Protocol-independent multicast pricing

Tristan N.H. Henderson and Saleem N. Bhatti
Department of Computer Science, University College London
{T.Henderson,S.Bhatti}@cs.ucl.ac.uk *

Abstract

Current multicast pricing proposals are dependent on specific routing protocols or require changes to the existing IP multicast model. Given that multicast has only seen limited deployment thus far, such schemes may become redundant if the multicast service model changes. We believe that instead of altering multicast delivery methods to suit particular pricing schemes, a multicast pricing scheme should be designed around the generic concept of multipoint communication, without depending on the underlying transmission methods. We highlight limitations of existing work in this area and present a work-in-progress; a framework that enables this separation between transmission and pricing.

1 Introduction

Multicast has long been viewed as an attractive service for the Internet for enabling multiparty applications, and has become an important research area. More than a decade after its inception, however, there has been little multicast deployment by commercial Internet Service Providers (ISPs) and most end-users still lack multicast capabilities [2]. One impediment to the broader deployment of multicast has been the current service model's inability to provide an easy and effective pricing structure; UUnet's vice-president of marketing has said that "Pricing multicast is very non-intuitive because there's no correlation between the number of recipients and the cost for network services" [6]. Diot *et al.* [12] cite pricing as a major issue in their analysis of the requirements for further multicast deployment.

The main motivations for network accounting and pricing are understanding/influencing behaviour, measuring policy compliance, and rational cost allocation/recovery [22]. In other words, as well as providing a method for network providers to recover costs, pricing is intended to make users aware of the costs of their actions, which will encourage them to act in a socially-optimal way. This will enable the introduction of different qualities of service (QoS), since without responsibility for costs, users will always choose the highest QoS available. A multicast pricing scheme should therefore provide incentives for both end-users and network providers to adopt and use multicast efficiently. As well as this, the scheme should be simple to implement if it is to be adopted by network providers.

Existing work into multicast pricing has been closely tied to a specific routing protocol, required changes to the existing multicast service model, or been limited to a specific session type such as single-source. *We believe that a multicast pricing scheme must be independent of any underlying network protocols.* It is unclear whether multicast will become universally deployed in its current form, as research into new methods for multicast shows [14, 29]. Multipoint communication, however, is likely to become popular irrespective of the method of transmission. Appropriate incentives and requirements for pricing should thus be determined from user behaviour in multipoint scenarios.

Not limiting the scheme to any specific transmission method minimises changes to existing infrastructure, makes it more flexible and reduces the possibility of redundancy in the event of protocols being obsoleted. Given the variety of multicast routing protocols, a protocol-independent pricing scheme can handle inter-domain transmissions more easily, and be used by all ISPs. If the same scheme can be used for single and multiple-source, or dynamic and static membership sessions, different sessions will be comparable in price, i.e., more stable and predictable prices for end-users. The INDEX project [8] has shown such stability to be desirable.

In Section 2 we list some of the existing work in this area and the issues that a multicast pricing scheme should consider. Section 3 outlines the framework that we are developing, and in Section 4 we present our conclusions.

2 Related work and issues

2.1 Terminology

Internet pricing is studied by computer scientists and economists, each of whom use their own, often distinct,

*Tristan Henderson is funded by an HP/EPSRC CASE studentship.

terms to describe similar concepts. We have chosen to use definitions from a variety of sources. We use the terms *pricing*, *charging* and *billing* as in Stiller *et al.* [28]. *Multicast* is sending to a set of hosts using a single address [7], whereas *multipoint* is "any type of communication within a group" [11]. Both *flow* and *session* refer to the transmission of a sequence of multicast packets. A flow is considered to be a single media stream from a single source, while a session may consist of one or more flows grouped by context, e.g., the audio and video streams of a film.

2.2 Pricing

Herzog *et al.* [17] is the earliest work on multicast pricing. Their scheme, ELSD (Equal Link Split Downstream), splits costs amongst downstream receivers and allocates no costs upstream. This is shown to be an optimal cost allocation for single-source sessions with a source-rooted tree, but requires changes to the IP multicast service. ELSD is modified in [16] for multiple-source sessions, at the expense of some scaleability. It is not clear, however, that ELSD works for dynamic sessions.

