

# Sharing the “Cost” of Multicast Trees: An Axiomatic Analysis

Shai Herzog, *Member, IEEE*, Scott Shenker, *Senior Member, IEEE*, and Deborah Estrin, *Member, IEEE*

**Abstract**—Given the need to provide users with reasonable feedback about the “costs” their network usage incurs and the increasingly commercial nature of the Internet, we believe that the allocation of cost among users will play an important role in future networks. This paper discusses cost allocation in the context of multicast flows. The question we discuss is this: When a single data flow is shared among many receivers, how does one split the cost of that flow among the receivers? Multicast routing increases network efficiency by using a single shared delivery tree. We address the issue of how these savings are allocated among the various members of the multicast group. We first consider an axiomatic approach to the problem, analyzing the implications of different distributive notions on the resulting allocations. We then consider a “one-pass” mechanism to implement such allocation schemes and investigate the family of allocation schemes such mechanisms can support.

**Index Terms**—Cost allocation, cost sharing, Internet economics, multicast, network accounting, quality of service (QoS).

## I. INTRODUCTION

**B**ECAUSE of the long history of government subsidy in the United States and other countries and the relatively cooperative Internet user population, the data networking research community has paid little attention to the issue of cost allocation. In the United States, however, the Internet has become almost completely commercial (i.e., unsubsidized) and so issues of cost recovery and profit incentives, for better or worse, have become much more relevant.<sup>1</sup> We believe that cost allocation will become a very important issue in this environment.

By cost allocation we mean the assignment to various users of some measure of the network resources they are consuming. This paper discusses cost allocation in the context of multicast flows. The question we discuss is this: when a single data flow is shared among many receivers, how does one *split* the cost of

Manuscript received September 1996; revised January 1997. This work was supported in part by the Advanced Research Projects Agency under Ft. Huachuca, Contract DABT63-91-C-0001, and by National Science Foundation under Grant SES-93-20481. An earlier version of this paper appeared in the *Proc. ACM SIGCOMM'95*.

S. Herzog is with IPHighway, San Jose, CA 95110 USA.

S. Shenker is with Xerox PARC, Palo Alto, CA 94304 USA.

D. Estrin is with the Computer Science Department/ISI, University of Southern California, Los Angeles, CA 90089 USA.

Publisher Item Identifier S 1063-6692(97)08935-8.

<sup>1</sup>In addition, as argued in [4], [5], [24], when data networks offer multiple qualities of service (QoS), which is fast approaching on the horizon, one must consider user incentives to ensure that these QoS features are used appropriately. By appropriate we mean behavior that leads to the efficient use of network resources, where *efficiency* is defined in [5], [24]. These incentive issues are also relevant to controlling usage even in a single QoS environment (see [11], [12]).

that flow among the receivers?<sup>2</sup> The whole point of multicast is to increase network efficiency by sharing a delivery tree. We ask: How are these savings allocated among the various members of the multicast group?

There are many possible ways to allocate costs; we use an axiomatic approach to discriminate among the various possibilities. We first discuss several desirable properties that a cost allocation scheme should have, such as anonymity (essentially, symmetry among group members). Using these properties as a basic set of axioms, we analyze the family of cost allocation functions that are consistent with these axioms. Thus, this axiomatic approach allows us to explore the implications of various notions of fairness and equity in cost allocation. Of course, one must not consider only such idealized properties but must also consider how to implement such a cost allocation strategy. We discuss one approach to implementing cost allocation strategies, called the “one-pass” approach, and analyze to what extent this approach can be made consistent with the basic axioms.

Our investigation is in the context of a very simple network model (we consider generalizations in Section VI). There is a single static distribution tree (i.e., routing and group membership does not change) with a single source and multiple receivers. The network provides a single quality of service to all members of the multicast group. There are costs associated with each link traversed by the flow. These costs could be tied to usage or reservations or both; for the purposes of our model it doesn't matter. The problem is to fully allocate the cost of the flow among the receivers.

This paper has six sections. We start in Section II by clarifying some aspects of our approach and briefly discussing related work. We then present the basic model in Section III and discuss various cost allocation axioms in Section IV. We turn to issues of implementation in Section V where we discuss the ability of our simple “one-pass” approach to satisfy the axioms presented in Section IV. We conclude in Section VI by discussing generalizations of our simple model. The proofs of the various theorems are presented in the Appendix.

## II. BASIC APPROACH AND RELATED WORK

This paper addresses a very specific and technical problem—that of sharing the cost of a multicast flow in the context of a very simple network model. However, the costs we allocate are not necessarily financial. If users are cooperative

<sup>2</sup>Our treatment here only considers a flow with a single source; the generalization to multiple senders is discussed in Section VI.

and they want to make the right decision about the wise use of network resources, they need to be informed about the “cost” their usage imposes on the network. This cost can be in terms of congestion-induced performance degradation imposed on other users and/or in terms of the actual capital and maintenance cost of the network facilities themselves.<sup>3</sup>

Moreover, it is crucial to distinguish between cost allocation and pricing. Pricing is the form and amount of payment extracted from the end user; pricing schemes are often based on many factors besides cost, such as demand elasticity, market structure, and price stability. Cost allocation is how the network assigns, internally to itself, portions of the total cost to various users. There is not necessarily any direct relationship between the allocated cost and the prices charged to individual users.

However, even though not charged directly, these allocated costs can serve as useful input to some pricing schemes. For instance, network users might buy a monthly quota of network service, with the amount debited from this quota based on the allocated costs. The user would have the benefit of extreme price stability (in terms of what they paid every month) but yet would also gain from the increased efficiency of multicast (see [25] for a fuller discussion).<sup>4</sup>

We are making the fundamental assumption that costs are assigned to the receivers of the multicast group and not to the sources. This is for two reasons. First, it is completely trivial to assign a fixed portion of the costs to the source and this changes almost nothing in our analysis except to complicate notation; so for convenience, we do not explicitly treat this case. Second, multicast membership is typically receiver initiated (see [6]) and so responsibility for data flowing over links rests primarily with the receivers. There are some applications where a more sophisticated cost sharing arrangement between sources and receivers is needed, but we consider this to be a higher layer issue and do not address it here.

We now discuss some related work. There is a rapidly growing literature on pricing in computer networks (see [2], [3], [5], [12], [21], [24], [25] for a few representative examples). However, this literature does not typically relate the prices charged to any underlying cost of network usage. The emphasis is on the strategic (utility maximizing) behavior of users, and the profit maximizing behavior of firms, rather than on the equitable distribution of costs that we consider here.

There is a large economics literature on cost allocation (see [19] and references therein, for a brief overview). Cost allocation is usually treated as a special case of a *cooperative game*. A cooperative game is one in which there are  $N$  players and there is some value function  $V$  that assigns a value to each coalition or subset of players. The question addressed by

<sup>3</sup>Calculating the costs for a unicast flow is a difficult problem. How does one quantify congestion costs? How does one relate capital and maintenance costs to the cost assigned to users? We do not address these questions. Instead, no matter how they are calculated, we address the issue of how to split them among the users.

<sup>4</sup>We should also point out that even though costs are allocated to individual users, the recovery of those costs might occur at a much higher level of granularity. For instance, USC might be responsible for all of the costs allocated to its students and staff.




Fig. 1. A simple example of a network.

this literature is how to assign to each player a share such that the sum of the shares equals the value of the complete set (see [18], [20]). Axiomatic approaches are often, in fact almost exclusively, used to analyze such cooperative games. The axioms presented in this paper are very similar to those used in the theory of cooperative games.

Megiddo [13] considered the computational complexity of computing various cost sharing formulas in a family of cooperative games which have the same structure as the problem we consider here. There is also a very extensive literature on minimal cost spanning tree games, which differ from what we consider here in that the routing is always chosen to minimize the total cost. See [23] for a lucid overview of these results and for a much broader discussion of network models in economics. In a more recent paper, Henriot and Moulin [8] considered a network model where costs are incurred for connecting to a central switch and the pairwise traffic matrix is known; the spirit of the analysis is very similar to ours, but the network model is rather different. While there is a rather active literature on these network cost allocation problems, we are not aware of any treatment of the mechanisms needed to implement such cost allocation schemes (see our discussion in Section V). While the presentation in this paper is geared toward the networking research community, we have included several footnotes that relate our work to the economics literature.

### III. THE COST ALLOCATION MODEL

In this section, we first present our abstract model and then illustrate it by discussing several simple examples of cost allocation policies.

#### A. Basic Definitions

In this section, we present an abstract model of the cost-sharing problem. The abstract model defines a formal description of the *network technology* for multicast transmission over a network and the structure of *incurred costs* of such a multicast transmission. The model we describe below is illustrated in Fig. 1 for a simple network scenario.

The *network technology* model has three components: network topology, group membership, and a routing function. *Network topology* describes the physical connectivity of the network in terms of vertices (nodes) and links.  $N = (V, L, T)$  describes a network  $N$ , with nodes  $v_i \in V$ , *directed* links  $(v_i, v_j) \in L$ , and a routing function  $T$ . Let  $M$  describe the set of *multicast group members*; we denote individual

members by  $m_\alpha \in M$ . Each packet sent to this multicast group by any source, not necessarily a member of the group, will go to all members  $M$  of the group. We will use the terms *member* and *receiver* somewhat interchangeably. The mapping of individual members to specific network nodes is described by the location function  $loc: M \mapsto V$ . The *routing function* abstraction describes the path taken by packets when a sender transmits to the group. In *unicast routing*, the routing function  $L' = T(N, R, v_i)$  takes a network  $N$  along with a source (root)  $R$  and a receiver  $v_i$  and defines the set of directed links that establish the path between  $R$  and  $v_i$ . While we make no particular assumptions about the optimality of this unicast route, we do assume that the set of unicast routes from a single root  $R$  to all other nodes forms a tree.<sup>5</sup> In this paper, we assume that multicast routing is based on loop-free unicast routing; that is, multicast routing creates a source rooted distribution tree created from the union of the unicast paths between the root and the set of receivers:  $T(N, R, V') = \cup_{v_j \in V'} T(N, R, v_j)$  for all subsets  $V' \subseteq V$ .

