Sharing the "Cost" of Multicast Trees: An Axiomatic Analysis *

Shai Herzog
herzog@isi.edu
Computer Science Department/ISI
University of Southern California
Los Angeles, CA 90089

Scott Shenker
shenker@parc.xerox.com
Xerox PARC
3333 Coyote Hill Road
Palo Alto, CA 94304

Deborah Estrin
estrin@usc.edu
Computer Science Department/ISI
University of Southern California
Los Angeles, CA 90089

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 reduces network usage 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.

1 Introduction

Because 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, we are now approaching an almost completely commercial (i.e., unsubsidized) Internet environment. This means that issues of cost recovery and profit incentives, for better or worse, become much more relevant. [1]

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 cost of the network resources they are consuming.

It is important to note that these costs are not necessarily financial (hence the quotes around the term in the title), nor are these costs necessarily extracted from the users in the form of payments.[2] When users are cooperative, they want to make the right decision about the wise use of network resources. Such decisions can only be made if users are informed about the "cost" their usage imposes on the network. Feedback about these costs is important if we wish to promote the efficient use of network resources. 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.[3]

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?[4] The whole point of multicast is to reduce network usage by sharing a delivery tree. We ask: how are these savings allocated among the various members of the multicast group? We should note that we are making the fundamental assumption that costs are assigned to the receivers of the multicast group, and not to the sources. This is because multicast membership is typically receiver initiated (see [5]). Some applications may impose an arrangement to share costs between the sender

*This research was supported by the Advanced Research Projects Agency under Ft. Huachuca contract number DABT63-91-C-0001, entitled "Gigabit Network Communications Research". The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Advanced Research Projects Agency, the Department of the Army, or the U.S. Government.

To appear in ACM SIGCOMM'95 Conference, August 1995, Cambridge, Massachusetts, USA

[1] In addition, as argued in [4, 3, 20], 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. Appropriate use leads to the efficient use of network resources, where *efficiency* is defined in [4, 20]. These incentive issues are also relevant to controlling usage even in a single QoS environment (see [8, 9]).

[2] 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.

[3] 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.

[4] Our treatment here only considers a flow with a single source; the generalization to multiple senders is discussed in Section 5.

and receiver, but we consider this to be a higher-layer issue.

We use an axiomatic approach to the problem of cost allocation. 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.

To focus attention on the basic issues, our investigation is in the context of a very simple network model (we consider generalizations in Section 5). 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. We do not discuss issues of security or of network dynamics.

This paper has 5 sections. We present the basic model in Section 2, and discuss various cost allocation axioms in Section 3. We turn to issues of implementation in Section 4, where we discuss the ability of our simple one-pass approach to satisfy the axioms presented in Section 3. We conclude in Section 5 with a discussion of various directions for future work.

Before turning to the basic model, we first discuss some related work. There is a rapidly growing literature on pricing in computer networks. See References [8, 4, 17, 2, 20] 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, which is what we consider here. This is an important difference. In theory, proper pricing requires knowledge of the users' utility functions (at least in aggregate, such as in a demand function); in practice, pricing also involves many other non-technical marketing factors. We consider none of those issues here. We instead focus exclusively on the allocation of costs.[5]

There is a large economics literature on cost allocation; see Reference [15] 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 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 References [16, 14]). Axiomatic approaches are often, in fact almost exclusively, used to analyze such cooperative games. The axioms presented here are very similar to those used in the theory of cooperative games. Megiddo [10] considered 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 [19] for a lucid overview of these results and for a much broader discussion of network models in economics. We are not aware, however, of any treatment of the mechanisms needed to implement such cost allocation schemes; see our discussion in Section 4. While the presentation in this paper is geared towards the networking research community, we have included several footnotes which relate our work to the economics literature. Also, to aid the presentation, we defer all proofs to the appendix.[6]

2 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.

2.1 Basic Definitions

In this section we present an abstract model of the cost-sharing problem. This model will be the basis for our axiomatic analysis of cost allocation. The abstract model defines a formal description of the **network technology** for multicast transmission over a network. The abstract model also describes the structure of **incurred costs** of such a multicast transmission. The model we describe below is illustrated in Figure 1 for a simple network scenario.

The **network technology** model has three components: **network topology**, **group membership** and **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*[7] links $(v_i, v_j) \in L$ and a routing function (protocol) T. Let M describe the set of **multicast group members**; we denote individual members by $m_\alpha \in M$. Members of a multicast group can join and leave dynamically and can attach themselves to any node. 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

[5] This is not to say that costs and prices are completely unrelated. For instance, in a competitive market, prices are often set at the true marginal costs of usage. In a regulated market, often prices are legislated to be linked to true costs.

[6] In the conference version of the paper we will refer readers to an online version of the appendix.

[7] Notice that links are directed and therefore $(v_i, v_j) \neq (v_j, v_i)$ for all nontrivial links (i.e., $v_i \neq v_j$).

the set of unicast routes from a single root R to all other nodes forms a tree.[8] 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 $V' \subseteq V$.