EXPRESS [18] modifies IP multicast for large-scale single-source applications. Multiple-source sessions are catered for by a "session relay" approach, where all sources' traffic is sent via one node. This master node becomes a single point of failure, and may not be able to cope with delay-intolerant applications.

Carle *et al.* [5] examine multicast in an ATM intserv environment. Charges are determined based on requested and received QoS, with the charging protocol encapsulated in RSVP. Since prices are determined a posteriori, they are unpredictable, and the scheme depends on RSVP.

In "split-edge pricing" [3], both sender and receiver initially pay a share of the cost of a transmission, and claims over the value of the transmission are settled later. Network providers agree prices for offering each subscription level to their neighbours, and these charges are summed to create the price for a complete transmission. It is shown that for multicast pricing both sender and receiver need to pay, because otherwise an incentive exists for downstream providers to lie about the number of receivers. This problem still exists with split-edge pricing, however, since senders and receivers settle claims after transmission, and no mechanism is provided for verifying the number of downstream receivers. This is the "collusion prevention" axiom of Herzog's thesis [16], which also states that this problem cannot be solved through cost allocation alone.

Einsiedler *et al.* [13] propose assigning weights to each link in a network, to represent the "cost" of that link. These weights can be derived from the congestion along the link, the costs of maintenance, or inter-domain costs for links that traverse ISPs. An extra Internet Group Management Protocol (IGMP) message or IPv6 header extension is used to store the weight information. Costs are determined, as in ELSD, by splitting the costs at each branching point in the tree. This scheme also violates the collusion prevention axiom, since the charge depends on who is paying, which may create incentives for senders or receivers to always pay, depending on which is cheaper.

Chuang and Sirbu [9] analyse the cost efficiency of multicast over unicast. Multicast costs are capped at the unicast level, and it is assumed that new joins to an existing branch in a tree represent a zero marginal cost. A pricing scheme based on this would thus be incentive-incompatible, since it would only charge the first user in a branch.

Ravindran and Gong[25] analyse multicast costs to determine the "cheapest" tree topology for a session in terms of network resources. Their method is protocol-independent, although pricing and user incentives are not considered.

There is also related work on multicast charging and billing. Kausar *et al.* [20] propose charging on a per-session basis, although this may prove difficult for highly dynamic sessions where each user's duration varies greatly. The Resource Negotiation and Pricing protocol (RNAP) [32] is a pricing protocol for delivering multiple levels of QoS, by integrating reservation, charging and pricing. Single-source multicast is treated the same as unicast, so that there is no incentive for receivers to choose multicast.

2.3 Routing

Almeroth [1] outlines how multicast has evolved since the introduction of the MBone, and the various routing protocols[1] that network providers can choose to use.

Multicast at the end-user level is still rare. Amongst those ISPs that have implemented multicast, Protocol Independent Multicast (PIM) is a popular protocol. Some ISPs, however, have chosen not to use IP multicast at all; UUnet's "multicast" service, UUcast [31] uses unicast proxies as well as sparse-mode PIM.

At the inter-domain level, the choice of routing protocol is up to the individual ISP. Given the limited scale of multicast deployment, statistics on protocol use might not be useful since the market is not large enough to create a de facto standard. Ongoing research into protocols such as PIM, Core-Based Trees (CBT), and Multicast extensions to Open Shortest Path First (MOSPF), indicates that they are all potentially in use.

In the core multicast backbone, which at present is mainly the MBone, the prevalent protocol is the Distance Vector Multicast Routing Protocol (DVMRP), with the Border Gateway Protocol (BGP) connecting Autonomous Systems (ASs). The Border Gateway Multicast Protocol

[1] The protocols mentioned in this section are described in [10]

(BGMP) [21] is designed to improve scaleability in inter-domain routing.