We now turn to the structure of incurred costs. Our abstract model considers the allocation of costs for each data flow individually where data flow refers to a particular source sending to a particular multicast group destination.

For reasons of simplicity, we associate all costs with the links that the packets traverse.<sup>6</sup> While these costs could arise from many different aspects of network usage (e.g., bandwidth, buffer usage, reservations), we lump all of these different costs into a single quantity. Because we assume that the same quality of service is delivered to all receivers, the link costs are independent of which receivers, or how many, are downstream.<sup>7</sup> Later, in Section VI, we discuss generalizations of the model to incorporate multiple qualities of service.

Moreover, we should note that our abstract model is intended to address the problem of cost allocation among a given set of group members with a given network topology; this paper does not explicitly address dynamics, either in membership or in topology. However, one can apply our approach to allocating cost for time slices and presume sufficiently static behavior within those time slices to render the resulting cost shares a reasonable approximation.

The set of link costs for a given network  $N = (V, L, T)$  is defined as  $c: L \mapsto \mathfrak{R}_+$  where we use the symbol  $\mathfrak{R}_+$  to denote the nonnegative reals. The *cost function* of a distribution tree comprised of links  $L'$  is merely the sum of all the directed links in  $L'$ :  $cf(L', c) = \sum_{(v_i, v_j) \in L'} c(v_i, v_j)$ . In this paper,

<sup>5</sup>This assumption follows if whenever the route from  $R$  to node  $v_i$  passes through node  $v_j$ , the route from  $R$  to  $v_j$  is a subset of the route from  $R$  to  $v_i$ . We also are assuming that routing is only a function of the location of the receiver. Some routing protocols allow receivers to request different quality of routes, which would violate our assumption.

<sup>6</sup>Although most reservation protocols reserve link resources and not intranode resources, we acknowledge that some intranode resources may be costly (like internal buffers, CPU cycles, etc.). Intranode costs can be attributed to links by ascribing the intranode cost associated with processing incoming data to the incoming link and the overhead for outgoing data to the appropriate outgoing links.

<sup>7</sup>The downstream portion of a tree relative to a link contains all the nodes (and members) that incorporate this link in their unicast path from the root. The upstream portion is merely defined as the portion of the tree that is not downstream.

we explore how to allocate this cost among the individual members.

We denote a *cost allocation function* by  $af(N, R, M, loc, c)$ ; this represents a specific cost allocation strategy. Given a network technology defined by  $N = (V, L, T)$ , a group membership defined by  $M$  and  $loc$  and a cost structure defined by the link costs  $c$ , the allocation function  $af$  defines the cost that will be individually assigned to members of the multicast group;  $af_\alpha(N, R, M, loc, c)$  denotes the allocation to member  $\alpha$ .

For any distribution there may be numerous ways of allocating costs to receivers and our focus in this paper is on providing a rationale for discriminating among these allocation policies. We restrict ourselves to policies that fully allocate the costs incurred. We require that

$$\sum_{m_\alpha \in M} af_\alpha(N, R, M, loc, c) = cf\{T[N, R, loc(M)], c\} \\ \forall N, R, M, loc, \text{ and } c.$$

We call this equality the *balanced budget* condition. Moreover, we restrict ourselves to allocating *costs*, which means that all allocations are nonnegative (i.e., no receiver is *paid* for using the network)

$$af_\alpha(N, R, M, loc, c) \geq 0 \quad \forall \alpha, N, R, M, loc, \text{ and } c.$$

## B. Examples

The previous section described our abstract model of cost allocation and here we present a few examples of allocation strategies. We use the sample problem depicted in Fig. 1 to illustrate these example of cost allocation strategies.

The simplest approach to allocating costs is to merely divide the total cost equally among all receivers; we call this the equal tree split (ETS) scheme. In the example depicted in Fig. 1, all members are allocated the cost  $[c(v_1, v_2) + c(v_2, v_3) + c(v_2, v_4) + c(v_1, v_5)]/9$ .

The ETS policy does not discriminate between those receivers far from the source and those close to the source and thus does not attempt to hold receivers accountable for the costs their individual membership incurs. However, the cost of a particular link is incurred because there is at least one downstream receiver and so while all downstream receivers can be considered equally responsible for the cost, all other receivers are not responsible at all. This leads to a different approach to allocating costs, one where the cost of each link is split equally among only the downstream receivers. We call this the equal link split among downstream members (ELSD) scheme.<sup>8</sup> For instance, in the example above, the cost allocated to  $m1$  is zero, since she is not downstream of any links, and the cost allocated to  $m4$  is given by  $c(v_1, v_2)/7 + c(v_2, v_3)/3$ .

Of course, all multicast cost allocation policies should be compared to the unicast costs. The total cost of separate unicast transmissions is higher than the total cost of the multicast transmission and different allocation policies distribute this

<sup>8</sup>For those familiar with the concept, this is the Shapley value [22] of this problem.

surplus differently. Unicast is not a real cost allocation strategy (in the sense of our model) since it allocates to members more than the network costs; however, it does provide a useful point of comparison.<sup>9</sup> The unicast cost of member  $\alpha$  is given by  $\sum_{(v_i, v_j) \in T[N, R, \text{loc}(m_\alpha)]} c(v_i, v_j)$ . For the example in Fig. 1, the unicast cost of  $m_4$  is  $c(v_1, v_2) + c(v_2, v_3)$ .

#### IV. AXIOMATIC ANALYSIS

We analyze this abstract model of cost sharing through the use of axioms. We use, as axioms, properties that desirable cost-sharing formula might have and then investigate the implications of assuming these properties. What additional properties are implied? What forms of cost-sharing rules are allowed?

We first, in Section IV-A, identify three *basic* axioms that describe the process of allocating cost. These axioms describe which aspects of the problem are relevant to the cost-sharing formula. We assume that the basic axioms apply throughout the rest of Section IV. We discuss the implications of these basic axioms in Section IV-B and then present a few examples cost allocation policies which satisfy them in Section IV-C. In Section IV-D, we identify a few additional axioms that express different policy objectives of cost-sharing and we explore the implications of each of these axioms when combined with the three basic axioms.

##### A. Basic Axioms

The three basic axioms describe aspects of the cost sharing problem that should be irrelevant to the actual cost sharing formula.

We start by observing that a cost-sharing formula should be invariant under arbitrary relabeling of equivalent members; it should not have an intrinsic and arbitrary asymmetry built into it.<sup>10</sup>

*Axiom 1—Anonymity:*<sup>11</sup> Name labels, by which members are identified, are irrelevant to cost allocation.

*Formal definition:* Given any network  $N = (V, L, T)$ , root  $R$ , link cost function  $c$ , member set  $M$ , two placement functions  $\text{loc}$  and  $\text{loc}'$  and any two members  $m_\alpha$  and  $m_\beta \in M$  **if**

$$\begin{aligned} \text{loc}(m_\alpha) &= \text{loc}'(m_\beta) & \text{and} \\ \text{loc}(m_\beta) &= \text{loc}'(m_\alpha) & \text{and} \\ \text{loc}(m_\gamma) &= \text{loc}'(m_\gamma) & \forall m_\gamma \in M - m_\alpha - m_\beta \end{aligned}$$

<sup>9</sup>Note that if the allocated multicast cost is greater or equal to the unicast cost there may be no incentive to use multicast.

<sup>10</sup>This does not mean that there cannot be systematic asymmetries (e.g., charging professors more than students on a campus network), but these asymmetries should be described as a nonequivalency of members (e.g., introducing two classes of members into the formalism).

<sup>11</sup>Notice that this axiom implies that two members located at the same node must be allocated the same cost:  $\text{loc}(m_\alpha) = \text{loc}(m_\beta) \implies af_\alpha(\dots) = af_\beta(\dots)$ .

**then**

$$\begin{aligned} af_\alpha(N, R, M, \text{loc}, c) &= af_\beta(N, R, M, \text{loc}', c) & \text{and} \\ af_\beta(N, R, M, \text{loc}, c) &= af_\alpha(N, R, M, \text{loc}', c) & \text{and} \\ af_\gamma(N, R, M, \text{loc}, c) &= af_\gamma(N, R, M, \text{loc}', c) \\ & \forall m_\gamma \in M - m_\alpha - m_\beta. \end{aligned}$$

Link costs can come from many sources such as per-packet costs, per-hour costs, and per-connection costs. They can also come from link-transmission costs, buffering costs, etc. In order to avoid the additional complexity and overhead of independently allocating all such costs, the cost allocation should not depend on whether the accounting method allocates these costs separately or jointly. That is, it should not matter if the various components of the link costs are combined into one single cost to be allocated or if they are separately allocated. This is expressed in the following axiom.

*Axiom 2—Additivity:* Given any two sets of link costs, the sum of their respective cost allocation functions is equal to the cost allocation function of the sum of the two sets.