Each link in a distribution tree creates a distinction, relative to itself, between upstream and downstream portions of the tree: the downstream portion 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 now turn to the structure of incurred costs. Our abstract model considers the allocation of costs for each data flow individually (i.e., we do not consider the joint allocation of the costs of multiple flows), where data flow refers to a particular source sending to a particular multicast group destination. Such a flow can be explicit (created by resource reservation) or implicit.[9]

For reasons of simplicity, we associate all costs with the links that the packets traverse.[10] These costs could arise from many different aspects of network usage; for instance there could be costs associated with each packet and/or costs associated with resource reservations (bandwidth, etc.). We lump all of these different costs into a single quantity.[11] 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. Later, in section 5, 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 address dynamics, either in membership or in topology.[12]

The set of link costs for a given network $N = (V,L,T)$ is defined as $c : L \mapsto \Re_+$. The **cost function** of a distribution tree comprised of links L' is merely the sum of all the directed links in L': $cf(L',c) = \sum_{(v_i,v_j)\in L'} c(v_i,v_j)$. This function expresses the total cost incurred by a distribution tree. In this paper, we explore how to allocate this cost among the individual members.

[8] 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 (QoR), which would violate our assumption.

[9] Data packets can be classified implicitly into flows by examining their source/destination addresses and other significant fields in their headers

[10] Although most reservation protocols reserve link resources and not intra-node resources, we acknowledge that some intra-node resources may be costly (like internal buffers, CPU cycles, etc.). Intra-node costs can be attributed to links by attributing the intra-node cost associated with processing incoming data to the incoming link and the overhead for outgoing data to the appropriate outgoing links.

[11] These costs are in no way related to the link-metrics used by routing algorithms to calculate routes in the network.

[12] Of course, our approach in not limited to the completely static case. One can apply our approach to allocating cost for time slices, and assuming static behavior within those time slices. The costs associated with transitioning (either membership or topology) is not addressed here.

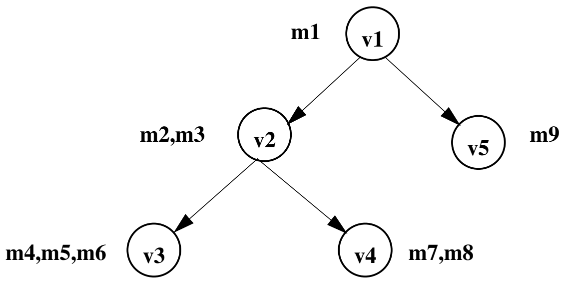

Figure 1: A simple example of a network.

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 attributed to members of the multicast group; $af_\alpha(N,R,M,loc,c)$ denotes the allocation to member α.

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:

$$\sum_{m_\alpha \in M} af_\alpha(N,R,M,loc,c) \;=\; cf(T(N,R,loc(M)),c) \qquad \forall N,\ R, M, loc,\ and\ c$$

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

$$af_\alpha(N,R,M,loc,c) \geq 0 \qquad \forall\ \alpha\ N,\ R,\ M,\ loc\ and\ c$$

2.2 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 Figure 1 to illustrate these example 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 Figure 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.[13] 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 $c(v_1, v_2)/7 + c(v_2, v_3)/3$.

An approach that is midway between the egalitarian approach of ETS and the emphasis on individual accountability of ELSD is to assign the cost of a link equally to all the next hop links that are part of the distribution tree (with all receivers at that node being treated as part of one other link)[14] This is motivated by the idea that we pass costs on to each downstream next-hop (rather than allocate them to the downstream receivers themselves) and then allocate the costs of those downstream links in a recursive fashion. These costs are recursively forwarded until there are no more downstream nodes and costs are fully allocated. We call this the **Equal Next-Hop Split (ENHS)** scheme. For the example in Figure 1, the cost allocated to $m4$ is given by $c(v_1, v_2)/9 + c(v_2, v_3)/3$ and the cost allocated to $m2$ is $c(v_1, v_2)/6$

Of course, all of these 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 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.[15] The unicast cost of member α is given by $\sum_{(v_i,v_j)\in T(N,R,loc(m_\alpha))} c(v_i, v_j)$. For the example in Figure 1, the unicast cost of $m4$ is $c(v_1, v_2) + c(v_2, v_3)$.

3 Axiomatic Analysis

We analyze this abstract model of cost-sharing through the use of axioms. We use as axioms some possible properties that desirable cost-sharing formula might have. We 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 3.1, 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 3. We discuss the implications of these basic axioms in Section 3.2 and then present a few examples cost allocation policies which satisfy them in Section 3.3. In Section 3.4 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.

[13] For those familiar with the concept, this is the Shapley Value of this problem.

[14] We give a more precise definition of this scheme in Section 4.

[15] Note that if the allocated multicast cost is greater or equal to the unicast cost there may be no incentive to use multicast.

3.1 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 relabelling of equivalent members; it should not have an intrinsic and arbitrary asymmetry built into it.[16]

Axiom 1 *Anonymity:* [17]

Name labels, by which members are identified, are irrelevant to cost allocation.

Formal definition:

Given any network $N = (V, L, T)$, any link cost function c, any member set M, any two placement functions loc and loc' and any two members m_α and $m_\beta \in M$:

if

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

then

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

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. It is important, in order to avoid the additional complexity and overhead of independently allocating all such costs, that the cost allocation 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 $N = (V, L, T)$, any member set M, any placement function loc, and any two link cost functions $c1$ and $c2$:

[16] 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 non-equivalency of members (e.g., introducing two classes of members into the formalism).