It is not clear that multicast, if it is to ever reach "critical mass", will continue to be based on the current IP multicast model. Francis [14] lists some of the problems with IP multicast and presents a new architecture that enables replicates content in the end-hosts. The REUNITE (REcursive UNicast TreE) scheme [29], on the other hand, allows multicast using unicast instead of class D addressing.

2.4 Session control

IP multicast is a network-level transmission mechanism, and routing protocols are designed with this in mind, attempting to be optimal according to a network-level metric. Multicast applications, however, are designed around the concept of sessions. The need for session-based control has led to the development of session-based signalling protocols, and, as with routing, there is a choice of these, such as H.323 [19] and the Session Initiation Protocol (SIP) [15]. SIP is still a work-in-progress and so it is difficult to assume the deployment of either of these protocols.

2.5 Security

Attempting to charge for multicast necessitates the exclusion of those users unwilling or unable to pay. Any pricing scheme must therefore provide some sort of access control, to ensure that only those who have paid or agreed to pay are allowed to join or transmit to multicast groups. This violates the current IP multicast model, where any user can send to any existing session.

Canetti *et al.* [4] list multicast security concerns, including access control. For pricing, however, many of these are non-essential. As a minimum, the only requirement may be authentication to ensure that data is only transmitted to and from authorised users. Such authentication may appear to the user to be at the session-level (i.e., they are permitted to join the session), but this requires authentication services both at the session-level (by using, e.g., application-level key distribution) and at the packet-level, to prevent, for example, spoofing source addresses of legitmate users.

2.6 User behaviour

End-user behaviour is interesting from a pricing perspective in a number of ways. The join/leave and duration characteristics of individual users' sessions can affect price stability. On the other hand, prices can affect user behaviour depending on the incentives created. A pricing scheme also needs to reflect user expectations accurately in order to be acceptable to end-users.

There has been little work on how users respond to the presence of multicast, with most studies concentrating on network characteristics. Almeroth [2] analyses join and leave behaviour for MBone sessions over a 4.5 year period. Due to the limited size of the MBone, however, there are only between 100 and 200 active receivers at most times, and so it is difficult to draw strong conclusions from this data. Most of these sessions are single-source, and we know of no analysis of multiple-source sessions.

2.7 User heterogeneity

As well as being able to cope with multiple and dynamic sources, an ideal pricing scheme should acknowledge user heterogeneity. A user might be accessing a session through a limited medium such as dial-up or wireless, or might wish to limit the amount of bandwidth used by a particular session.

3 Our Work

We are currently developing an application-level framework that will allow network providers to charge for multicast usage from a specific point in their network. This *charging point (CP)* is chosen to reflect a specific location in the network at which charging is required for senders or receivers. The CP is chosen as an administrative convenience but we expect it to be closely linked to the paths of the multicast packets, i.e., it can be seen as the point from which a user obtains access to a multicast flow or where a multicast flow is transmitted onto the network. Our framework can be applied at any point in the network, for example, at an ISP's dial-up server, or within the core of the network as part of a peering agreement between ASs.

3.1 The components of pricing

We assume that there are several components to the price paid for a multicast transmission. This applies to both senders and receivers. The exact nature of the pricing is specific to a particular network operator. The charging of a particular session may depend on factors such as content. For example, a film may incur a single payment, whereas a news channel may have a monthly subscription. Such matters are policy-specific, and we only consider recovery of transmission costs and other overheads for a flow (such as providing the network service or generating bills). We call this the *multicast recovery cost,* C_{RM}. C_{RM} is then used as a component of the final charge for a particular multicast session. A session may consist of several flows, and it is up to the network provider to collect the correct accounting information in order to construct the user's final session charge.

We assume, like Chuang and Sirbu, that C_{RM} will have a maximum cost equivalent to the recovery cost for transmission of a unicast flow, which we label C_{RU}. There will also be a minimum value for C_{RM}, specific to particular network providers, which may include recovery of costs for functions such as the costs of accounting for network usage and generating bills, similar to the way that some telecoms providers stipulate a minimum charge for telephone calls.