*Formal definition:* Given any network  $N = (V, L, T)$ , root  $R$ , member set  $M$ , placement function  $\text{loc}$ , and any two link cost functions  $c_1$  and  $c_2$

$$\begin{aligned} af(N, R, M, \text{loc}, c_1 + c_2) &= \\ af(N, R, M, \text{loc}, c_1) + af(N, R, M, \text{loc}, c_2). \end{aligned}$$

The cost allocation formula must depend on the underlying network topology. However, it is preferable for this topological dependency to be restricted to factors related to cost. For instance, if we took a single link and artificially broke it into two links (spreading the cost of the link between the two links), the cost allocation should not change. A more general form of this intention is that the cost allocation should only depend on the set of costs incurred by each subgroup. If two networks incur the same costs for every subgroup of members, then they should allocate the cost in the same way. This condition, that two networks incur the same costs for serving each subgroup of members, is actually quite strong; in particular, every unicast cost must be the same, every cost for a pair of receivers must be the same, etc. Any changes in the network that preserve this property might well be considered irrelevant to cost allocation.<sup>12</sup>

*Axiom 3—Equivalency:* Consider two networks with the same single group of members; if the cost of serving any subgroup of members is identical in both networks, then the allocated costs must be the same.

*Formal definition:* Given a single set of members  $M$  and two different network scenarios  $N1, R1, \text{loc}1, c1$ , and  $N2, R2, \text{loc}2, c2$

**if**

$$\begin{aligned} cf\{T[N1, R1, \text{loc}1(M')], c1\} &= \\ cf\{T[N2, R2, \text{loc}2(M')], c2\} & \quad \forall M' \subseteq M \end{aligned}$$

<sup>12</sup>This axiom implies that the only relevant aspect of the problem is the induced cooperative game; all other aspects of the topology are irrelevant.

then

$$\begin{aligned} af(N1, R1, M, loc1, c1) = \\ af(N2, R2, M, loc2, c2). \end{aligned}$$

### B. The Canonical Form of Cost Allocation Formulas

The three basic axioms presented in the previous section greatly reduce the scope of allowable cost-allocation policies. In this section, we show that all cost-allocation policies satisfying these three basic axioms can be expressed in a very simple canonical form. We consider functions  $F: \{Z_+^2 - (0, 0)\} \mapsto \mathbb{R}_+^2$ , where we use the symbol  $Z_+$  to denote the nonnegative integers. We define the family of functions  $\mathcal{F}$  to be those functions  $F: \{Z_+^2 - (0, 0)\} \mapsto \mathbb{R}_+^2$  that satisfy the properties<sup>13</sup> 1)  $z_1 F_1(z_1, z_2) + z_2 F_2(z_1, z_2) = 1$  and 2)  $F_1(0, z) = F_2(z, 0) = 1$ . Note that the set  $\mathcal{F}$  is convex:  $F, G \in \mathcal{F} \implies [\mu F + (1 - \mu)G] \in \mathcal{F}$  for all  $0 \leq \mu \leq 1$ . Also, we know that  $F_2(0, z) = F_1(z, 0) = 1/z$  for  $z > 0$ .<sup>14</sup>

We now use these functions to define cost-allocation strategies. Consider a link  $(v_i, v_j)$  with a cost  $c(v_i, v_j)$  where there are  $n_d > 0$  receivers downstream<sup>15</sup> and  $n_u \geq 0$  receivers upstream. We use a function  $F = (F_u, F_d) \in \mathcal{F}$  to determine the fraction of the link cost that is allocated to each upstream or downstream receiver. Members are allocated the following costs:

$$\text{For each upstream member: } F_u(n_u, n_d) * c(v_i, v_j)$$

$$\text{For each downstream member: } F_d(n_u, n_d) * c(v_i, v_j).$$

The total cost allocated to a member is the sum, over all links, of its share of each link cost.

Cost-allocation strategies that can be expressed in terms of a function  $F \in \mathcal{F}$  are called *canonical* strategies. Note that the first condition above in the definition of  $\mathcal{F}$  ensures that costs are fully allocated:  $n_u F_u(n_u, n_d) + n_d F_d(n_u, n_d) = 1$ . It is fairly clear that all canonical strategies satisfy the three basic axioms. More interestingly, all cost-allocation strategies satisfying the basic axioms are canonical.

*Theorem 1:* A cost allocation formula satisfies the basic axioms if and only if it is a canonical strategy.

### C. Examples of Canonical Cost Allocation Strategies

In this section, we present several examples of functions  $F \in \mathcal{F}$ . Two of the functions were already introduced in Section III-B. Since we know that  $F_d(0, n_d) = 1/n_d$  for all  $F \in \mathcal{F}$ , we merely describe the functions on the set  $Z_{++}^2$  below (where  $Z_{++}$  denotes the positive integers). Since the set  $\mathcal{F}$  is convex, all linear combinations of the examples below are also in  $\mathcal{F}$ .

<sup>13</sup>The second condition in the definition of  $\mathcal{F}$  is unrelated to allocations, but is merely used to simplify the expression of Theorem 11.

<sup>14</sup>For our use in this section we need only have defined  $F$  on  $Z_+ \times Z_{++}$  (where  $Z_{++}$  denotes the positive integers) but we will make use of the family of functions again in Section V where we need the larger domain of  $\{Z_+^2 - (0, 0)\}$ .

<sup>15</sup>If there are no receivers downstream then this link is not part of the distribution tree and we can ignore its costs.

#### 1) Equal Link Split Downstream (ELSD)

$$F_u(n_u, n_d) = 0; \quad F_d(n_u, n_d) = \frac{1}{n_d}.$$

#### 2) Equal Tree Split (ETS)

$$F_u(n_u, n_d) = F_d(n_u, n_d) = \frac{1}{n_u + n_d}.$$

#### 3) Cost is charged to upstream receivers

$$F_u(n_u, n_d) = \frac{1}{n_u}; \quad F_d(n_u, n_d) = 0.$$

#### 4) Cost is charged relative to the number of receivers upstream/downstream

$$F_u(n_u, n_d) = \frac{n_u}{n_u^2 + n_d^2}; \quad F_d(n_u, n_d) = \frac{n_d}{n_u^2 + n_d^2}.$$

### D. Additional Axioms

The basic axioms address the issue of what factors can be considered relevant to cost allocation. These axioms have substantial reductive power in that they define a class of canonical cost allocation strategies. However, as the examples above show, one can allocate all costs to upstream nodes, to downstream nodes, anywhere in between. Thus, this family of canonical cost allocation strategies incorporates a wide variety of distributive notions. We use the phrase *distributive notion* to mean standards of equity or justice that allows one to discriminate between allocation policies. Our next step is to examine some additional axioms that express particular distributive notions. These axioms can be used to select a subset of canonical allocation strategies.

*Stand-Alone and Related Axioms:* One such distributive notion is that a member's cost should reflect the benefits of multicast. Just as the total network cost of a multicast flow is less than the sum of the costs of unicast flows to each member, one might require that each individual allocated cost in a multicast flow never be greater than the cost incurred by the corresponding unicast flow. This yields the following axiom.

*Axiom 4. Stand Alone:* The unicast cost of a member is an upper limit on her cost allocation.

*Formal definition:*

$$\begin{aligned} af_\alpha(N, R, M, loc, c) \leq cf\{T[N, R, loc(m_\alpha)], c\} \\ \forall N, R, M, loc, c, m_\alpha \in M. \end{aligned}$$

The stand-alone axiom protects the individual; every individual receiver is guaranteed that joining a group can never cost more than her unicast cost. Assuming users have the power of choice in their network activities and assuming self-interest guides them, users would not want to join a shared group of receivers if they risk an increase in their own allocated costs.<sup>16</sup> Insisting upon the stand-alone axiom when

<sup>16</sup>Of course, in a cooperative environment receivers may choose to risk having increased costs if the total cost distributed to the group decreases.

combined with the basic axioms means that there is one and only one applicable cost allocation strategy.<sup>17</sup>

*Theorem 2:* A cost allocation function satisfies the basic and Stand-Alone axioms if and only if it is the Equal Link Split Downstream (ELSD) function.

A stronger form of the Stand-Alone axiom is the “Sharing-is-Good” axiom. This axiom embodies the distributive notion that sharing a multicast tree with more members always benefits everybody.

*Axiom 5. Sharing-is-Good:* The cost allocated to a member never increases when another member joins.

*Formal definition:*

$$af_{\alpha}(N, R, M + m_{\gamma}, \text{loc}, c) \leq af_{\alpha}(N, R, M, \text{loc}, c) \\ \forall N, R, M, \text{loc}, c, m_{\gamma}, \text{ and } m_{\alpha} \in M.$$

The ELSD scheme satisfies the sharing-is-good axiom, since the share of costs from each link strictly decreases with the number of downstream members (and is independent of the number of upstream members). Clearly any cost allocation scheme obeying the sharing-is-good axiom also obeys the stand-alone axiom. These two axioms both describe an upper bound on the cost that can be allocated to a particular member. However, we might also be concerned about the problem of *free riders*, who are members who do not pay their fair share. According to the stand-alone axiom, the most a member should pay is her unicast cost, and the sharing-is-good axiom requires that the allocations decrease as members join. How much can a member benefit without being a free rider? If all members are located at the same node, then they all pay  $1/|M|$ th of their unicast cost. We suggest that any member paying less than this should be considered a free rider.<sup>18</sup>

*Axiom 6— No-Free-Rider:* The cost allocated to a member is never less than  $1/|M|$ th of her unicast cost.

