly out-of-date overview of this proposed architecture.

B Ramsey Prices in a Simple Model

In this appendix we explore the behavior of Ramsey prices in a simple network model. We consider a facility providing network service charging a price p for every unit of usage and an attachment cost q; we assume there are no usage-dependent costs associated with the facility, and for convenience we consider only nonnegative prices. The user population is a continuum labeled by α, with $\alpha \in [0,1]$. The usage of each user is denoted by x_α. We consider utility functions of the form $U_\alpha = V_\alpha - px_\alpha - q$, where V_α represents the valuation of usage, and assume users can detach (yielding a utility $U_\alpha = 0$) if prices are too high. The total welfare is given by $W = \int d\alpha\ V_\alpha$, the total usage $Y = \int d\alpha\ x_\alpha$, and the total revenue by $R = pY + q \int d\alpha$, where the integrals run over all attached users.

The functions V_α take the form $V_\alpha = \lambda_\alpha x_\alpha - x_\alpha^2$. Each attached user sets their $x_\alpha = \frac{\lambda_\alpha - p}{2}$. We consider three cases for the λ_α: (1) a homogeneous case (where all users have the same V_α), (2) a heterogeneous case (where users have different V_α) without network externalities, and (3) a heterogeneous case with network externalities (where one user's valuation depends on the number of other attached users).

In the homogeneous case we set $\lambda_\alpha = 1$ for all α. The total welfare is $W = \frac{1-p^2}{4}$ as long as $0 \le p \le 1$ and $0 \le q \le \frac{(1-p)^2}{4}$ (otherwise all users detach and $W = 0$). Thus, welfare is maximized when $p = 0$, $q \le \frac{1}{4}$ and each $x_\alpha = \frac{1}{2}$. Setting $p = 0$ and $q = \frac{1}{4}$ raises maximal revenue in this case. In this homogeneous case, attachment prices are indeed the optimal way to raise additional revenue.

We now consider a heterogeneous case where $\lambda_\alpha = \alpha$. For a given p and q, all users with $\alpha > A(p, q) = p + 2\sqrt{q}$ remain attached. The total welfare is given by $W = \frac{1}{12}(1 - A^3) - \frac{p^2}{4}(1 - A)$, and the total revenue is given by $R = q(1 - A) + \frac{p}{4}(1 - A^2) - \frac{p^2}{2}(1 - A)$. The total welfare is maximized when $p = 0$ and $q = 0$. The curve of Ramsey prices, the points that maximize W for a fixed R, is given by $q = p^2$ for $0 \le p \le .2$. The point $p = .2$ and $q = .04$ maximizes the revenue R. The quadratic nature of the Ramsey curve means that increases in usage prices dominate (over increases in attachment prices) close to the origin.

We introduce network externalities by allowing the constants λ_α to depend on the number of other attached users. Let $\lambda_\alpha = \alpha(1 - A)$ where, as above, A is the critical value of α such that all users with $\alpha > A$ are attached and no users with with $\alpha < A$ are attached. Then $A = \frac{1}{2}(1 - (1 - 8\sqrt{q} - 4p)^{0.5})$. In addition, $W = (1 - A)\left(\frac{(1-A^3)(1-A)}{12} - \frac{p^2}{4}\right)$ and $R = (1 - A)\left(\frac{p(1-A^2)}{4} - \frac{p^2}{2} + q\right)$. The point $p = 0$ and $q = 0$ maximizes W. The point $p = 0.16$ and $q = 0$ maximizes R. The Ramsey prices fall along the line segment between these two points: $q = 0$ and $0 \le p \le 0.16$. For this model, increasing revenue is best done through increasing usage charges only.

The simple model we considered here is extremely unrealistic, and neglects important aspects of the problem such as congestion. However, it does illustrate the basic point that when one has a heterogeneous population containing users who derive marginal benefit from attachment, then raising the attachment prices alone is not necessarily a Ramsey price.

C Capacity Filters

The ability to express the capacity in terms an arbitrary filter provides substantial flexibility for accommodating user needs. In this appendix we give a concrete example of a sophisticated usage filter. A *token bucket* filter is parameterized by a rate r and a bucket size b. Usage complies with this capacity as long at the cumulative number of units sent in any time interval of length t (for any such t) is bounded above by $rt + b$. This allows bursts of size b, but bounds the long-term average to be no greater than r. A filter might be the composition of three different token buckets, one with $b = 0$, one with $b = \infty$ (in actual practice, this value of b will be chosen to be large but finite, since an infinite sized b imposes no constraints) and one with an intermediate value of b_i. The values of r associated with the two extremal bucket sizes control very different aspects of the traffic: r_0 describes the allowable peak rate and r_∞ describes the allowable long-term average rate (the actual size of the large but finite

b value used in practice determines, along with the associated rate, the time interval over which this long-term average is applied). The intermediate parameters r_i, b_i describes some intermediate allowable burst rate and size. These different filters should be thought of as constraining the flow on different time scales: the bigger the b the longer the time scale. In general, one might describe a filter as a nonincreasing function $b(r)$; for every r there is a token bucket with parameters $r, b(r)$ applied to the flow. By adjusting these parameters appropriately one can provision the capacity for various levels of web-browsing or video consumption.