3.2 The SCALE pricing framework

Our framework uses the basic concepts of transmission and duplication common to all multipoint mechanisms. A transmission is sent from a host and is duplicated at various points in the network in order to reach all the receivers. Where a multicast packet is duplicated — transmitted on more than one interface — this incurs an additional cost. The costs for transmitters and receivers are assessed separately, and we assume that these costs are determined using the costs of transmission from the CPs.

3.2.1 Receivers

Figure 1 shows a multicast flow entering CP_1 and being duplicated on three outgoing interfaces $I_{\{1..3\}}$. CP_1 could, for example, be a multicast router but could also be another network element or even an application-level entity. The cost of the flow up to CP_1 is C, and the costs of duplication and transmission onto the the other three interfaces are $C_{\{1..3\}}$, as shown. These costs are all multicast recovery costs.

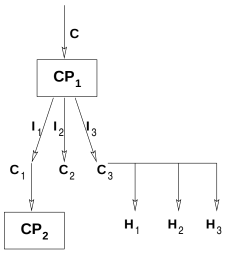

Fig. 1. Packet duplication at a charging point

Where several downstream users can receive the same data, they can effectively share C_{RM} between them, i.e., they could share the cost of the flow from the last CP. For example, hosts $H_{\{1..3\}}$, on a LAN receiving from I_3 on CP_1 would all share the cost C_3. Part of C_3 would comprise part of C, as would C_1 and C_2. In this way, at any point in the network, receivers need only know the cost of reception with respect to their CP.

Note that a CP could actually be any distribution point for a flow. It could be, for instance, an ISP's dial-up server, in which case the flow would be duplicated to each dial-up port and it is likely that $C_{RM} \approx C_{RU}$. The CP could also be an H.323 multipoint control unit (MCU) or part of an H.323 Gatekeeper. In Figure 1, CP_1 and CP_2 could be border routers between two ASs.

The CP does not have to be a physical network element; Figure 2 shows an ISP's network being used as a CP. Here, the users pay for a flow that is received from a network service provider. In reality, the physical access point to the network may be a dial-up server or leased line. Access to the CP is through interaction with a *Session Charging, Access and Logging Entity (SCALE)*.

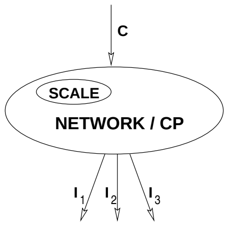

Fig. 2. A network as a charging point

3.2.2 The SCALE

The SCALE is the function responsible for allowing users to receive or transmit a multicast flow. An end-system must contact the SCALE before it can receive or transmit a multicast flow that is subject to a charge. The SCALE function sits within the provider network and gives access to a CP. The SCALE and CP need not be co-located, but they need to communicate. The SCALE function may form part of a bandwidth broker [24] or an H.323 Gatekeeper, for example. This allows an end-system to receive a flow from many different providers or many different points in the network by communicating with the right SCALE, so heterogeneous receivers can connect to the SCALE best suited to their capabilities or willingness to pay.

The functions of a SCALE are to provide:

- a signalling point for users to access multicast transmissions
- access control instructions to CPs to allow users to access multicast sessions, authenticate receiving users and enable appropriate filtering to be set up at CPs or other network elements
- information to users and network providers on the costs of a multicast transmission
- logging of multicast usage for accounting and billing

Consider a receiver, R_1, wishing to have access to a multicast session. R_1 obtains information about the session, for example, by using an *sdr*-like mechanism provided locally by the network service provider. This session information includes a session charge, which would typically be specified as a range based on the maximum and minimum values of C_{RM} of each flow in the session. Once R_1 has selected the session, it contacts the SCALE and goes through an authentication process, which includes an exchange that establishes R_1's willingness and ability to pay. The SCALE then contacts various points of the network (such as the CP) and sets up appropriate filtering to allow the transmission to be received by R_1. Access control takes place on a per-session basis, whilst charging is performed per-flow.