*Formal definition:*

$$|M| * af_{\alpha}(N, R, M, \text{loc}, c) \geq cf\{T[N, R, \text{loc}(m_{\alpha})], c\} \\ \forall N, R, M, \text{loc}, c, \text{ and } m_{\alpha} \in M.$$

Eliminating free riders does not pick out a specific allocation scheme, but does narrow the range of possibilities.

<sup>17</sup>This result is closely related to the standard axiomatization of the Shapley value (see [20], [22]) in economics. Members attached at the root can be considered *dummy members* because adding them to a group does not increase the total cost incurred. We can define a *dummy member* axiom that says that no member located at the root can be allocated a nonzero cost

$$\text{loc}(m_{\alpha}) = R \implies af_{\alpha}(N, R, M, \text{loc}, c) = 0 \\ \forall N, R, M, \text{loc}, c \text{ and } m_{\alpha} \in M.$$

Theorem 2 continues to hold if we replace the stand-alone axiom with the much weaker dummy member axiom. The equivalency axiom means that only the cooperative game matters (i.e., topology is irrelevant aside from the cooperative game it induces). The basic result due to Shapley is that there is one and only one budget balanced cost allocation formula satisfying the additivity, anonymity, and dummy axioms, and this formula is now known as the Shapley value.

<sup>18</sup>This is much like the unanimity bound in economics (see [14], [15]).

*Theorem 3:* A cost allocation function satisfying the basic axioms and the no-free-rider axiom must satisfy

$$F_u(n_u, n_d) \leq F_d(n_u, n_d) \quad \forall (n_u, n_d) \in Z_{++} \times Z_+.$$

*Subset Monotonicity* The essential guiding principle behind the equivalency axiom is that the cost allocations should depend only on the costs incurred by the various subsets or coalitions of members. Another distributive notion that arises from this principle is that the cost allocated to a particular member should be monotonic with respect to these subset costs. More precisely, if we consider two cost structures  $c1$  and  $c2$  (that is, we consider the network  $N$  and the members  $M$  fixed and we merely consider two sets of link costs), then if for every subset  $M' \subseteq M$  the total cost of serving  $M'$  is no greater under cost  $c1$  (compared to  $c2$ ), then one might require that the allocated costs under  $c1$  would not be greater than those under  $c2$ . This yields the following axiom.

*Axiom 7— Subset Monotonicity:* No cost allocation can increase when subset costs all decrease or stay the same.

*Formal definition:* Consider an arbitrary tree

$$L' = T[N, R, \text{loc}(M)] \text{ and two link costs } c1 \text{ and } c2 :$$

**if**

$$cf\{T[N, R, \text{loc}(M')], c1\} \geq cf\{T[N, R, \text{loc}(M')], c2\} \\ \forall M' \subseteq M$$

**then**

$$af_{\alpha}(N, R, M, \text{loc}, c1) \geq af_{\alpha}(N, R, M, \text{loc}, c2) \\ \forall m_{\alpha} \in M.$$

It turns out that this axiom, when combined with the basic axioms, determines a unique cost allocation strategy; in fact, only two of the three basic axioms are needed for this result.

*Theorem 4:* A cost allocation formula satisfies equivalency, anonymity, and subset monotonicity if and only if it is the equal tree split (ETS) formula.

### Collusion Prevention

Another aspect of allocation that is important to consider is collusion. Whenever a cost is shared among clients, it may be possible for several clients to unite and be represented by a single client who then forward the data on to them. This is analogous to the classic “copy and distribute” security problem. Collusion among some receivers may increase the cost allocated to the other receivers and may decrease the efficiency of sharing the transmission. We would prefer that a cost allocation scheme not encourage collusion among the members. We, therefore, propose the following axiom.

*Axiom 8— Collusion Prevention:* The cost allocation scheme does not yield benefits for colluding members.

*Formal definition:* Consider an arbitrary network  $N = (V, L, T)$ , a root  $R$ , a set of link costs  $c$ , a set of members  $M$  and their location function  $\text{loc}$ . For each subset  $M' \subseteq M$  and  $m_\gamma \in M'$ :

$$\sum_{m_\alpha \in M'} af_\alpha(N, R, M, \text{loc}, c) \leq af_\gamma(N, R, M - M' + m_\gamma, \text{loc}, c) + cf\{T[N, \text{loc}(m_\gamma), \text{loc}(M' - m_\gamma)], c\}.$$

Obviously, collusion prevention is a desirable property for cost allocation formulas. Unfortunately, we can prove that no allocation strategy satisfies this axiom.<sup>19</sup>

*Theorem 5:* No cost allocation formula can satisfy the collusion prevention axiom.

### E. Discussion

We started this section with three basic axioms, which narrowed the space of cost allocation strategies to the canonical ones. Within this class we discussed how various distributive notions pointed toward different choices. Eliminating free riders restricts us to schemes that allocate more to downstream members than to upstream members. Subset monotonicity leads us to the ETS scheme, while the stand-alone axiom suggests the ELSD scheme. Choosing between axioms is purely subjective, but charging a nonzero amount for a member located at the root seems generally inappropriate. The only canonical allocation scheme that always allocates zero to members located at the root is the ELSD scheme, and so perhaps it is the most natural choice.

Our treatment here is completely static. Consider, for a moment, the dynamic policy of allocating to each member the incremental cost of adding them to the distribution tree. The resulting allocations depend on the order in which members joined, which seems rather unfair. It might seem appropriate to then average these incremental costs over all arrival orders. Indeed, this averaging produces the ELSD allocations.<sup>20</sup>

In this section, we discussed various cost allocations schemes from an axiomatic perspective. The next section discusses mechanisms for implementing these cost allocation schemes.

## V. ACCOUNTING MECHANISMS

In this section, we use the term *accounting schemes* to denote mechanisms for implementing cost allocation schemes. First, we discuss the general structure of one class of accounting schemes we call *one-pass* schemes. We then describe two different models of implementation of such schemes, models 1 and 2, which differ in the information provided about downstream members. Each of these models is examined according to the central question raised throughout this section: What forms of cost allocation schemes can this family of accounting mechanisms support?

<sup>19</sup>Note that this result applies for all cost allocation formulas, not just those obeying the basic axioms.

<sup>20</sup>This is the standard motivation for the Shapley value.

### A. One-Pass Accounting Schemes

We have made the fundamental assumption that costs are allocated among receivers and not just assigned to the sender. Because the number of receivers can become quite large and widely dispersed geographically, the key concern in designing accounting mechanisms is scalability. It is important that the traffic load imposed by the accounting mechanism on any particular link should not increase (without bound) with the size (in terms of numbers or in geographical dispersion) of the multicast group. Thus, scalability concerns rule out any form of centralized accounting.

In this paper, we consider only the family of *one-pass* mechanisms whereby the accounting control messages make a single pass from the source down the multicast tree to all receivers. While this is not the only scalable accounting mechanism one might imagine,<sup>21</sup> it certainly seems among the most natural. In this one-pass method of accounting, nodes allocate costs to members as the *accounting message* passes through them. The information used to make the allocation decisions comes from two sources. The first source of information is multicast routing (and perhaps the reservation establishment protocol if the costs are related to reservations), which provides information about the downstream links. Traditional multicast routing only provides information about whether or not there exist members downstream of a particular link. We call this *model 1*. It is possible, perhaps only because it enables better allocation of costs, that multicast routing could provide the exact number of members downstream of each link. We call this *model 2*.<sup>22</sup> If the link costs are tied to reservations then this information about the number of downstream members making reservations could be provided by the reservation establishment protocol. We will consider both models in what follows. As we shall see, there is an important difference in the functionality that can be achieved in the two models.

The second source of information is the accounting message itself. The design of a one-pass accounting mechanism essentially reduces to the question of what information is carried in the accounting message sent downstream. To ensure scalability, this information cannot grow with the size of the multicast group, nor with the number of links traversed. This prevents us from carrying detailed information about each upstream link cost and each member's allocated cost. Instead, we choose to carry only a single piece of information in the accounting message: the unallocated or residual cost passed down from the upstream node. While this is not the most general form of accounting message it seems a natural and simple choice.

Thus, with this form of accounting message, the costs are allocated with the following process. The accounting message

<sup>21</sup>For instance, one could design an accounting method that involved two passes of control messages, one downstream and the other back upstream. Also, one could use an iterative accounting method where the control messages continued to circulate until an equilibrium had been reached.

<sup>22</sup>To do any form of cost allocation we must be able to identify all members at a node, if for no other reason than to give them the feedback about their allocated costs. The question is then which protocol will carry these numbers upstream. It appears to be simple to modify most multicast routing protocols to carry the cumulative membership numbers upstream once the number of local members is available.

arrives at a node on the incoming link from an upstream neighbor, carrying the upstream residual cost; we will call this the *input cost* to the node and let  $\text{in}(v_j)$  denote the input cost arriving at the downstream node  $v_j$ . The cost allocation function determines how much of this cost is allocated to each of the local members and how much is passed down to each of the downstream next-hops. We will call the costs that are passed on the *residual output* from a node, and let  $\text{out}(v_i, v_j)$  denote the costs that are passed on from node  $v_i$  to downstream node  $v_j$ . The sum of all residual outputs plus the sum of all locally allocated costs must be equal to input cost as a result of the balanced budget rule. When an accounting message is forwarded to the downstream neighbor, the cost of the link connecting the two nodes is added to the residual costs and this sum is carried in the accounting message as input costs to the next-hop node: thus,  $\text{in}(v_j) = c(v_i, v_j) + \text{out}(v_i, v_j)$  when  $v_j$  is a next-hop downstream of  $v_i$ . At leaf nodes, all costs are allocated to the local members. At nodes with no local members, all costs are passed down to the downstream next-hops in the accounting message.