[17] Notice that this axiom implies that two members located at the same node must be allocated the same cost: $loc(m_\alpha) = loc(m_\beta) \Longrightarrow af_\alpha(\cdots) = af_\beta(\cdots)$.

$$af(N,R,M,loc,c1+c2) = af(N,R,M,loc,c1)+af(N,R,M,loc,c2)$$

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 charges for every subgroup of members, then they should allocate the cost in the same way. This condition, that the networks charge the same for 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.[18]

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 $N1$, $R1$, $loc1$, $c1$ and $N2$, $R2$, $loc2$, $c2$:

if

$$cf(T(N1,R1,loc1(M')),c1) = cf(T(N2,R2,loc2(M')),c2) \quad \forall\ M' \subseteq M$$

then

$$af(N1,R1,M',loc1,c1) = af(N2,R2,M',loc2,c2) \quad \forall\ M' \subseteq M$$

3.2 The Canonical Form of Cost Allocation Formulae

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 $F : \{Z_+^2 - (0,0)\} \mapsto \Re_+^2$, where we use the symbol Z_+ to denote the nonnegative integers and $\Re_+$ to denote the nonnegative reals.[19] We define the family of functions $\mathcal{F}$ to be those functions $F : \{Z_+^2 - (0,0)\} \mapsto \Re_+^2$ that satisfy the properties (1) $z_1F_1(z_1,z_2) + z_2F_2(z_1,z_2) = 1$ and (2) $F_1(0,z) = F_2(z,0) = 1$. Note that the set $\mathcal{F}$ is convex: $F,G \in \mathcal{F} \Longrightarrow (\mu F + (1-\mu)G) \in \mathcal{F}$ for all $0 \le \mu \le 1$. Also, we know that $F_2(0,z) = F_1(z,0) = \frac{1}{z}$ for $z > 0$.[20]

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[21] and $n_u \ge 0$ 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:

$F_u(n_u,n_d) * c(v_i,v_j)$ for each upstream member.
$F_d(n_u,n_d) * c(v_i,v_j)$ for each downstream member.

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: $n_uF_u(n_u,n_d)+n_dF_d(n_u,n_d) = 1$.[22] 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.*

3.3 Examples of Canonical Cost Allocation Strategies

The set $\mathcal{F}$ is infinite, and so there are an infinite number of canonical cost allocation strategies. In this section we present several examples of functions $F \in \mathcal{F}$. Some strategies may seem intuitive and others may look less than that (like charging upstream receivers), but they all obey the three basic axioms. Two of the functions were already introduced in section 2.2. Since we know that $F_d(0,n_d) = \frac{1}{n_d}$ 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 $\mathcal{F}$.

1. **Equal Link Split Downstream (ELSD)**
   The cost is equally split among the downstream receivers and no cost is allocated upstream
   $$F_u(n_u,n_d) = 0$$
   $$F_d(n_u,n_d) = \frac{1}{n_d}$$
2. **Equal Tree Split (ETS)**
   The total tree cost is equally split among all the receivers
   $$F_u(n_u,n_d) = F_d(n_u,n_d) = \frac{1}{n_u+n_d}$$

[18] This axiom implies that the only relevant aspect of the problem is the induced cooperative game; all other aspects of the topology are irrelevant.

[19] In case our notation is confusing, the domain of this function is all pairs of nonnegative integers except the pair (0,0).

[20] 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 4 where we need the larger domain of $\{Z_+^2 - (0,0)\}$.

[21] If there are no receivers downstream then this link is not part of the distribution tree and we can ignore its costs.

[22] The second condition in the definition of $\mathcal{F}$ is is unrelated to allocations but is merely used to simplify the expression of Theorem 12.

3. **Cost is charged to upstream receivers**
   $F_u = \frac{1}{n_u}$
   $F_d = 0$

4. **Cost is charged relative to the number of receivers upstream/downstream**
   $F_u = \frac{n_u}{{n_u}^2 + {n_d}^2}$
   $F_d = \frac{n_d}{{n_u}^2 + {n_d}^2}$

3.4 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, or to downstream nodes, or 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 appropriate 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:

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

The Stand-Alone axiom protects the individual; every individual receiver is guaranteed that joining a group can never can never cost more than her unicast cost. Assuming users have the power of choice in their network activities, and assuming some (even minimal) amount of self interest guides them, it is hard to imagine why any user would want to join a shared group of receivers if she risks an increase in her allocated costs. Of course, in a cooperative environment receivers may choose to risk having increased costs if the total cost distributed to the group decreases.

Insisting upon the Stand-Alone axiom, when combined with the basic axioms, means that there is one and only one applicable cost allocation strategy:[23]

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:

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

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.[24]

Axiom 6 *No-Free-Rider:*

The cost allocated to a member is never less than $1/|M|$'th of her unicast cost.

Formal definition:

$|M| * af_\alpha(N,R,M,loc,c) \geq$

[23] This result is closely related to the standard axiomatization of the Shapley Value (see [18, 16]) 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 charged:
$loc(m_\alpha) = R \implies af_\alpha(N,R,M,loc,c) = 0 \quad \forall N,R,M,loc,c$ and $m_\alpha \in M$
Theorem 2 continues to hold if we replace the Stand-Alond 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.

[24] This is much like the Unanimity bound in economics; see [11, 12].

$$\sum_{(c_i,c_j)\in T(N,R,loc(m_\alpha))} c(v_i,v_j) \quad \forall N,R,M,loc,c, \textit{ and } m_\alpha \in M$$

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

Theorem 3 *A cost allocation function satisfying the basic axioms and the No-Free-Rider axiom must satisfy:*

$F_u(n_u,n_d) \le F_d(n_u,n_d) \quad \forall (n_u,n_d) \in Z^2_{++}$.