Cost information is transferred between CPs by the interaction of the SCALEs representing those CPs. The interaction is a provider-user relationship, and so the same SCALE functionality can be used at many different types of CP, both at the edge of the network and within the core.

3.2.3 Transmitters

A potential transmitter must interact with a SCALE to:

- establish the cost of the transmission of a session
- enable multicast forwarding through the CP, which again requires user authentication and establishing ability to pay

A potential transmitter, T_1, who wishes to initiate a new session declares the scope of this session to a SCALE (the choice of SCALE is left to the user and market). The SCALE then interacts with other SCALEs within this scope to determine the cost information for the session, using the CPs that represent this scope. From these queries the SCALE is able to provide maximum and minimum charges for the session. A transmitter T_2 who wishes to send to T_1's existing session also contacts a SCALE. The same cost information may be used, but this will depend on routing and policy; asymmetric tree topologies might exist, or some domains might implement policy-based forwarding.

Whether senders or receivers pay for transmission is policy-specific, and depends on who is responsible for the traffic in a given session. It might also be possible for senders and receivers to share the charge; a sender could declare session scope and any receivers outside this scope who wish to join the session would be responsible for the charge from a CP within the scope to their nearest CP.

3.3 Determining prices

In principle pricing policy is left to individual ISPs; they are free to advertise whatever prices they wish via the SCALEs in their domains. All that is required is that they must be able to furnish the SCALEs with sufficient information to allow them to determine a maximum and minimum price for a given session. There are some considerations that ISPs may need to make, however, when deciding what prices to advertise.

If we assume a competitive multicast service provision market, we may expect that providers will compete on price, decreasing prices until they equal marginal cost in a Bertrand equilibrium [30]. This is a necessary assumption; if service providers are unregulated monopolists (e.g., there are no alternative routing paths or access methods), it is unlikely that they will be interested in efficient network accounting or pricing, since they are price-setters. If price competition exists, then where flow-sharing is possible, this cost-saving will be reflected in receivers' prices. Similarly, as users join and leave the session, the cost, and price, for senders will vary. Users will thus have incentives to use multicast when it is cheaper to do so, from a provider's point of view. There might be minor distortion between cost and price, since we have assumed that C_{RM} is bounded at C_{RU}, which might not be true for very small sessions.

When pricing multicast, there is a danger that prices can beome unstable. Assume that users will only stay in a multicast session as long as their received utility exceeds the price, and that both utility and price are related to the number of members in a session (utility since a session is only useful if data is transmitted between two or more parties, and price since the cost of receiving data might be shared amongst receivers). A user leaving a session thus leads to a decrease in utility and an increase in price for the remaining members. This may lead to other users leaving the session, if the new higher price exceeds their new lower utility, and so on. This is similar to the layered multicast synchronisation problem that the synchronisation points of Receiver-driven Layered Congestion control (RLC) [26] are designed to solve. To limit the scale of this problem we restrict price changes to particular points in time, called *Session Cost Re-evaluation Points (SCRAPs)*. Prices are determined using the session membership at the previous SCRAP. The frequency of SCRAPs is context-dependent, e.g., SCRAPs in a news flow might occur more often than in a pay-per-view event. User behaviour might also play a part in determining SCRAP frequency — there may be incentives to join at particular times if users know when a SCRAP will occur.

4 Conclusion and further work

We have argued that existing multicast pricing schemes have been tied to particular service models or transmission protocols. We have presented a framework for multicast pricing that is independent of the underlying network topology and protocols. We believe that this is useful because rather than adapting service models to a particular pricing scheme, a scheme should adapt to the method of delivery.

Unlike most proposed schemes, it is possible for both sender and receiver to be made responsible for the cost. This is because a multicast pricing scheme needs to be flexible enough to deal with the multitude of possible pricing policies. It is foreseeable that multicast will be popular with advertisers, since they will be able to reach a large audience at low cost. Responsibility for such transmissions lies ultimately with the sender, and so a mechanism for charging senders at least part of the cost is required. Even for non-advertising data, a sender is implicitly willing to pay something, since they must receive a positive utility from making their information available. For senders to always pay, however, ignores the case where receivers are to blame for traffic, or are willing to pay to access a session for which they are out of scope.

