

# The Impact of Multicast Layering on Network Fairness

Dan Rubenstein, *Member, IEEE*, Jim Kurose, *Fellow, IEEE*, and Don Towsley, *Fellow, IEEE*

**Abstract**—Many definitions of fairness for multicast networks assume that sessions are single-rate, requiring that each multicast session transmits data to all of its receivers at the same rate. These definitions do not account for multi-rate approaches, such as layering, that permit receiving rates within a session to be chosen independently. We identify four desirable fairness properties for multicast networks, derived from properties that hold within the max-min fair allocations of unicast networks. We extend the definition of multicast max-min fairness to networks that contain multi-rate sessions, and show that all four fairness properties hold in a multi-rate max-min fair allocation, but need not hold in a single-rate max-min fair allocation. We then show that multi-rate max-min fair rate allocations can be achieved via intra-session coordinated joins and leaves of multicast groups. However, in the absence of coordination, the resulting max-min fair rate allocation uses link bandwidth inefficiently, and does not exhibit some of the desirable fairness properties. We evaluate this inefficiency for several layered multi-rate congestion control schemes, and find that, in a protocol where the sender coordinates joins, this inefficiency has minimal impact on desirable fairness properties. Our results indicate that sender-coordinated layered protocols show promise for achieving desirable fairness properties for allocations in large-scale multicast networks.

**Keywords**—Multicast, fairness, congestion control

## I. INTRODUCTION

THE current Internet has few internal mechanisms to regulate the rates at which sessions should transmit data. How to achieve *fairness* within such a network, in effect allowing sessions to share bandwidth in a manner that satisfies some set of network utilization criteria, remains a challenging research problem. The problem is further complicated in networks that support both unicast and multicast delivery services. Current definitions of multicast fairness [3], [6], [15], [23], [25] typically assume that sessions are *single-rate*, requiring all receivers within a multicast session to receive data at a uniform rate. However, *layered multicast* permits *multi-rate* transmission: different receivers within a session can receive data at different rates. This is accomplished by layering data among several multicast groups and allowing each receiver to determine the subset of layers (i.e., multicast groups) it joins. Protocols have used a layered approach to support multicast applications ranging from live multimedia [1], [10], [11], [13] to reliable data transfer [4], [16], [24]. These protocols have the appealing property that the transmission rate to each receiver is constrained only by the bandwidth availability on the receiver's own data-path from the data source, and is not limited by other receivers' rate limitations in the same session. What is lacking in this previous work

is a formal study that examines the impact that layering has on fair allocations within a large-scale multicast network.

In this paper, we contribute to the formal understanding of how layering impacts fairness in multicast networks. In particular, we focus on how layering affects properties of multicast max-min fairness in an environment in which each session has a single sender. We have chosen to use max-min fairness as our fairness criterion since its formal definition is a well-accepted criterion for fairness, allowing us to proceed directly to an examination of the properties of a fair allocation. We believe that with other definitions of fairness, layered approaches will yield similar fairness advantages, and expect this work to stimulate interest in examining the impact of layering in the context of these other definitions.

Our examination begins with a theoretical and idealized model of multi-rate sessions that does not account for the current practical limitations of layered approaches (e.g., limited number of multicast groups, pre-determined rates of layers) to achieve multi-rate max-min fairness. Hence, the results based on this model demonstrate the *potential* fairness benefits that can be gained through the use of multi-rate sessions. A utility function is associated with each session that maps each receiver's receiving rate to a utility for that receiver. The fairness of an allocation is measured by comparing the relative utilities obtained by receivers in the network. We show that in such networks, allowing multicast sessions to be multi-rate instead of single-rate leads to additional desirable fairness properties within the max-min fair allocation. One such property is that receiver utilities should be equal for two receivers whose data transmission paths from their respective senders traverse an identical set of links. We examine multicast max-min fair allocations under the definition given by Tzeng and Siu [23], that requires that all sessions be single-rate, and find that several of these fairness properties do not necessarily hold within the max-min fair allocation (the two receiver utility example presented above being one such property). We extend the multicast max-min fair definition to permit multi-rate sessions, and formally prove that, when all sessions in a network are multi-rate, all of our identified fairness properties hold for the max-min fair allocation. We also consider networks in which not all sessions are multi-rate (e.g., a session may have an application-specific requirement that requires it to be single-rate), and examine the effect on fairness properties of the max-min fair allocation as single-rate sessions are "replaced" by identical multi-rate sessions (i.e., same session members, same topology). Using our identified set of fairness properties and a mathematical lexicographic ordering relation of allocations that indicates an allocation's "level" of max-min fairness, we demonstrate that increasing the set of "replaced" sessions results in an increase in the "level" of max-min fairness and that more fairness properties hold for max-min