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 $L' = T(N,R,loc(M))$ and two link costs $c1$ and $c2$:
>
> **if**
>
> $$cf(T(N,R,loc(M')),c1) \ \ge cf(T(N,R,loc(M')),c2) \quad \forall M' \subseteq M$$
>
> **then**
>
> $$af_\alpha(N,R,M,loc,c1) \ \ge af_\alpha(N,R,M,loc,c2) \quad \forall m_\alpha \in M$$

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 forwards 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 set of link costs c, a set of members M and their location function loc. For each subset $M' \subseteq M$ and $m_\gamma \in M'$:
>
> $$\sum_{m_\alpha \in M'} af_\alpha(N,R,M,loc,c) \ \le \ af_\gamma(N,R,M-M'+m_\gamma,loc,c) \ + \sum_{m_\beta \in M'-m_\gamma} af_\beta(N,loc(m_\gamma),loc(M'-m_\gamma),c)$$

Obviously, collusion prevention is a desirable property for cost allocation formulae. Unfortunately, we can prove that no canonical allocation strategy satisfies this axiom:

Theorem 5 *No cost allocation formula satisfying the basic axioms can satisfy the no collusion axiom.*

3.5 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 towards 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 rather outlandish. The only canonical allocation scheme which 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.[25]

[25] This is the standard motivation for the Shapley Value.

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

4 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?

4.1 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,[26] 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 kind of centralized accounting, and so we must turn to a more decentralized approach. 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,[27] 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**.[28] 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.[29] 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,[30] 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 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 $in(v_j)$ 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 $out(v_i, v_j)$ 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.[31] 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, $in(v_j) = c(v_i, v_j) + out(v_i, v_j)$ 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: $out(v_i, v_j) = out(v_i, v_k)$ 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 imple-

[26] The geographic dispersion is important because it means the amount of information needed to describe the tree topology is also growing.

[27] 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.

[28] 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 numbers is available.

[29] From the point of view of implementation of the accounting schemes themselves, model 1 and model 2 are essentially equivalent.

[30] One could, for instance, include information of the previous five upstream links.

[31] It is straightforward to show that if there is one node with a locally unbalanced budget then one cannot guarantee that the overall cost budget is balanced.

mented 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:[32]

Theorem 6 *Given a tree $T(N, R, loc(M))$, a set of link costs c, and a 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.*

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.

4.2 Model 2

As we stated above, we are interested in the extent to which these accounting mechanisms can support cost allocation formulae that obey our previous axioms. As we show below, model 2 can support the basic axioms presented in Section 3.1. 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 6 and 2. How is the ELSD formula implemented? Consider some node v_i in the distribution tree; we let $tmem(v_i)$ denote the number of members located in the sub-tree rooted at v_i and recall that $in(v_i)$ denotes the input costs in the accounting message arriving at v_i.[33] The cost allocated to each local member is $in(v_i)/tmem(v_i)$. For each link (v_i, v_j) in the distribution tree, $out(v_i, v_j) = tmem(v_j)/tmem(v_i) * in(v_i)$ and so $in(v_j) = tmem(v_j)/tmem(v_i) * in(v_i) + c(v_i, v_j)$. 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.

[32] A stronger result holds. No subset $M' \subseteq M$ can be allocated a cost that is greater than their subtree cost: $\sum_{m_\alpha \in M'} af_\alpha(N, R, M, loc, c) \leq cf(T(N, R, loc(M')), c)\ \forall N, R, M, loc, M' \subseteq M$, and c. This means that all one-pass accounting schemes produce results that are in the *core* (see [19] for a definition).

[33] A node v_i knows the value of $tmem(v_j)$ for all immediate downstream members because this is a basic property of model 2. The node can calculate $tmem(v_i)$ by adding up all the $tmem(v_j)$ for all immediate downstream members and then adding the number of local members.

4.3 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 are there.

4.3.1 Reduced Basic Axioms

We would like to invoke the basic axioms presented in Section 3.1. 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*

This theorem can be intuitively explained by observing that the model 1 accounting schemes are *physical-topology dependent* in contrast with the *physical-topology independence* of the Equivalency axiom. Thus, for model 1 we will require that the *reduced* basic axioms of Anonymity and Additivity still hold. In the next section we develop a canonical form for model 1 allocation schemes that obey these axioms. In passing we should mention that these reduced basic axioms lead to the same impossibility result about collusion:

Theorem 10 *No model 1 cost allocation formula satisfying the reduced basic axioms can satisfy the no collusion axiom.*

4.3.2 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. At nodes with no local members, all costs much 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 $in(v_i)$. If there are n_l local receivers and n_r next hop nodes. The allocated costs are:

For a local member on v_i: $F_l(n_l, n_r) * in(v_i)$
For a next hop node v_j : $out(v_i, v_j) = F_r(n_l, n_r) * in(v_i)$

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

Theorem 11 *A cost allocation formula satisfies the reduced basic axioms if and only if it can be expressed in the one-pass canonical form.*

While the one-pass canonical form appears very similar to the canonical form discussed in Section 3.2, 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. 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.

4.3.3 Examples

The one-pass canonical form allows for an infinite set of possible allocation strategies. Below we list a few examples. As before, some strategies may seem intuitive and others may look less than that (like charging upstream receivers). Since we know that $F_l(n_l, 0) = \frac{1}{n_l}$ and $F_r(0, n_r) = \frac{1}{n_r}$ for all $n_l, n_r > 0$, we merely describe the allocations on the set Z^2_{++}.