Our scheme requires a change to the IP multicast model, in that senders are not allowed to transmit unless they join a session. As mentioned in Section 2.5, however, this change is necessary, again because of the possibility of advertising. To allow unauthorised transmissions permits advertisers to transmit without taking into account the cost of their advertising, both in terms of transmission costs and users' disutility from receiving unwanted spam. Any pricing scheme which wishes to charge receivers cannot be implemented, therefore, without some accompanying security measures. We have yet to determine the other security details of the SCALEs, e.g., user authentication, verification of ability to pay and perhaps per-packet authentication.

Since CPs and SCALEs can be located anywhere in the network, the scheme should theoretically scale well. Heterogeneous receivers can be catered for by different SCALEs, and SCALEs could also be incorporated into existing application-level entities such as H.323 Gatekeepers. There might be problems in practice, however, stemming from session dynamics. If join and leave behaviour is such that prices are constantly changing, or that SCRAPs need to be held every second, this scheme might be impractical. Given the lack of suitable high-membership dynamic multicast sessions, we are currently analysing the join and leave behaviour of online games such as Quake. Although these games are unicast, they are multipoint communication from a user perspective.

To make prices predictable we bound a user's charge and limit the amount of possible price changes. There is a trade-off between efficiency and predictability in pricing; completely stable prices in a dynamic multicast session would indicate a disparity between cost and price, since the cost of a session depends on the composition of the session membership, which varies over time. An optimally-efficient pricing scheme, however, might be an unrealistic goal [27]. We believe that informing users of the range of charges that they can expect is an acceptable solution, and one currently in use by some online retailers such as MobShop [23].

The framework as presented here is concerned only with best-effort multicast. Further work will consider multiple levels of QoS and congestion control. In particular we intend to examine user behaviour in layered multicast sessions; incentives may exist to subscribe to lower levels of QoS, since if a user subscribes to a higher level, the cost of transporting this additional layer to other users located nearby in the tree becomes lower.

Acknowledgement

We would like to thank Jon Crowcroft and the reviewers for useful comments.

References

[1] K. C. Almeroth. The evolution of multicast: From the MBone to inter-domain multicast to internet2 deployment. *IEEE Network*, 14(1):10–20, Jan. 2000.

[2] K. C. Almeroth. A long-term analysis of growth and usage patterns in the multicast backbone. In *Proc. INFOCOM 2000*, Tel Aviv, Israel, Mar. 2000.

[3] B. Briscoe. The direction of value flow in connectionless networks. In *Proc. NGC '99*, Pisa, Italy, Nov. 1999.

[4] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas. Multicast security: A taxonomy and some efficient constructions. In *Proc. INFOCOM '99*, New York, NY, Mar. 1999.

[5] G. Carle, M. Smirnov, and T. Zseby. Charging and accounting architecture for IP multicast integrated services over ATM. In *Proc. Internetworking '98*, Ottawa, Canada, July 1998.

[6] J. Caruso and K. Gerwig. Internet's congestion makes streaming iffy. *InternetWeek*, Mar. 02, 1998.

[7] D. R. Cheriton and S. E. Deering. Host groups: A multicast extension for datagram internetworks. In *Proc. the 9th symposium on Data communications*, pages 172–179, British Columbia, Canada, Sept. 1985.

[8] K. Chu and J. Altmann. Demand for different qualities of service for internet access: A review of INDEX findings. In *Proc. Royal Society Discussion Meeting on Network modelling in the 21st century*, London, UK, Dec. 1999.

[9] J. C. Chuang and M. A. Sirbu. Pricing multicast communication: A cost-based approach. In *Proc. INET '98*, Geneva, Switzerland, July 1998.