This material was supported in part by the National Science Foundation under Grant No. NCR-9508274, NCR-9527163, CDA-9502639, and by DARPA under Grant No. N66001-97-C-8513. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

D. Rubenstein is with the Department of Electrical Engineering at Columbia University, New York, NY, 10027 (e-mail: danr@ee.columbia.edu).

J. Kurose and D. Towsley are with the Department of Computer Science, University of Massachusetts, Amherst, MA, 01003 (e-mail: kurose@cs.umass.edu, towsley@cs.umass.edu).

fair allocations. Our results extend easily to the rate-based definition of max-min fairness instead of the utility-based definition presented in this paper by simply setting a receiver's utility to be equal to its rate.

Next, we explore how permitting the use of different utility functions by receivers within the same session impacts our results. We examine the properties of the max-min fair allocation in these scenarios, and find that our desirable fairness properties either do not hold or are not well-defined if we permit the use of different utility functions by receivers within the same session.

Next, we examine the impact of some current practical limitations of layering on the fairness properties of multi-rate sessions. We show that if each receiver's fair rate is restricted to what can be obtained by joining some fixed set of layers, a max-min fair allocation need not even exist. However, we do demonstrate that receivers can achieve an *average* rate that matches their fair rate by using carefully timed joins and leaves. These joins and leaves must be tightly coordinated among receivers in the same session (i.e., correlating their sets of received packets) in order to prevent excess bandwidth utilization on a shared link. To quantify bandwidth usage, we introduce the notion of *link-overuse*,<sup>1</sup> a ratio of bandwidth used in practice by a session on a shared link to the theoretical lower bound needed on that link to deliver fair rates to downstream receivers. While several works have indirectly identified the negative implications of link-overuse, a link-overuse measure has never been formally defined, and its effect on fair allocations within a network has never been studied directly. We show that increased link-overuse leads to a decrease in the "level" of max-min fairness, to a decrease in the number of fairness properties that hold for the max-min fair allocation, and, usually, to a decrease in receivers' fair rates. We study how the ideas in [10], [13], [24], that coordinate joins of receivers within a session, significantly reduce the negative effects of link-overuse. The study is performed via analytical modeling and simulation of max-min fair congestion control protocols in which receivers join and leave layers based on congestion observations. Within the model, we present three protocols that differ in the degree to which the layer joins are coordinated among session receivers. We find that, although link-overuse is still not optimal, coordinated joins reduce link-overuse most significantly when the correlation in loss among receivers is high, and that a protocol with sender coordination keeps link-overuse at low enough levels to allow layered multicast to achieve fairness within a multi-rate multicast network while making more efficient use of network bandwidth.

This paper makes two fundamental contributions to network protocol design. First, it formally demonstrates the theoretical benefits in terms of fairness of using multi-rate (i.e., layered) sessions, and that these benefits also exist in networks that support a mix of multi-rate and single-rate sessions. Second, we formally identify and define the link-overuse of layered protocols, and demonstrate the drawbacks (in terms of fairness and efficiency of using available bandwidth) of having high link-overuse. This suggests that future research geared toward improving layered protocol performance for multicast (e.g., new

<sup>1</sup>What we now call *link-overuse* was called *redundancy* in previous versions of this work. We have changed terminology to avoid overloading the term redundancy, that has a different meaning in an information-theoretic context.