1. **Local members pay nothing**

   $F_l(n_l, n_r) = 0$
   $F_r(n_l, n_r) = 1/n_r$

2. **Local members pay everything**

   $F_l(n_l, n_r) = 1/n_l$
   $F_r(n_l, n_r) = 0$

3. **All locals are considered as one next hop (ENHS):**

   $F_l(n_l, n_r) = \frac{1}{n_l*(n_r+1)}$
   $F_r(n_l, n_r) = \frac{1}{n_r+1}$

4. **Local members and Next Hops are allocated identical costs:**

   $F_l(n_l, n_r) = F_r(n_l, n_r) = 1/(n_l + n_r)$

5. **Equal Split between all the local members and all the next-hops:**

   $F_l(n_l, n_r) = 1/(2 * n_l)$
   $F_r(n_l, n_r) = 1/(2 * n_r)$

6. **Majority Loses:**

   if $n_l > n_r$ $F_l(n_l, n_r) = 1/n_l$ and $F_r(n_l, n_r) = 0$
   if $n_l < n_r$ $F_r(n_l, n_r) = 1/n_r$ and $F_l(n_l, n_r) = 0$
   if $n_l = n_r$ $F_l(n_l, n_r) = F_r(n_l, n_r) = 1/2n_l$

4.3.4 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 12 *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 $F_l(n_l, n_r)$ are nonincreasing on $\{Z^2_+ - (0, 0)\}$.*

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

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 13 *There is no model 1 cost allocation formula that satisfies the reduced basic axioms and the No-Free-Rider axiom.*

As long as our functions F_l, F_r are nonincreasing (so that the Sharing-is-Good axiom is obeyed) we have little to guide us in our choice of canonical model 1 allocation schemes. Allocating everything to the local members is unfair, as is passing all costs on to the downstream links. In a very rough intuitive sense, fairness seems to dictate that local members individually are allocated no more than is passed on to the individual links ($F_l(n_l, n_r) \leq F_r(n_l, n_r)$) and the set of local members as a whole are allocated at least as much as is passed on to the individual links ($n_l F_l(n_l, n_r) \geq F_r(n_l, n_r)$). There is wide latitude between these two extremes of treating local members each as a downstream link, and treating the collection of them as a downstream link. How can we choose among the various possibilities that lie within the spectrum $\frac{F_r(n_l,n_r)}{n_l} \leq F_l(n_l, n_r) \leq F_r(n_l, n_r)$?

One possibility is to require that the cost allocated to members not depend not on the exact number of members at all other nodes, but only on the set of nodes where there are members. This condition makes little sense for model 2, because there we had the information about individual nodes downstream. However, model 1 does not provide this information. Moreover, the contexts in which model 1 might be deployed may not even have this information at the local level; the costs, rather than being assigned to individuals, might really be assigned to the nodes (subnets) at which members reside. For instance, all costs allocated to members on an ethernet at USC would be assigned to USC rather than to the individual members. This would allow current multicast protocols, which do not keep track of individual members, to be consistent with cost allocation. Given that the number of members at a particular node is not known, the allocation to other members should not depend on it. This requirement means that $F_r(n_l, n_r) = F_r(1, n_r)$ for all $n_l > 0$, and thus narrows the spectrum $\frac{F_r(n_l,n_r)}{n_l} \leq F_l(n_l, n_r) \leq F_r(n_l, n_r)$ down to the single point $\frac{F_r(n_l,n_r)}{n_l} = F_l(n_l, n_r)$, which is merely the ENHS scheme.

We did not embody the considerations in the previous two paragraphs in axioms, because the distributive notions were significantly less general and compelling than our previous axioms. We fully admit that this line of reasoning, which leads to ENHS, is much weaker than our previous results, but we do not see any other general principles at our disposal. This is largely because the allocations that result from model 1 have unavoidably poor properties (free rider, no equivalency, etc.) and so it is hard to formulate desirable distributive notions that are achievable in this context.

4.4 Discussion

In Section 3 we 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. We do present an intuitive line of reasoning that suggests that 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.[34] 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 local 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. 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, and we plan to do so, we believe that it is essentially impossible to achieve Equivalency if allocations must be done without knowledge of 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.

[34] 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.

5 Future Work

This paper used a very simple model to illustrate some of the basic issues involved in cost allocation. In this section we briefly discuss several directions in which our model could be made more general.

1. **Multicast Distribution Model**

   This paper modeled multicast distributions as source rooted trees, computed from unicast routing. Although this is the most common form of multicast distribution, newer protocols use different models for their distribution. CBT [1] advocates using a shared tree for all sources (core rooted rather than source rooted trees). PIM (Protocol Independent Multicast) [6], in its sparse mode, enables mixing source rooted, shortest path (SP) routes for some groups and shared trees for others. Our theory only assumes that the route taken from a particular source to a particular receiver is independent of the group membership (i.e., the route will be the same to a specific group member regardless of who else has joined the group). This remains true for both CBT and PIM, and thus our results apply.[35] Another aspect of routing which can invalidate our theory is the ability to request alternate routes. Typically these alternate routes are not automatically loop-free, and so when an a new alternate route is being established the route is followed until it hits the tree. This means that the resulting routes will depend on the current membership.

2. **Multiple Sources**

   Our model considered only a single source. If the "costs" being allocated are tied to a per-packet metric, then this is sufficient since each packet comes from a single source. This is not sufficient, however, if the "costs" are tied to resource reservations (e.g., bandwidth). RSVP [21] uses the idea of reservation styles, and two of these styles, *dynamic filter* and *wildcard filter*, allow a single reservation to be shared by several sources. In terms of our theory, this would require the root R to be a set of nodes rather than a single node. This, by itself, introduces no nontrivial changes to our theory. However, if the costs on a link depended on the particular set of upstream sources (say, for instance, that the bandwidth reserved was the maximal upstream source bandwidth) then our theory needs to be extended.

[35] Although this assumes that the route used when computing the stand-alone cost is chosen using multicast, rather than reverting to a unicast route.