[10] J. Crowcroft, M. Handley, and I. Wakeman. *Internetworking Multimedia*. Taylor & Francis, London, UK, 1999.

[11] C. Diot, W. Dabbous, and J. Crowcroft. Multipoint communication: A survey of protocols, functions, and mechanisms. *IEEE Journal of Selected Areas In Communications*, 15(3):277–290, Apr. 1997.

[12] C. Diot, B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen. Deployment issues for the IP multicast service and architecture. *IEEE Network*, 14(1):78–88, Jan. 2000.

[13] H. J. Einsiedler, P. Hurley, B. Stiller, and T. Braun. Charging multicast communications based on a tree metric. In *Proc. 1st Workshop on Multicast-Protokolle und Anwendungen*, Braunschweig, Germany, May 1999.

[14] P. Francis. Yallcast: Extending the internet multicast architecture. unrefereed report, Sept. 1999.

[15] M. Handley, H. Schulzrinne, E. Schooler, and J. Rosenberg. SIP: Session initiation protocol, Mar. 1999. RFC 2543.

[16] S. Herzog. *Accounting and Access Control for Multicast Distributions: Models and Mechanisms*. PhD thesis, University of Southern California, Los Angeles, CA, Aug. 1996.

[17] S. Herzog, S. Shenker, and D. Estrin. Sharing the "cost" of multicast trees: An axiomatic analysis. In *Proc. SIGCOMM '95*, pages 315–327, Cambridge, MA, Aug. 1995.

[18] H. W. Holbrook and D. R. Cheriton. IP multicast channels: EXPRESS support for large-scale single-source applications. In *Proc. SIGCOMM '99*, pages 65–78, Cambridge, MA, Aug. 1999.

[19] ITU-T Recommendation H.323. Packet-based multimedia communications systems, Feb. 1998.

[20] N. Kausar, B. Briscoe, and J. Crowcroft. A charging model for sessions on the internet. In *Proc. ISCC '99*, Sharm El Sheikh, Egypt, July 1999.

[21] S. Kumar, P. Radoslavov, D. Thaler, C. Alaettinoğlu, D. Estrin, and M. Handley. The MASC/BGMP architecture for inter-domain multicast routing. In *Proc. SIGCOMM '98*, pages 93–104, Vancouver, Canada, Aug. 1998.

[22] C. Mills, D. Hirsh, and G. Ruth. Internet accounting: Background, Nov. 1991. RFC 1272.

[23] MobShop. http://www.mobshop.com.

[24] K. Nichols, V. Jacobson, and L. Zhang. A two-bit differentiated services architecture for the internet, Sept. 1999. RFC 2638.

[25] K. Ravindran and T.-J. Gong. Cost analysis of multicast transport architectures in multiservice networks. *IEEE/ACM Transactions on Networking*, 6(1):94–109, Feb. 1998.

[26] L. Rizzo, L. Vicisano, and J. Crowcroft. TCP-like congestion control for layered multicast data transfer. In *Proc. INFOCOM '98*, San Francisco, CA, Mar. 1998.

[27] S. Shenker, D. Clark, D. Estrin, and S. Herzog. Pricing in computer networks: Reshaping the research agenda. *ACM Computer Communication Review*, 26(2):19–43, Apr. 1996.

[28] B. Stiller, G. Fankhauser, B. Plattner, and N. Weiler. Charging and accounting for integrated internet services - state of the art, problems, and trends. In *Proc. INET '98*, Geneva, Switzerland, July 1998.

[29] I. Stoica, T. E. Ng, and H. Zhang. REUNITE: A recursive unicast approach to multicast. In *Proc. INFOCOM 2000*, Tel Aviv, Israel, Mar. 2000.

[30] J. Tirole. *The Theory of Industrial Organization*. MIT Press, Cambridge, MA, 1988.

[31] UUcast. http://www.uu.net/products/uucast/.

[32] X. Wang and H. Schulzrinne. RNAP: A resource negotiation and pricing protocol. In *Proc. NOSSDAV '99*, Basking Ridge, NJ, June 1999.