TABLE I  
VARIABLES USED IN THE NETWORK MODEL

$N$	a network, $(G, \{S_1, \dots, S_m\}, \tau, \sigma)$
$G$	A network graph with $n$ links.
$l_j, 1 \leq j \leq n$	The $j$ th link of $N$
$c_j, 1 \leq j \leq n$	The capacity of link $l_j$
$S_i, 1 \leq i \leq m$	The $i$ th session in $N$
$\sigma()$	a mapping onto each session $S_i$ that indicates the session's type ( $\mathcal{M}$ = multi-rate or $\mathcal{S}$ = single-rate)
$r_{i,k}$	The $k$ th receiver in session $S_i$
$x_i$	the single sender for session $S_i$
$\tau()$	A topology mapping that maps session members onto network nodes.
$R_{i,j}$	The set of receivers in $S_i$ whose route traverses $l_j$
$R_j$	The set of receivers over all sessions whose route traverses $l_j$
$d_{i,k}$	The data rate for transmission to receiver $r_{i,k}$
$\mu_i()$	Session $S_i$ 's utility function
$a_{i,k}$	Receiver $r_{i,k}$ 's utility, where $a_{i,k} = \mu_i(d_{i,k})$
$r_i$	The receiver in a unicast session $S_i$
$d_i$	The data rate in a unicast or single-rate session
$a_i$	Unicast or single-rate session $S_i$ 's utility, where $a_i = \mu_i(d_i)$
$m_i$	The maximum desired rate for session $S_i$ , $0 < \alpha_i \leq \infty$
$\alpha_i$	The maximum session utility, where $\alpha_i = \mu_i(m_i)$
$f_{i,j}$	The link rate for session $S_i$ on link $l_j$
$f_j$	The link rate for link $l_j$ (i.e., $\sum_i f_{i,j}$ )
Defined in Section III:	
$\rho$	The aggregate rate of the "single-layer"
$v_i$	A more general session link rate function

layering or routing protocols) should aim to keep the level of link-overuse low.

The paper proceeds as follows. Section II presents theoretical results for multicast utility-based max-min fairness with multi-rate sessions. Section III introduces the notion of link-overuse, and Section IV examines the effects of join coordination in several simple congestion control protocols. Section V presents related work, and we conclude in Section VI.

## II. MULTI-RATE MULTICAST MAX-MIN FAIRNESS

In this section, we present the formal network model used to examine a utility-based max-min fairness of multicast sessions, and identify a set of desirable fairness properties derived from a set of desirable properties exhibited within the max-min fair allocations for unicast networks. We then show that in this network model, the max-min fair allocation will always achieve all of these desirable properties only if the sessions are multi-rate. For the reader's convenience, a list of all the variables used is provided in Table I.

A session,  $S_i$ , is a tuple  $(x_i, \{r_{i,1}, \dots, r_{i,k_i}\})$  of session members:  $x_i$  is the session sender that transmits data within a network; each  $r_{i,k}$  is a receiver that receives data from  $x_i$ . Each session contains exactly one sender and at least one receiver. We write  $r_{i,s} \in S_i$  to indicate that receiver  $r_{i,s}$  is a member of session  $S_i$ .<sup>2</sup> We consider two types of sessions:

<sup>2</sup>We assume that each receiver is a member of a single session. A receiver that is a member of two sessions can simply be viewed as two distinct receivers.

- If  $S_i$  is a *single-rate* session, then data must be transmitted to all receivers in  $S_i$  at the same rate.
- If  $S_i$  is a *multi-rate* session, then receivers within  $S_i$  can receive data at multiple (arbitrary) rates.