3. **Multiple Qualities of Service**

   In this paper, all members are equivalent (except for location). However, if the network supports multiple qualities of services (say, for instance, several priority levels or several levels of real-time service) and these different QoS levels incur different network costs, then the members are no longer necessarily equivalent. This raises an issue that is in some sense orthogonal to the one considered here. In this paper we considered 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.[36] Clearly, for these different QoS requests to be merged onto a single multicast distribution tree, there must be at least a partial set of ordering relations between the QoS's, and the QoS installed on the link must be greater than or equal to all QoS requests. 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 [7]) 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 [13]) and appears to be a rather natural choice. If the set of QoS's is not completely ordered but only partially ordered, then the problem becomes much more complicated.

   Multiple qualities of service also present some problems for 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, if there is a continuum of QoS levels (e.g., bandwidth) then we must rethink the form of the accounting message.

4. **Other Issues**

   There are several other areas that we have not yet explored. We did not discuss the issue of dynamics. Clearly in a large multicast group the membership can change rather rapidly, and these changes affect the costs allocated to individual members. Is there some principled way to describe and ameliorate these changes? Also, we have not at all considered the issue of incentives. Given some model of strategic behavior, what equilibria result from the various forms of cost allocation methods proposed here?

References

[1] A.J. Ballardie, P.F. Francis and J. Crowcroft Core Based Trees (CBT). *Proceedings of ACM SIGCOMM '93, San-Francisco*, 1993.

[2] H.W. Braun and K. Claffy. A Framework for Flow-Based Accounting on the Internet. *Proceedings of SICON '93, IEEE Singapore*, 1993.

[3] D.D. Clark, S. Shenker and L. Zhang Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism. *Proceedings of ACM SIGCOMM '92*, August 1992.

[4] R. Cocchi, S. Shenker, D. Estrin, and L. Zhang Pricing in computer networks: Motivation, Formulation, and Example. *IEEE/ACM Transactions on Networking*, December 1993.

[5] S. Deering and D. Cheriton Multicast Routing in Datagram Internetworks and Extended LANs. *ACM Transactions on Computer Systems*, pages 85–111, May 1990.

[6] S. Deering, D. Estrin, D. Farinacci, V. Jacobson, C.G. Liu and L. Wei An Architecture for Wide-Area Multicast Routing. *ACM SIGCOMM '94, London*, September 1994.

[7] S.C. Littlechild and G. Owen A Simple Expression for the Shapley Value in a Special Case. *Management Science, 20, 370-2*, 1973.

[8] J.K. MacKie-Mason and H.R. Varian. Pricing the Internet. *In B. Kahin and J. Keller (Eds.),* Public Access to the Internet. *Prentice-Hall, Englewood Cliffs, New Jersey. Available from ftp://gopher.econ.lsa.umich.edu/pub/Papers*, May 1993.

[9] J.K. MacKie-Mason and H.R. Varian. Pricing Congestable Network Resources. *Preprint*, October 1994.

[10] N. Megiddo Computational Complexity of the Game Theory Approach to Cost Allocation for a Tree. *Mathematics of Operations Research, 3, 189-96*, 1978.

[11] H. Moulin Uniform Externalities: Two Axioms for Fair Allocation. *Journal of Public Economy 43 (1990) p305-326.*

[12] H. Moulin Welfare Bound in the Cooperative Production Problem. *Games and Economic Behavior Volume 4, (1992) p 373-401.*

[13] H. Moulin and S. Shenker Average Cost Pricing versus Serial Cost Sharing: An Axiomatic Comparison. *Journal of Economic Theory 64(1), (1994) p. 178-201.*

[14] R.B. Myerson Game Theory. *Harvard University Press, Cambridge, Massachusetts*, 1991.

[15] H. Peyton and I. Young (Eds.). Cost Allocation: Methods, Principles, Applications. *Elseviers Science Publishers, B.V., Amsterdam, The Netherlands and New York, N.Y., U.S.A.*, 1985.

[16] A.E. Roth (Ed.). The Shapley Value, an Essay in Honor of Lloyd S. Shapley. *Cambridge University Press, Cambridge*, 1988.

[17] J. Sairamesh, D.F. Ferguson and Y. Yemini. An Approach to Pricing, Optimal Allocation and Quality of Service Provisioning in High-Speed Packet Networks. *To appear in Infocom '95*, 1995.

[36] Combining the two issues to have different QoS's at different locations can only be done after one has solved the two problems independently.

[18] L.S. Shapley A Value for N-Person Games. *In 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 University Press, 307-17,* 1953.

[19] W.W. Sharkey Network Models in Economics. *Handbook of Operations Research and Management Science: Networks*, to appear.

[20] S. Shenker. Service Models and Pricing Policies for an Integrated Services Internet. *In B. Kahin and J. Keller (Eds.),* Public Access to the Internet*. Prentice-Hall, Englewood Cliffs, New Jersey.*

[21] L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala RSVP: A New Resource Reservation Protocol. *IEEE Networks Magazine*, September 1993.