One-pass accounting is a distributed accounting scheme. Independent cost allocation decisions are made by each individual node based on the information provided to it by multicast routing (either model 1 or model 2) and the accounting message. We assume that no other information about topology or group members can be factored into the allocation decision. For model 1, since a node can make no meaningful distinctions between downstream links, we require that the residual costs passed on to each next-hop are the same:  $\text{out}(v_i, v_j) = \text{out}(v_i, v_k)$  for all links  $(v_i, v_j)$  and  $(v_i, v_k)$  in the distribution tree. We further assume, in both model 1 and model 2, that all nodes must implement the same allocation rules. In order to achieve a consistent allocation scheme, all nodes, if given the same information, must produce the same allocation.

We will refer to cost allocation schemes that can be implemented with a model 1 one-pass accounting mechanism as model 1 allocation schemes, and similarly for model 2. The basic one-pass structure of cost accounting imposes some significant restrictions on what cost allocation formulae can be supported. In particular, such one-pass accounting schemes can only support cost allocation schemes that satisfy the stand-alone axiom.<sup>23</sup>

*Theorem 6:* For any one-pass accounting mechanism, no member can be allocated a cost greater than her unicast cost.

We are interested in how many of our original axioms are consistent with our one-pass family of accounting mechanisms. Before we consider models 1 and 2 separately, we can rule out one axiom that does not apply to either of them.

*Theorem 7:* No cost allocation scheme implemented with a one-pass accounting mechanism can satisfy the subset monotonicity axiom.

<sup>23</sup>A stronger result holds. No subset  $M' \subseteq M$  can be allocated a cost that is greater than their subtree cost  $\sum_{m_\alpha \in M'} af_\alpha(N, R, M, \text{loc}, c) \leq cf\{T[N, R, \text{loc}(M')], c\} \forall N, R, M, \text{loc}, M' \subseteq M$ , and  $c$ . This means that all one-pass accounting schemes produce results that are in the *core* (see [23] for a definition).

We first discuss model 2, because our treatment of it is closer to our previously developed results. We then will return to model 1, where we will need to modify our basic set of axioms.

### B. Model 2

As we stated above, we are interested in the extent to which these accounting mechanisms can support cost allocation formulas that obey our previous axioms. As we show below, model 2 can support the basic axioms presented in Section IV-A. In fact, there is only one model 2 cost allocation formula that obeys the basic axioms.

*Theorem 8:* ELSD is the only cost allocation formula that obeys the basic axioms and that can be implemented with a model 2 one-pass accounting mechanism.

This follows trivially from Theorems 2 and 6. How is the ELSD formula implemented? Consider some node  $v_i$  in the distribution tree; we let  $\text{tmem}(v_i)$  denote the number of members located in the subtree rooted at  $v_i$  and recall that  $\text{in}(v_i)$  denotes the input costs in the accounting message arriving at  $v_i$ .<sup>24</sup> The cost allocated to each local member is  $\text{in}(v_i)/\text{tmem}(v_i)$ . For each link  $(v_i, v_j)$  in the distribution tree,  $\text{out}(v_i, v_j) = \text{in}(v_i) * \text{tmem}(v_j)/\text{tmem}(v_i)$  and so  $\text{in}(v_j) = \text{in}(v_i) * \text{tmem}(v_j)/\text{tmem}(v_i) + c(v_i, v_j)$ . Notice that the residual costs passed on to next-hops is proportional to the number of receivers downstream. The cost of the connecting link is passed on fully to the downstream next-hop.

### C. Model 1

We now consider accounting mechanisms in the context of model 1, where multicast routing only indicates which links have downstream members but not how many members are there.

1) *Reduced Basic Axioms:* We would like to invoke the basic axioms presented in Section IV-A. Additive and anonymous cost-allocation schemes can be supported (as we shall see in examples) by model 1 accounting schemes. However, we find that the equivalency axiom is not consistent with model 1.

*Theorem 9:* No cost allocation formula that is implemented by a model 1 one-pass accounting scheme can satisfy the equivalency axiom.

Thus, model 1 accounting schemes are necessarily dependent on the physical topology, in contrast with the topological independence of the equivalency axiom. For model 1 we will still require that the *reduced* basic axioms of anonymity and additivity still hold, but must relinquish the equivalency axiom. In the next section, we develop a canonical form for model 1 allocation schemes that obey these two axioms.

2) *The One-Pass Canonical Form:* Consider any model 1 allocation scheme that obeys the reduced basic axioms. At leaf nodes, all costs must be allocated equally to local members.

<sup>24</sup>A node  $v_i$  knows the value of  $\text{tmem}(v_j)$  for all immediate downstream members because this is a basic property of model 2. The node can calculate  $\text{tmem}(v_i)$  by adding up all the  $\text{tmem}(v_j)$  for all immediate downstream members and then adding the number of local members.

TABLE I  
EXAMPLES OF MODEL 1 COST ALLOCATION STRATEGIES

Allocation Strategy	$F_l(n_l, n_r)$	$F_r(n_l, n_r)$
Local members pay nothing	0	$\frac{1}{n_r}$
Local members pay all	$\frac{1}{n_l}$	0
All locals are considered as one next hop (ENHS)	$\frac{1}{n_l(n_r + 1)}$	$\frac{1}{n_r + 1}$
Local members and next-hops are allocated identical costs	$\frac{1}{n_l + n_r}$	$\frac{1}{n_l + n_r}$
Majority Loses	$n_l > n_r$	$\frac{1}{2n_l}$
	$n_l < n_r$	0
	$n_l = n_r$	$\frac{1}{2n_l}$

At nodes with no local members, all costs must be passed on equally to all downstream links. Thus, the only design freedom left is how much of the residual cost to allocate to the local members and how much to pass on to downstream members when both are present. We can express the family of possible design choices with the *one-pass canonical form* of cost allocation formulae (supporting the model 1 one-pass accounting scheme).

Each one-pass canonical form is associated with a function  $F = (F_l, F_r) \in \mathcal{F}$  in the following way. Consider a node  $v_i$  with an input cost of  $\text{in}(v_i)$ . If there are  $n_l$  local receivers and  $n_r$  next hop nodes, the allocated costs are

$$\begin{aligned} \text{For a local member on } v_i: & F_l(n_l, n_r) * \text{in}(v_i) \\ \text{For a next hop node } v_j: & \text{out}(v_i, v_j) = F_r(n_l, n_r) * \text{in}(v_i) \end{aligned}$$

All model 1 allocation schemes obeying the reduced basic axioms can be expressed in this form.

*Theorem 10:* A model 1 cost-allocation formula satisfies the reduced basic axioms if and only if it can be expressed in the one-pass canonical form.

Table I lists a few examples of possible allocation strategies expressed in the one-pass canonical form. While the one-pass canonical form appears very similar to the canonical form discussed in Section IV-B, there are important differences. The previously discussed canonical form expressed how the cost of a particular link was allocated to all upstream and downstream members. Here, the canonical form only describes the allocation to local members and to downstream links. Since we know that  $F_l(n_l, 0) = 1/n_l$  and  $F_r(0, n_r) = 1/n_r$  for all  $n_l, n_r > 0$ , this table only describes the allocations on the set  $Z_{++}^2$ . To find the resulting allocation to all members we must recursively iterate this formula down the tree. Thus, it is much harder to understand what allocations will result from a particular one-pass canonical form.

3) *Additional Axioms:* We already know, from Theorem 6, that all the one-pass allocation schemes must satisfy the stand-alone axiom. However, not all such cost allocation schemes obey the sharing-is-good axiom.

*Theorem 11:* A model 1 cost allocation formula satisfies the reduced basic axioms and the sharing-is-good axiom if and only if the functions  $F_l(n_l, n_r)$  and  $F_r(n_l, n_r)$  are nonincreasing on  $\{Z_{++}^2 - (0, 0)\}$ .

All model 1 allocation schemes obeying this mild restriction ensure that the benefits of multicast are shared among all receivers. Note that of our examples in Table I, only the last one fails this test.

Recall that in Theorem 3 a wide variety of canonical cost allocation policies satisfy the no-free-rider axiom. In particular, the ELSD scheme, which is the only model 2 allocation policy which obeys the basic axioms, satisfies the no-free-rider axiom. However, no model 1 allocation policies satisfy the no-free-rider axiom.

*Theorem 12:* There is no model 1 cost allocation formula that satisfies the reduced basic axioms and the no-free-rider axiom.

#### D. Discussion

In Section IV, discussed a general axiomatization of allocation policies. The ELSD scheme emerged as the most attractive scheme. We then turned, in this section, to issues of implementation. We discussed two different models in this section. Model 2 can implement the ELSD scheme, and in fact this is the only model 2 scheme consistent with the basic axioms. So, if we adopt model 2 in our implementations, there seems little question that ELSD would be the most appropriate allocation policy.

However, when we use model 1 we are faced with a much more confusing situation. We cannot achieve the desired degree of topological independence, nor can we prevent free riders. There are few distributive notions besides sharing-is-good and stand-alone that we can achieve. The ENHS scheme has the property that its allocations depend only on the set of nodes at which there are receivers; the allocations are not dependent of the number of receivers at the other nodes. While this is hardly a ringing endorsement, perhaps ENHS is the best of this rather sorry lot of allocation policies.