A *network graph*,  $G$ , consists of a set of nodes connected together by  $n$  links in some arbitrary fashion. The links are labeled  $l_1, \dots, l_n$ . Each link  $l_i$  has a *capacity*,  $c_i$ , that limits the aggregate rate of flow it can transmit in either direction between the two nodes it connects.<sup>3</sup> We define a network,  $N = (G, \{S_1, \dots, S_m\}, \tau, \sigma)$  to be a tuple containing a network graph,  $G$ , a set of sessions,  $\{S_1, \dots, S_m\}$ , a mapping,  $\tau$ , that maps each member of each session to a node in the network graph, and a second mapping,  $\sigma$ , that maps each session  $S_i$  to its type. We write  $\sigma(S_i) = \mathcal{S}$  to indicate that session  $S_i$  is single-rate, and  $\sigma(S_i) = \mathcal{M}$  to indicate that session  $S_i$  is multi-rate.

The mapping,  $\tau$ , of a session onto the network graph has one restriction: no two members of a single session are mapped to the same node. However, there is no restriction that forbids two members of different sessions to be mapped to the same node. The network employs a routing algorithm, such that for each receiver  $r_{i,k} \in S_i$ , there is a sequence of links  $(l_{j_1}, \dots, l_{j_n})$  that carries data from  $x_i$  to  $r_{i,k}$ . We refer to this set of links in this sequence as the receiver's *route*. The route for a session is defined to be the set of all links that carry data to any receiver within the session.

For a network  $N$ , we define  $R_{i,j}$  to be the set of receivers in session  $S_i$  whose route includes link  $l_j$ , and define  $R_j$  to be the set of all receivers whose route includes link  $l_j$ , i.e.,  $R_j = \cup_i R_{i,j}$ . An *allocation* is an assignment of receiver rates within a network. Once an allocation has been determined, we use  $d_{i,k}$  to represent the rate at which data is transmitted to receiver  $r_{i,k}$  (that equals the rate at which the data is received by  $r_{i,k}$ , barring loss). We let  $f_{i,j}$  represent an absolute measure of bandwidth (e.g., in bytes/sec) used by session  $S_i$  on link  $l_j$  to transmit data to its receivers, and *linkrate $j$*  the amount of bandwidth used by all sessions across link  $j$ ,  $f_j = \sum_{i=1}^m f_{i,j}$ . We refer to  $f_{i,j}$  as the *session link rate* of  $l_j$  for session  $S_i$ , and  $f_j$  simply as the *link rate* of  $l_j$ . Since bandwidth for each flow is non-negative, we have  $0 \leq f_{i,j} \leq f_j$ . We say a link is *fully utilized* if the total bandwidth used by all sessions across the link matches its capacity, i.e.,  $l_j$  is fully utilized iff  $f_j = c_j$ .

We require that  $f_{i,j} \geq d_{i,k}$  whenever  $r_{i,k} \in R_j$ , i.e., any bandwidth received by a receiver must traverse its route. In this section, we make an additional assumption that  $f_{i,j} = \max\{d_{i,k} : r_{i,k} \in R_{i,j}\}$ , which is the minimum value for  $f_{i,j}$  that satisfies the above requirement. The reader should note that, if there is no restriction on the number of layers that a session can use, such a session link rate is easily achieved using a layered approach. In later sections, we examine the implications if  $f_{i,j}$  is larger than this value. The assumption also allows us to model a unicast session as either a multi-rate session with a single receiver, or as a single-rate session with a single receiver. Thus, any results given in this section for networks containing a mix of single-rate and multi-rate sessions also holds for networks that contain a mix of single-rate, multi-rate, and unicast sessions.

<sup>3</sup>Assigning capacity per direction is a simple extension: simply extend a bidirectional link into two unidirectional links.

We note that different network applications can have differing bandwidth requirements to support a given "level" of quality for an application. In practice, there is currently no easy way to compare this "level" of quality among all applications that might utilize a network. However, in theoretical studies such as the one discussed in this section, one common approach to dealing with this issue is to associate a *utility function* with each session that maps the session's transmission rate to a *utility* [7], [8], [21], [1], [5]. The utility is a uniform measure of quality across all sessions of the network. It facilitates comparisons of the "level" of quality across various types of sessions in the network: we say that session  $A$  derives a higher level of satisfaction than session  $B$  if session  $A$ 's utility is larger than session  $B$ .