6 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 $L' = T(N, R, loc(M))$. 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

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

and so

$$\sum_{(i,j):(v_i,v_j)\in L} (af_\alpha(N, R, M, loc, c_{i,j}) * c(v_i, v_j)) = af_\alpha(N, R, M, loc, c) \qquad \forall N, R, loc, M \text{ and } \alpha \in M$$

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:

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

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. For example, Figure 2 shows nodes v_1, v_2, v_3 merge to a single node v_2 in the reduced network. The anonymity condition requires that equivalent members be charged the same; in particular, 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 non-negative, 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.[37] QED

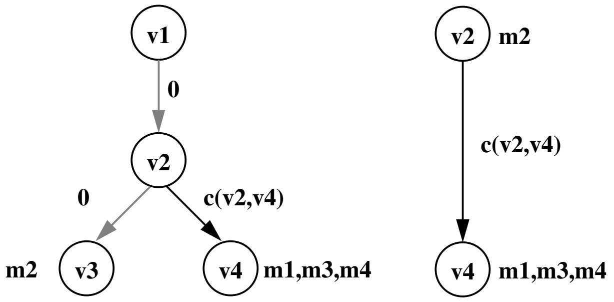

Figure 2: Merging nodes with no cost link: the Equivalency axiom

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 $F_d(n_u, n_d) = 1/n_d \ \ \forall n_u > 0$ 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. QED

Theorem 3 *A cost allocation function satisfying the basic axioms and the No-Free-Rider axiom must satisfy:*

$$F_u(n_u, n_d) \le F_d(n_u, n_d) \qquad \forall (n_u, n_d) \in Z^2_{++}.$$

Consider any canonical form F. The Free Rider axiom is obeyed if and only if $\frac{1}{n_u+n_d} \le F_d(n_u, n_d)$ 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 $n_u(F_d(n_u, n_d) - F_u(n_u, n_d)) \ge 0$. Thus, whenever $(n_u, n_d) \in Z^2_{++}$, we must have $F_d(n_u, n_d) \ge F_u(n_u, n_d)$. QED

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. Figure 3 provides an example that follows throughout the proof. Consider any network $N = (V, L, T)$ with a root R, set of members M, location function loc, and cost function c (case a). From it, we buld a second network (illustrated in case b), by adding a new link $l' = (R', R)$ with no cost:

[37] The conditions $F_u(0, n_d) = 1$ and $F_d(n_u, 0) = 1$ in the definition of $\mathcal{F}$ are irrelevant to the actual allocations.

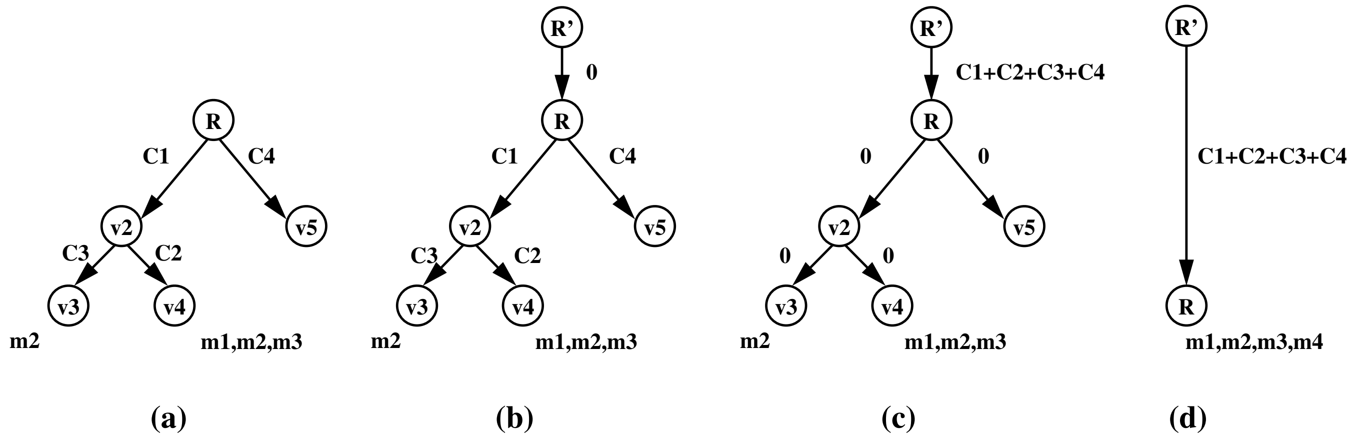

Figure 3: Equal Tree Split transformations: cases a,b,c,d

$N' = (V + R', L + l', T')$
$L'' = T'(N', R', loc(M)) = L' + l'$
$c' = c + \{c'(R', R) = 0\}$.

Clearly all of the subset costs are the same under case a and case b, so allocations must not change (due to the Equivalency axiom). We now keep the network fixed and consider a new cost function (illustrated as case c) where the cost of the whole tree is on (R',R) while the remaining links have zero costs.

$$\begin{array}{rcll} c''(R',R) & = & cf(L',c') & \\ c''(v_i,v_j) & = & 0 & \forall (v_i,v_j) \in L' \end{array}$$

Note that no subset cost have decreased, therefore the subset monotonicity condition implies that no individual allocated costs can decrease. Since the total costs have remained the same, the balanced budget implies that total allocated costs must remain unchanged. Under these two constraints, it must be that the allocations for this new cost function must be the same as for the old cost function. Case d is created by merging all the nodes downstream into R, which, according to the equivalency axiom does not change the allocated costs. With all the members in one node (R), the anonymity axiom implies that they all share the cost of the link equally, so, each member is allocated $c''(R',R)/|M| = cf(L',c)/|M|$. This is the ETS policy. QED

Theorem 5 *No cost allocation formula satisfying the basic axioms can satisfy the no collusion axiom.*