How to choose between these two models? From the perspective of the accounting protocol needed to realize them, the difference in implementations between models 1 and 2 is very small. The key difference between the two models is in the availability of the exact number of local members; once that number is available, it seems fairly straightforward to provide it upstream either through multicast routing directly or through the reservation establishment protocol, or even through a separate set of accounting control messages.<sup>25</sup> If costs are tied to reservations, then model 2 is quite practical since the number of local members is already known to receiver-initiated reservation establishment protocols. However, most multicast routing protocols do not determine the number of lo-

<sup>25</sup> If the number of local members is available, but multicast routing does not propagate these numbers upstream, then one could add such a function directly to the accounting mechanism and this would necessitate a second set of accounting control messages traversing up the distribution tree. This can be viewed as an alternative implementation of model 2 since, as we observe below, the basic allocation process can be adequately handled with a one-pass mechanism.

cal members and so if costs are applied more generally we are faced with the tradeoff between the increased implementation difficulty of model 2 with the correspondingly better allocation policy and the significantly easier implementation of model 1, which comes with a seriously flawed allocation policy. An intermediate point is to keep track of the cumulative number of domains with internal members instead of the number of individual host members. Only border routers would propagate the member-domain counts upstream, eliminating the need for interior-router modifications and thereby sidestepping much of the implementation difficulties.

The discussion in this section focused exclusively on the one-pass accounting mechanism. There are a wide variety of other approaches available; why narrow consideration to this particular family? From a purely mechanistic perspective, the one-pass mechanism has several desirable properties: simple, easy to implement, and scalable (in that the state carried in the accounting message does not grow with the number of members or with the size of the distribution tree). In addition, in the context of model 2, the one-pass accounting mechanism can implement the ELSD policy which, on purely axiomatic grounds, was identified as the most natural allocation policy.<sup>26</sup> The most obvious drawback with the one-pass accounting mechanism is that when combined with model 1, it cannot implement allocation policies that obey the equivalency axiom. While we do think it important to explore other accounting mechanisms, we expect that achieving equivalency will require information that is equivalent to knowing the number of downstream members. That is, we believe that the key factor preventing equivalency is the difference between model 1 and model 2, not the features of the one-pass mechanism itself.

## VI. GENERALIZING THE MODEL

To more clearly illustrate the basic issues inherent in sharing the cost of multicast trees, we considered an extremely simple network model. In this section, we briefly discuss some directions in which this model can be generalized. A fuller discussion of these generalizations can be found in [9].

### A. Different Qualities of Service

By assigning a single cost for traversing a link, our basic model implicitly assumed that all receivers request the same quality of service (QoS). It is clearly unrealistic to require such uniformity of service; there will likely be cases where some receivers request best-effort service and others establish reservations with an assured quality of service. We can extend the model to accommodate multiple QoS's by introducing a link cost that depends on the level of service. This, in turn, affects the canonical form that was introduced in Section IV-B. With  $l$  distinct levels of QoS, the family of functions  $\mathcal{F}$  becomes

<sup>26</sup>We should note that we have also analyzed a two-pass accounting mechanism, where the first pass goes upstream from receivers to senders and the second pass goes downstream from senders to receivers. In the context of model 2, this two-pass mechanism is capable of implementing all canonical cost allocation strategies. However, since the ELSD policy seems the most natural policy and it does not require the extra complication of the two-pass mechanism we will not discuss the two-pass mechanism here.

those functions  $F(z_u, z_d) = [F_u(z_u, z_d), F_d(z_u, z_d)]$  that<sup>27</sup> satisfy: 1)  $\sum_{i=1}^l z_u^i F_u^i(z_u, z_d) + z_d^i F_d^i(z_u, z_d) = 1$  and 2)  $z_d^i = 0 \Rightarrow F_d^i(z_u, z_d) = 1$  and  $z_u^i = 0 \Rightarrow F_u^i(z_u, z_d) = 1$ . This function  $F(z_u, z_d)$  determines the fraction of the link cost allocated to each upstream or downstream receiver based on the number of upstream and downstream receivers of each QoS level:  $n_u^1, n_d^1, \dots, n_u^l, n_d^l \dots n_u^l, n_d^l$ . Given a link  $(v_i, v_j)$ , the following costs are allocated to a member with QoS level  $i$ :

$$\begin{aligned} \text{For each upstream member:} & \quad F_u^i(z_u, z_d) * c(v_i, v_j) \\ \text{For each downstream member:} & \quad F_d^i(z_u, z_d) * c(v_i, v_j). \end{aligned}$$

Multiple qualities of service also requires some changes in implementation. If there are a discrete set of QoS levels, then models 1 and 2 can easily be changed to include information about the presence (or exact number) of members requesting each level. However, a continuum of QoS levels (e.g., reserved bandwidth) may result in as many QoS levels as the number of receivers.

Our original network model addressed the issue of how to share costs between equivalent users at different locations. Having multiple QoS's raises the issue of how to share the cost between several members who are in the same place but request different QoS levels.<sup>28</sup> We do not address this problem here, since it is a local issue. In contrast with the need for consistency in how costs are shared between users at different locations, the allocation of costs between different service levels at a given link is a decision that can be made locally.

### B. Different Reservation Models

Since multicast packets follow delivery trees, any "costs" associated with individual packet delivery can be shared according to our analysis here. However, for costs associated with reservations, the single source aspect of the tree must be generalized. For instance, RSVP supports the sharing of reservations by multiple senders with its wild-card and shared-explicit reservation styles. For these cases, the reservation established (i.e., the level of service) is independent of the number of upstream sources. It can be shown (see [9]) that one can use the same canonical formalism as before (in Section IV-B), except now the relevant distinction is not whether a receiver is upstream or downstream or a particular source, but whether it is downstream of *any* source.

### C. Different Multicast Distribution Models

Our model for multicast distribution, based on source-rooted trees computed from unicast routing is representative of the most widespread multicast routing protocols (e.g., PIM's

<sup>27</sup> $F_u: \{Z_+^{2l} - 0^l\} \mapsto \mathbb{R}_+^l$  and  $F_d: \{Z_+^{2l} - 0^l\} \mapsto \mathbb{R}_+^l$ .

<sup>28</sup>If the set of qualities of service is completely ordered then the problem reduces to what has been called in the literature the *airport game* (see [10]) where the cost of the link is the cost associated with the highest QoS requested. Here, the Shapley value of this game is easy to compute; it is very much like the ELSD scheme in that every user shares equally the incremental cost of all levels less than or equal to their requested level. This cost allocation formula for this special form of the problem can be axiomatized in a variety of ways (see [16]) and appears to be a rather natural choice.

SPT’s [7], MOSPF [17], DVMRP [26]). Some newer multicast routing protocols use different models for their distribution. CBT [1] uses a shared tree for all sources (core-rooted rather than source-rooted trees). PIM [7], in its sparse mode, enables mixing, within the same multicast group, source-rooted shortest path (SPT) routes for some senders and a single core-rooted shared tree for all other senders to the group. These developments do not require modification of our analysis, since we only assumed that the route taken from a particular source to a particular receiver is independent of the group membership (i.e., the route a data packet will traverse to reach a specific group member does not depend on who else has joined the group). This remains true for both core-rooted trees as well as source-rooted trees and, thus, our results apply CBT and PIM.<sup>29</sup>

One modification of routing that would invalidate our analysis is the ability to request alternate routes. There are proposals [27] to allow receivers to request a route that is different than the default unicast-based route. When an alternate route is requested, the route is followed until it hits the shared tree. Thus, the route will depend on the current multicast membership, which violates one of the crucial assumptions of our routing model.<sup>30</sup>

## VII. APPENDIX

*Theorem 1:* A cost allocation formula satisfies the basic axioms if and only if it is a canonical strategy.

It is straightforward to verify that any canonical form satisfies the basic axioms. We now show that any cost allocation formula that obeys the basic axioms can be expressed in the canonical form. Consider any cost allocation formula that obeys the basic axioms. We begin with the most general case of a network  $N = (V, L, T)$  and a tree  $L' = T[N, R, \text{loc}(M)]$ . Since the allocation function is additive, we can restrict our attention to cost functions  $c_{i,j}$  which have nonzero cost only on the link  $(v_i, v_j)$ , and have unit cost on that link. More specifically, we know that

$$\sum_{(i,j):(v_i,v_j) \in L} [c_{i,j}(v_i, v_j) * c(v_i, v_j)] = c$$

and so

$$\begin{aligned} & \sum_{(i,j):(v_i,v_j) \in L} [af_\alpha(N, R, M, \text{loc}, c_{i,j}) * c(v_i, v_j)] \\ &= af_\alpha(N, R, M, \text{loc}, c) \quad \forall N, R, \text{loc}, M, \alpha : m_\alpha \in M. \end{aligned}$$

We must now show that the cost allocations that result from cost functions  $c_{i,j}$  can be expressed in terms of the canonical form.

Consider the *reduced* network with the single link  $(v_i, v_j)$  and the following loc function:

$$\text{loc}'(m_\beta) = \begin{cases} v_j: & m_\beta \text{ is downstream to } (v_i, v_j). \\ v_i: & \text{Otherwise} \end{cases}$$

<sup>29</sup> Although this assumes that the route used when computing the stand-alone cost is chosen using multicast rather than reverting to a unicast route.

<sup>30</sup> Alternate path routing is not yet deployed in the Internet, but the expectation is that even when deployed the use of alternate paths would always represent a very small percentage of the routes in use.