A *utility function*,  $\mu_i$ , is associated with each session,  $S_i$ , which maps the receiver's receiving rate,  $d_{i,k}$ , to its utility, which we define as  $a_{i,k}$ . More formally,  $a_{i,k} = \mu_i(d_{i,k})$ . We assume that the utility functions are monotonically increasing, such that any increase in a receiver's rate results in an increase in its utility. As such, the functional inverse,  $\mu_i^{-1}$  is well-defined, and  $d_{i,k} = \mu_i^{-1}(a_{i,k})$  holds. We account for the fact that session  $S_i$  might have a maximum rate constraint,  $m_i$ , at which it will transmit data ( $m_i$  can be infinite). The maximum desired utility for a session is written as  $\alpha_i = \mu_i(m_i)$ . Note that if we assume that each receiver's utility equals its receiving rate, i.e.,  $a_{i,k} = d_{i,k}$  and  $\alpha_i = m_i$ , then results presented here reduce to the results of our earlier work in [17].

An allocation is *feasible* if each receiver  $r_{i,k}$  is assigned a rate  $0 \leq d_{i,k} \leq m_i$ , and all receivers can receive at these rates without overutilizing any link's capacity in the network, i.e.,  $\forall i, k, 0 \leq d_{i,k} \leq m_i$ , and  $\forall j, f_j \leq c_j$  (Hence, in this section, we require  $f_j = \sum_i f_{i,j} = \sum_i \max_{\{d_{i,k} : r_{i,k} \in R_{i,j}\}} d_{i,k} \leq c_j$ ). The additional requirement imposed on each single-rate session  $S_i$  that all of its receivers' rates must be equal means that for any pair of receivers,  $r_{i,k}, r_{i,k'} \in S_i$ , when  $\sigma(S_i) = \mathcal{S}$ , then  $d_{i,k} = d_{i,k'}$ . When  $S_i$  is a single-rate session, or a session of either type containing a single receiver (i.e., a unicast session), we can write the single rate at which all receivers within the session receive data simply as  $d_i$ . We stress that in this section, receiver rates in multi-rate sessions are not constrained by practical limitations of layering. In effect, one can assume that a multi-rate session has at its disposal an unlimited supply of multicast groups, and can configure the rates on the layers to the exact needs and desires of its receivers. We say that a set of receivers' utilities is feasible if the associated set of rates produces a feasible allocation.

Note that the feasibility of a particular allocation of receiver rates is a function of the link capacities of the network graph,  $G$ , the mapping  $\tau$ , and also of the mapping  $\sigma$ . The dependence of an allocation's feasibility on  $\sigma$  is important: we will be examining how varying  $\sigma$  (i.e., varying sessions' types between single-rate and multi-rate) affects which allocation within a network is maximum fair.

**Definition 1: ([Multicast] Util-Max-min Fairness)** An allocation of receiver rates is said to be **util-max-min fair** if it is feasible, and for any alternative feasible allocation of rates (where for each receiver  $r_{i,k}$  we define  $\bar{d}_{i,k}$  as an alternative feasible rate and  $\bar{a}_{i,k}$  as its utility for that rate) where  $\bar{a}_{i,k} > a_{i,k}$ , there is some other receiver  $r_{i',k'} \neq r_{i,k}$  such that  $a_{i,k} \geq a_{i',k'} > \bar{a}_{i',k'}$ .

In other words, if any receiver  $r_{i,k}$ 's rate is increased beyond

its util-max-min fair rate to obtain some other feasible allocation, then there is some other receiver whose util-max-min fair utility is no larger than that of  $r_{i,k}$ , and whose adjusted utility (to account for the increase in  $r_{i,k}$ 's rate) must be decreased.