Consider any canonical form F and assume that is satisfies the no collusion axiom. Consider the network with a single link, with n_u members at the root and n_d members downstream. Because $5(F_u(5,5) + F_d(5,5)) = 1$ we either have $F_u(5,5) \geq 1/10$ or $F_d(5,5) \geq 1/10$. Since our proof works with either case, let us assume without loss of generality that $F_u(5,5) \geq 1/10$. The no collusion condition requires that $F_u(n_u - 1, n_d) \geq 2 * F_u(n_u, n_d)$ for any n_u, because otherwise one element would collude with another. Recursing this inequality, we find: $F_u(n_u - k + 1, n_d) \geq 2^{k-1} * F_u(n_u, n_d)$. Choose $n_u = n_d = k = 5$, and then we find that $F_u(1,5) \geq 2^4 * F_u(5,5) \geq 1.6$. This is a violation of the canonical form, thus a contradiction. QED

Theorem 6 *Given a tree $T(N,R,loc(M))$, a set of link costs c, and a one-pass mechanism, no member can be allocated a cost greater than her unicast cost.*

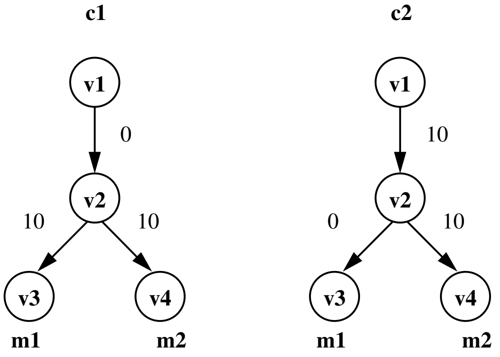

Figure 4: Subset monotonicity vs. One-Pass model

From the definition of model 1 we know that: (1) nodes are budget balanced so their output and allocations are bounded by their input, and (2) $in(v_j) = c(v_i,v_j) + out(v_i,v_j)$ for each v_j downstream to v_i. Thus, $in(v_j) \leq c(v_i,v_j) + in(v_i)$. If we iterate this inequality for each node along the path from R to any node v_k, and collapse the inequalities, we get the following result: $in(v_k) \leq cf(T(N,R,v_k),c)$. QED

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 Figure 4 of a single tree $T(N,R,loc(M))$ with two different link costs: $c1$, $c2$. The cost allocations for $c1$ are straightforward: $af_1() = af_2() = 10$. 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 $c1$ and $c2$. To achieve this allocation for $c2$ it must be that $out(v_2,v_3) = 10$. But we must have $out(v_2,v_3) = out(v_2,v_4)$ which violates the local budget balance. Contradiction. QED

Theorem 9 *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 figure 5, where $c(v_1,v_2) > 0$ and

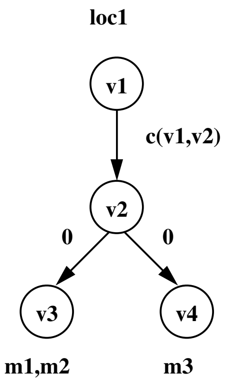

Figure 5: One-Pass, model 1 vs. Equivalency

$c(v_2, v_3) = c(v_2, v_4) = 0$. The Equivalency axiom implies that the cost allocations to members on v_3 and v_4 should be the same: $af_1() = af_2() = af_3()$. However, the definition of model 1 implies that $in(v_3) = in(v_4) = \frac{c(v_1,v_2)}{2}$, Therefore, the cost allocations cannot be the same. Contradiction. QED

Theorem 10 *No model 1 cost allocation formula satisfying the reduced basic axioms can satisfy the no collusion axiom.*

This proof is identical to the one in theorem 5 replacing (F_u, F_d) by (F_l, F_r), n_u by n_l and n_d by n_r. QED

Theorem 11 *A 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 $L' = T(N, R, loc(M))$. 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 the 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:

$$\begin{array}{ll} \text{Local members} & : in(v_i) * F_l(n_l, n_r) \\ \text{Residual costs to node } v_k{}^{38} & : out(v_j, v_k) = in(v_i) * F_r(n_l, n_r) \end{array}$$

In model 1 the input costs to the next hop v_k include both residual cost and the cost of the link connecting them:

[38] In the One-Pass model, residual costs are costs that are allocated individually to each immediate downstream neighboring node

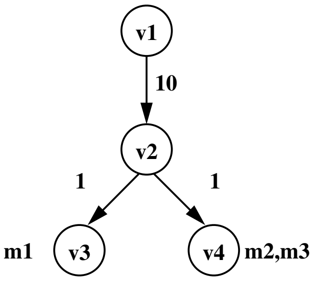

Figure 6: No-Free-Rider axiom in model 1

$in(v_k) = in(v_i) * F_r(n_l, n_r) + c(v_i, v_k)$. This is precisely the One-Pass canonical form. QED

Theorem 12 *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 $F_l(n_l, n_r)$ are nonincreasing on $\{Z_+^2 - (0,0)\}$.*

It is straightforward to show that the Sharing-is-Good axiom conditions are exactly identical to the no-increasing formulae conditions. When adding a local member to a node already on the distribution tree, we must have:

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

When adding a member to a node which was not on the distribution tree, at the nearest node on the tree we must have:

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

When we combine these equations with the second condition in the definition of $\mathcal{F}$ we get the non-increasing condition on all of $\{Z_+^2 - (0,0)\}$. QED

Theorem 13 *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 11 we have shown that all model 1 cost allocation formulae obeying the reduced basic axioms can be expressed in the one-pass canonical form. Consider the tree detailed in figure 6: with no members in v_2, $F_r(0,2) = 1/2$ for all formulae, resulting in $in(v_4) = 6$. With no next-hops downstream of $v4$, $F_l(2,0) = 1/2$ for all formulae, making the allocation to each member on v_4 exactly 3 which is less than 11/3 (unicast cost/$|M|$). QED