The subset costs of the reduced network problem are the same as the original problem, and the equivalency axiom requires that the cost allocations be the same. The anonymity condition requires that all members at the same node be allocated the same cost. Thus, the allocations in this reduced problem are characterized by two quantities, the allocation to the upstream members and the allocation to the downstream members. These allocations can depend on the number of upstream and downstream members, so they are expressed as functions  $F_u(n_u, n_d)$  and  $F_d(n_u, n_d)$ . These costs must be nonnegative, and the budget balance requirement means that  $n_u F_u(n_u, n_d) + n_d F_d(n_u, n_d) = 1$ . This is precisely the canonical form.<sup>31</sup> Q.E.D.

*Theorem 2:* A cost allocation function satisfies the basic and stand-alone axioms if and only if it is the equal link split downstream (ELSD) function.

Consider any canonical form  $F$  that obeys the stand-alone axiom. The stand-alone axiom implies that  $F_u(n_u, n_d) = 0 \forall n_u > 0$ . Combining this with the budget balance condition  $n_u F_u(n_u, n_d) + n_d F_d(n_u, n_d) = 1$  yields  $F_d(n_u, n_d) = 1/n_d \forall n_d > 0$  which is the ELSD formula. It is straightforward to verify that the ELSD formula satisfies the stand-alone axiom, so the converse holds as well. Q.E.D.

*Theorem 3:* A cost allocation function satisfying the basic axioms and the no-free-rider axiom must satisfy:

$$F_u(n_u, n_d) \leq F_d(n_u, n_d) \quad \forall (n_u, n_d) \in Z_{++} \times Z_+$$

Consider any canonical form  $F$ . The no-free-rider axiom is obeyed if and only if  $1/(n_u + n_d) \leq F_d(n_u, n_d)$  whenever  $n_d > 0$ . However, combining this with the budget balance condition  $n_u F_u(n_u, n_d) + n_d F_d(n_u, n_d) = 1$  yields  $n_u [F_d(n_u, n_d) - F_u(n_u, n_d)] \geq 0$ . Thus, whenever  $(n_u, n_d) \in Z_{++} \times Z_+$ , we must have  $F_d(n_u, n_d) \geq F_u(n_u, n_d)$ . Q.E.D.

*Theorem 4:* A cost allocation formula satisfies equivalency, anonymity, and subset monotonicity if and only if it is the equal tree split (ETS) formula.

Clearly the ETS scheme satisfies the anonymity, equivalency, and subset monotonicity axioms. We must now show that any allocation function which obeys the anonymity, equivalency, and subset monotonicity axioms must be the ETS policy. Consider any network  $N = (V, L, T)$  with a root  $R$ , set of members  $M$ , location function  $\text{loc}$ , and cost function  $c$ . Then, consider another network  $N'$  with a single link and the same set of members all located on the node downstream from the root. Set the cost of this single link to be the sum of all link costs in the original network:  $cf\{T[N, R, \text{loc}(M)], c\}$ . By anonymity, all assigned costs in this new network are equal and so are given by the ETS formula. The costs of all subsets in this new network are greater than or equal to the subset costs in original network (so all assigned costs must be greater than or equal) but the total cost is the same in the two trees; therefore, the assigned costs in both networks are the same. Q.E.D.

<sup>31</sup> The conditions  $F_u(0, n_d) = 1$  and  $F_d(n_u, 0) = 1$  in the definition of  $F$  are irrelevant to the actual allocations.




Fig. 2. Subset monotonicity versus one-pass model.

**Theorem 5:** No cost allocation formula can satisfy the collusion prevention axiom.

Consider a simple network consisting of a single link (with unit link cost) connecting two nodes  $v_1$  and  $v_2$ , with the root at  $v_1$ . Let  $f_\alpha(M)$  be the cost allocated to a member  $m_\alpha$  if we place the set of members  $M$  at node  $v_2$ . By budget balance, we must have  $\sum_{\alpha=1}^n f_\alpha(M) = 1$ . Assuming  $|M| \geq 5$  we can pick two members  $m_\alpha$  and  $m_\beta$  such that  $f_\alpha(M) + f_\beta(M) < 0.5$ . Removing all members but  $m_\alpha$  and  $m_\beta$ , without loss of generality, we can assume  $f_\alpha(m_\alpha + m_\beta) \leq 0.5$ . If the members of  $M' = M - m_\alpha - m_\beta$  collude with  $m_\alpha$  their total cost (including  $m_\alpha$ ) is just  $f_\alpha(m_\alpha + m_\beta) \leq 0.5$ , which is smaller than their original cost  $1 - f_\beta(M) > 0.5$ . Thus, collusion prevention cannot hold. Q.E.D.

**Theorem 6:** For any one-pass accounting mechanism, no member can be allocated a cost greater than her unicast cost.

From the definition of the one-pass model we know that: 1) nodes are budget balanced so their output and allocations are bounded by their input and 2)  $\text{in}(v_j) = c(v_i, v_j) + \text{out}(v_i, v_j)$  for each  $v_j$  downstream to  $v_i$ . Thus,  $\text{in}(v_j) \leq c(v_i, v_j) + \text{in}(v_i)$ . If we iterate this inequality for each node along the path from  $R$  to any node  $v_k$ , and collapse the inequalities, we find:  $\text{in}(v_k) \leq c_f[T(N, R, v_k), c]$ . Q.E.D.

**Theorem 7:** No cost allocation scheme implemented with a one-pass accounting mechanism can satisfy the subset monotonicity axiom.

Consider an accounting scheme that satisfies the subset monotonicity axiom. Consider the example in Fig. 2 of a single tree  $T[N, R, \text{loc}(M)]$  with two different link cost functions:  $c_1, c_2$ . The cost allocations for  $c_1$  are straightforward:  $af_1(\cdot) = af_2(\cdot) = 10$ . Because the total cost of the tree remains the same, while no subset cost decreases, subset monotonicity implies that the cost allocations must be the same for  $c_1$  and  $c_2$ . To achieve this allocation for  $c_2$  it must be that  $\text{out}(v_2, v_3) = 10$ . But we must have  $\text{out}(v_2, v_3) = \text{out}(v_2, v_4)$  which violates the local budget balance; contradiction. Q.E.D.

**Theorem 8:** No cost allocation formulae that is implemented by a model 1 one-pass accounting scheme can satisfy the equivalency axiom.

Consider a model 1 one-pass accounting scheme that satisfies the equivalency axiom. Consider the following example of a network illustrated in Fig. 3, where  $c(v_1, v_2) > 0$




Fig. 3. One-pass, model 1 versus equivalency.

and  $c(v_2, v_3) = c(v_2, v_4) = 0$ . The equivalency axiom implies that the cost allocations to members on  $v_3$  and  $v_4$  should be the same:  $af_1(\cdot) = af_2(\cdot) = af_3(\cdot)$ . However, the definition of model 1 implies that  $\text{in}(v_3) = \text{in}(v_4) = [c(v_1, v_2)]/2$ . Therefore, the cost allocations cannot be the same; contradiction. Q.E.D.

**Theorem 9:** A model 1 cost allocation formula satisfies the reduced basic axioms if and only if it can be expressed in the one-pass canonical form.

It is straightforward to verify that any one-pass canonical form satisfies the reduced basic axioms. We now show that any model 1 cost allocation formula that obeys the basic axioms can be expressed in the one-pass canonical form. Similar to the proof for Theorem 1, we begin with the most general case of a network  $N = (V, L, T)$  and a tree  $L' = T[N, R, \text{loc}(M)]$ . Since the allocation function is additive, we can restrict our attention to cost functions  $c_{i,j}$  which have nonzero cost only on the link  $(v_i, v_j)$ , and have unit cost on that link. We must now show that all cost allocations that result from cost functions  $c_{i,j}$  can be expressed in terms of the canonical form.

The anonymity axiom implies that all local member are allocated identical costs. The definition of model 1 requires that the residual costs passed on to all next hops are the same. These allocations can depend on the number of local members and downstream next-hops, so they are expressed as functions  $F_l(n_l, n_r)$  and  $F_r(n_l, n_r)$ . These costs must be nonnegative, and the budget balance requirement means that  $n_l F_l(n_l, n_r) + n_r F_r(n_l, n_r) = 1$ . Using additivity, we can show that the allocations must be proportional to the incoming costs so

$$\begin{aligned} \text{Local members:} & \quad \text{in}(v_i) * F_l(n_l, n_r) \\ \text{Residual costs to node } v_k: & \quad \text{out}(v_j, v_k) = \text{in}(v_i) * F_r(n_l, n_r). \end{aligned}$$

In model 1, the input costs to the next hop  $v_k$  include both residual cost and the cost of the link connecting them:  $\text{in}(v_k) = \text{in}(v_i) * F_r(n_l, n_r) + c(v_i, v_k)$ . This is precisely the one-pass canonical form. Q.E.D.

**Theorem 10:** A model 1 cost allocation formula satisfies the reduced basic axioms and the sharing-is-good axiom if and only if the functions  $F_l(n_l, n_r)$  and  $F_r(n_l, n_r)$  are nonincreasing on  $\{Z_+^2 - (0, 0)\}$ .




Fig. 4. No-free-rider axiom in model 1.