When all sessions within  $N$  are single-rate (i.e.,  $1 \leq i \leq m$ ,  $\sigma(S_i) = \mathcal{S}$ ), we say that  $N$  is a single-rate network, and the util-max-min fair allocation is called the single-rate util-max-min fair allocation. A similar naming convention holds when all sessions are multi-rate. The definition of max-min fairness in [23] holds only for single-rate networks in which each session's utility equals its rate ( $a_{i,k} = d_{i,k}$ ),<sup>4</sup> and involves a comparison of session rates rather than of receiver rates as in our definition. It is easy to show that when each session's utility equals its rate ( $a_{i,k} = d_{i,k}$ ), then the util-max-min fair allocation in a single-rate network is identical under both definitions. In a network that contains multi-rate sessions or where utility functions can vary for different sessions, the definition in [23] is not well defined.

1.  $T_0 = \{r_{i,k}\}; \forall r_{i,k}, a_{i,k}^0 = 0; \forall i, j, f_{i,j}^0 = 0, f_j^0 = 0; b = 0$
2. While  $|T_b| > 0$
3.  $t_{b+1} = \sup\{t : \forall j, f_j^b + \sum_i \phi_{i,j}(T_b) \mu_i^{-1}(t) \leq c_j \wedge \forall r_{i,k} \in T_b \Rightarrow a_{i,k}^b + t \leq \alpha_i\}$  where  $\phi_{i,j}(T) = \begin{cases} 1 & |R_{i,j} \cap T| > 0 \\ 0 & \text{otherwise} \end{cases}$
4.  $\forall r_{i,k} \in T_b, a_{i,k}^{b+1} = a_{i,k}^b + t_{b+1}$ . For all other  $r_{i,k}, a_{i,k}^{b+1} = a_{i,k}^b$ .
5.  $f_{i,j}^{b+1} = \sum_{r_{i,k} \in R_j} \mu_i^{-1}(a_{i,k}^{b+1}), f_j^{b+1} = \sum_i f_{i,j}^{b+1}$ .
6.  $T' = T_b - \{r_{i,k} \in T_b : a_{i,k}^{b+1} = \alpha_i \vee (\exists j, r_{i,k} \in R_{i,j} \wedge f_j^{b+1} = c_j)\}$
7.  $T_{b+1} = T' - \{r_{i,k} \in T' : \sigma(S_i) = \mathcal{S} \wedge \exists r_{i,k'} \notin T'\}$
8.  $b++$
9. end while
10.  $\forall r_{i,k}, a_{i,k} = a_{i,k}^b, \forall i, j, f_{i,j} = f_{i,j}^b, f_j = f_j^b$

Fig. 1. An algorithm that generates the util-max-min fair allocation.

Just as there is always one and only one unicast max-min fair allocation [2] and one and only one single-rate max-min fair allocation [23], there is one and only one multi-rate util-max-min fair allocation. In fact, for any choice of  $\sigma$ , the network has one and only one util-max-min fair allocation. An algorithm that yields the util-max-min fair allocation is given in Figure 1. We include the algorithm in the paper simply to demonstrate the process by which a multi-rate max-min fair allocation can be computed, and do not claim that the algorithm provides a practical means to generate max-min fair allocations in real networks. More recently, a distributed version of the algorithm was developed and presented in [19].

The algorithm iterates over a set of receivers, each step increasing those receivers' rates uniformly as much as possible without overutilizing any links in the network. A receiver is removed from this set once some link on its route reaches full capacity, or, if the receiver is part of a single-rate session, the route of some receiver in the session contains a link that has reached full capacity.

<sup>4</sup>[23] also permits a multicast session to consist of distinct unicast connections. We model this inherently via separate unicast sessions. Such a session differs significantly from a multi-rate session achieved through layering.

Step 3 of the algorithm selects the largest value of  $t$  such that all receivers' utilities in  $T_b$  are incremented by the same amount while maintaining feasibility of the allocation. Steps 4 and 5 apply this increase to the "current" receiver rates and link rates respectively. Step 6 removes any receivers from  $T_{b+1}$  whose rates cannot be incremented any further, or else they would be larger than the maximum session rate, or would cause overutilization of some link. Step 7 removes any receivers in single-rate sessions from  $T_{b+1}$ , given that some other receiver in that session has been removed (so that all receiver rates in this session remain identical). A utility-free version of the algorithm appears in the appendix of [17].

Existence and uniqueness of the util-max-min-fair allocation is given by the following theorem, whose proof appear in [18].

*Theorem 1:* The algorithm presented in Figure 1 constructs the unique util-max-min fair allocation.

### A. Fairness Properties

Let us first examine some desirable properties of a unicast util-max-min fair allocation. It is easy to extend those properties that are known to hold for a (non-utility) unicast max-min fair allocation (see [2]) to the unicast util-max-min fair allocation [5].

**Unicast Fairness Property 1: (Unicast-Max-min-Fairness)** For each session  $S_i$ ,  $1 \leq i \leq m$ , either  $a_i = \alpha_i$ , or else there is at least one fully utilized link,  $l_j$ , where for all  $1 \leq i' \leq m, 0 < a_{i'} \leq \alpha_i$  whenever  $r_{i'} \in R_j$  (or equivalently,  $0 < \mu_{i'}^{-1}(f_{i',j}) \leq \mu_i^{-1}(f_{i,j})$ ).

**Unicast Fairness Property 2: (Unicast-Same-Path-Receiver-Fairness)** If two unicast sessions,  $S_i$  and  $S_{i'}$ , within a unicast network have identical routes, then either  $a_i = \alpha_i < a_{i'}$ , or  $a_{i'} = \alpha_{i'} < a_i$ , or  $a_i = a_{i'}$ .

Let us consider what makes these fairness properties desirable. To do this, we consider two *perspectives* of fairness of an allocation. From a receiver perspective, an allocation should be fair to receiver utilities: a receiver's utility should be as large as possible without "stealing" bandwidth from receivers with lower utilities. This is guaranteed by Unicast Fairness Property 1: there is no unused available bandwidth since some link on the receiver's route is fully utilized. Also, there is a fully utilized link over which the receiver's utility is as high as that of any other receiver's utility whose route crosses the link. Increasing this receiver's rate further would result in "stealing" bandwidth, and hence "stealing" utility, from these other receivers sharing the link. From the perspective of a session, a link's capacity should be used "fairly" by sessions. In other words, a session's allocation on a link (and hence all of its downstream receivers' utilities) should be as large as possible without "stealing" bandwidth (and hence utility) from other sessions that utilize the link.

For a unicast network, the receiver and session perspectives are identical because a session's route is identical to its receiver's route, and the share of bandwidth used on each link by the session equals the receiving rate of its receiver. This is not always true in a multicast network: a receiver's route is only part of the session's route, and, in a multi-rate session, when two receivers within the session receive at different rates, there are at least two links that have differing session link rates for that session. Hence, an allocation might be "fair" from the session




Fig. 2. A sample network

perspective without being “fair” from the receiver perspective, or vice versa. One possibility is to only consider fairness properties from a single perspective (e.g., [23] considers only the session perspective). However, in this section we will assume that it is desirable to satisfy fairness properties from both perspectives. We extend Unicast Properties 1 and 2, to multicast networks from both a session and receiver perspective.

Before presenting the desirable fairness properties for multicast networks, we introduce an example network that will be used to illustrate these different properties. Figure 2 presents a simple network with three sessions; sender  $x_1$  in session  $S_1$  sends to a single receiver,  $r_{1,1}$ . In session  $S_2$ , sender  $x_2$  sends to two receivers  $r_{2,1}$  and  $r_{2,2}$ . In session  $S_3$ , sender  $x_3$  sends to two receivers,  $r_{3,1}$  and  $r_{3,2}$ . The receiving rate of a receiver,  $a_{i,k}$ , is indicated to the immediate right of the receiver. Each link  $l_j$  has its capacity indicated next to the link labeling, separated by a colon (e.g.,  $l_1 : 5$  means that  $c_1 = 5$ ). Adjacent to the link labeling for each  $l_j$  are the session link rates, appearing in the form,  $(f_{1,j} : f_{2,j} : f_{3,j})$ . For simplicity, we assume that in the example, all sessions utilize the same utility function,  $\forall i, \mu_i(x) = x$ , such that each receiver’s utility is its receiving rate, i.e.,  $\forall i, k, a_{i,k} = d_{i,k}$ .