We will show that the sharing-is-good axiom and the nonincreasing property yield the same set of conditions on  $F_l(n_l, n_r)$  and  $F_r(n_l, n_r)$ . When adding a local member to a node already on the distribution tree, sharing-is-good requires

$$\begin{aligned} F_l(n_l + 1, n_r) &\leq F_l(n_l, n_r) & \forall (n_l, n_r) \in Z_{++} \times Z_+ \\ F_r(n_l + 1, n_r) &\leq F_r(n_l, n_r) & \forall (n_l, n_r) \in Z_+ \times Z_{++}. \end{aligned}$$

When adding a member to a node, which was not on the distribution tree, at the nearest node on the tree sharing-is-good requires

$$\begin{aligned} F_l(n_l, n_r + 1) &\leq F_l(n_l, n_r) & \forall (n_l, n_r) \in Z_{++} \times Z_+ \\ F_r(n_l, n_r + 1) &\leq F_r(n_l, n_r) & \forall (n_l, n_r) \in Z_+ \times Z_{++}. \end{aligned}$$

When we combine these equations with the second condition in the definition of  $\mathcal{F}$  we get exactly the nonincreasing condition on all of  $\{Z_+^2 - (0, 0)\}$ . Q.E.D.

**Theorem 11:** There is no model 1 cost allocation formula that satisfies the reduced basic axioms and the no-free-rider axiom.

Assume there is a model 1 cost allocation formula that satisfies the reduced basic axioms and the no-free-rider axiom. In Theorem 9, we have shown that all model 1 cost allocation formulas obeying the reduced basic axioms can be expressed in the one-pass canonical form. Consider the tree detailed in Fig. 4: with no members in  $v_2$ ,  $F_r(0, 2) = 1/2$  for all formulas, resulting in  $\text{in}(v_4) = 6$ . With no next-hops downstream of  $v_4$ ,  $F_l(2, 0) = 1/2$  for all formulae, making the allocation to each member on  $v_4$  exactly three, which is less than  $11/3$  (unicast cost/ $|M|$ ). Q.E.D.

## REFERENCES

- [1] A. J. Ballardie, P. F. Francis, and J. Crowcroft, “Core based trees (CBT),” in *Proc. ACM SIGCOMM’93*, Aug. 1993.
- [2] H. W. Braun and K. Claffy, “A framework for flow-based accounting on the Internet,” in *Proc. IEEE SICON*, Singapore, 1993.
- [3] D. Clark, “Adding service discrimination to the Internet,” in *Proc. 23rd Annu. Telecommun. Policy Res. Conf.*, 1995; Available HTTP: <http://anaweb.lcs.mit.edu/anaweb/ps-papers/TPRC2-0.ps>
- [4] D. Clark, S. Shenker, and L. Zhang, “Supporting real-time applications in an integrated services packet network: Architecture and mechanism,” in *Proc. ACM SIGCOMM*, Aug. 1992.
- [5] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang, “Pricing in computer networks: Motivation, formulation, and example,” *IEEE/ACM Trans. Networking*, vol. 1, pp. 614–627, Dec. 1993.
- [6] S. Deering and D. Cheriton, “Multicast routing in datagram Internet networks and extended LAN’s,” *ACM Trans. Comput. Syst.*, pp. 85–111, May 1990.
- [7] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C. G. Liu, and L. Wei, “An architecture for wide-area multicast routing,” in *Proc. ACM SIGCOMM*, Sept. 1994.

- [8] D. Henriot and H. Moulin, “Traffic based cost allocation in a network,” *Rand J. Econ.*, vol. 27, no. 2, pp. 332–345, 1996.
- [9] S. Herzog, “Accounting and access control for multicast distributions,” Ph.D. dissertation, Univ. Southern California, Comput. Sci. Dept., Aug. 1996.
- [10] S. C. Littlechild and G. Owen, “A simple expression for the Shapley value in a special case,” *Manag. Sci.*, vol. 20, pp. 370–372, 1973.
- [11] J. K. MacKie-Mason and H. R. Varian, “Pricing congestible network resources,” *IEEE J. Selected Areas Commun.*, vol. 13, pp. 1141–1149, 1995; also, Available FTP: <ftp://goopher.econ.lsa.umich.edu/pub/Papers/pricing-congestible.ps.Z>
- [12] ———, “Pricing the Internet,” B. Kahin and J. Keller, Eds., *Public Access to the Internet*. Englewood Cliffs, NJ: Prentice-Hall, 1995; available HTTP: [http://www.spp.umich.edu/spp/papers/jmm/Pricing\\_the\\_Internet.ps.Z](http://www.spp.umich.edu/spp/papers/jmm/Pricing_the_Internet.ps.Z)
- [13] N. Megiddo, “Computational complexity of the game theory approach to cost allocation for a tree,” *Math. Operations Res.*, vol. 3, pp. 189–196, 1978.
- [14] H. Moulin, “Uniform externalities: Two axioms for fair allocation,” *J. Public Economy*, vol. 43, pp. 305–326, 1990.
- [15] ———, “Welfare bound in the cooperative production problem,” *Games Econ. Behavior*, vol. 4, pp. 373–401, 1992.
- [16] H. Moulin and S. Shenker, “Average cost pricing versus serial cost sharing: An axiomatic comparison,” *J. Econ. Theory*, vol. 64, no. 1, pp. 178–201, 1994.
- [17] J. Moy, “Multicast extensions to OSPF,” *Internet Request for Comments*, RFC-1584, Mar. 1994.
- [18] R. B. Myerson, *Game Theory*. Cambridge, MA: Harvard Univ. Press, 1991.
- [19] H. Peyton-Young, Ed., *Cost Allocation: Methods, Principles, Applications*. New York: Elsevier Sci., 1985.
- [20] A. E. Roth, Ed., *The Shapley Value, An Essay in Honor of Lloyd S. Shapley*. Cambridge, U.K.: Cambridge Univ. Press, 1988.
- [21] J. Sairamesh, D. F. Ferguson, and Y. Yemini, “An approach to pricing, optimal allocation and quality of service provisioning in high-speed packet networks,” in *Proc. Infocom’95*, Apr. 1995.
- [22] L. S. Shapley, “A value for N-person games,” H. W. Kuhn and A. W. Tucker, Eds., *Contributions to the Theory of Games, Vol. II, Annals of Mathematics Studies No. 28*. Princeton, NJ: Princeton Univ. Press, 1953, pp. 307–317.
- [23] W. W. Sharkey, “Network models in economics,” M. O. Ball *et al.*, Eds., *Handbook Operations Res. Management Sci.: Networks*, vol. 8, 1995, pp. 713–765.
- [24] S. Shenker, “Service models and pricing policies for an integrated services Internet,” B. Kahin and J. Keller, Eds., *Public Access to the Internet*. Englewood Cliffs, NJ: Prentice-Hall, 1995.
- [25] S. Shenker, D. Clark, D. Estrin, and S. Herzog, “Pricing in computer networks: Reshaping the research agenda,” *Telecommunications Policy*, vol. 20, 1996; also published in *Proc. 23rd Ann. Telecommun. Policy Res. Conf.*, 1995.
- [26] D. Waitzman, C. Partridge, and S. Deering, “Distance vector multicast routing protocol,” *Internet Request For Comments*, RFC-1075, Nov. 1988.
- [27] D. Zappala, “Multicast routing support for integrated services,” Ph.D. dissertation, Univ. Southern California, Comput. Sci. Dept., 1995; Available FTP: <ftp://catarina.usc.edu/pub/daniel/RSRR/quals.ps>.



**Shai Herzog** (M’96) received the LL.B. degree from Tel-Aviv University, in 1985; the M.Sc. degree from Bar-Ilan University, Israel, in 1991; and the M.S. and Ph.D. degrees from the University of Southern California in Los Angeles, in 1993 and 1996, respectively.

He is currently the Chief Technology Officer (CTO) of IPHighway Inc., San Jose, CA. Between 1993 and 1996, he was a Research Assistant at Information Sciences Institute, where he co-developed RSVP (ReSerVation Protocol). He later joined IBM T. J. Watson Research Center as a Research Staff Member (1996–1997) and recently co-founded IPHighway, a startup company specializing in service differentiation and Internet policy control. His most recent computer science research focuses on Internet policy control approaches; his interests include network accounting, access-control, cost allocation, billing, and other issues of policy control across public and private networks.

Dr. Herzog is a member of ACM.

**Scott Shenker** (M'87–SM'95) received the Sc.B. degree from Brown University, Providence, RI, in 1978, and the Ph.D. degree in theoretical physics from the University of Chicago, in 1983.

He is currently a Principal Scientist at the Xerox Palo Alto Research Center, Palo Alto, CA. He spent the 1983–1984 academic year at Cornell University as a Post-Doctoral Associate. His most recent computer science research focuses on the design of integrated services packet networks and the related issues of service models, scheduling algorithms, and reservation protocols. His recent economic research addresses incentive compatibility and fairness in various cost sharing mechanisms. Besides computer networks and theoretical economics, his other research interests include chaos in nonlinear systems, critical phenomena, distributed algorithms, conservative garbage collection, and performance analysis.

Dr. Shenker is a member of ACM.

**Deborah Estrin** (SM'78–M'78) received the B.S. degree from the University of California, Berkeley, in 1980, and the M.S. and Ph.D. degrees from the Massachusetts Institute of Technology (MIT), Cambridge, in 1982 and 1985, respectively.

She is currently an Associate Professor of Computer Science at the University of Southern California in Los Angeles, where she joined the faculty in 1986.

Dr. Estrin received the National Science Foundation's Presidential Young Investigator Award in 1987 for her research in network interconnection and security. She is a co-PI on the DARPA funded VINT and RSVP projects, and on the NSF Routing Arbiter project, all at USC's Information Sciences Institute. She is an active participant in the Inter-Domain Multicast Routing working group. She is a member of the ACM and AAAS. She has served on several panels for the National Science Foundation and National Research Council/CSTB, and is currently a member of DARPA's ISAT advisory group. She is currently an Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING and was Program Co-Chair for ACM's Sigcomm'96.