**Fairness Property 1: (Fully-Utilized-Receiver-Fairness)** A receiver’s utility  $a_{i,k}$  is fully-utilized-receiver-fair if either  $a_{i,k} = \alpha_i$ , or there is at least one fully utilized link,  $l_j$ , where  $r_{i,k} \in R_{i,j}$  and  $a_{i',k'} \leq a_{i,k}$  for all receivers  $a_{i',k'} \in R_j$ . A session’s allocation is defined to be fully-utilized-receiver-fair if the rate for each receiver in the session is fully-utilized-receiver-fair. An allocation of rates throughout the network is fully-utilized-receiver-fair if each session is fully-utilized-receiver-fair.

We define a receiver’s rate  $d_{i,k}$  to be fully-utilized-receiver-fair if its utility is fully-utilized-receiver-fair. This definition is introduced because at times in this section, it is more convenient to discuss fair allocations in terms of receiver rates than in terms of receiver utilities.

Fully-utilized-receiver-fairness is the multicast extension of Unicast Property 1’s prevention of “stealing bandwidth” from other receivers. For instance, in Figure 2, link  $l_3$  is fully utilized and lies on receiver  $r_{2,2}$ ’s route. Because  $r_{2,2}$  receives a utility that is no less than any other receiver whose route traverses  $l_3$ , its utility (and hence its rate) is fully-utilized-receiver-fair. Because all other receivers’ rates in  $S_2$  are fully-utilized-receiver-fair, session  $S_2$ ’s allocation of rates is fully-utilized-receiver-fair. Because  $S_1$ ’s and  $S_3$ ’s allocations are also fully-utilized-receiver-fair, the allocation (of rates for the entire network) is fully-utilized-receiver-fair.

**Fairness Property 2: (Same-Path-Receiver-Fairness)** A pair of receivers  $r_{i,k}$  and  $r_{i',k'}$  are same-path-receiver-fair if their routes traverse the same set of links ( $r_{i,k} \in R_j \iff r_{i',k'} \in R_j$ ), and either one receiver’s utility is constrained by its session’s maximum desired rate (i.e., either  $a_{i,k} = \alpha_i < a_{i',k'}$  or  $a_{i',k'} = \alpha_{i'} < a_{i,k}$ ), or else  $a_{i,k} = a_{i',k'}$ .

Same-path-receiver-fairness states that if two receivers’ routes traverse identical links, then the receivers should receive identical utilities (unless a receiver’s utility reaches its application’s maximum desired utility,  $\alpha_i$  or  $\alpha_{i'}$ ). In Figure 2, receivers  $r_{1,1}$  and  $r_{2,1}$  are a pair of receivers whose utilities (and hence their rates) are same-path-receiver-fair. The reader should note that for networks where all sessions’ utility functions are identical ( $\forall i, i' \mu_i = \mu_{i'}$ ), same-path-receiver-fairness is also a property of TCP-fairness [12], which states that a flow that traverses the same route as a TCP session should receive at a rate identical to the TCP session. If  $S_1$  is a unicast TCP session, then, in order for  $r_{2,1}$ ’s rate to be TCP-fair, same-path-receiver-fairness must hold for these two receivers.

**Fairness Property 3: (Per-Receiver-Link-Fairness)** A session  $S_i$ ’s allocation is per-receiver-link-fair if for each receiver  $r_{i,k} \in S_i$ , either 1)  $a_{i,k} = \alpha_i$ , or 2) there is a link  $l_j$  that is fully utilized ( $\exists j, r_{i,k} \in R_j, f_j = c_j$ ), and for other sessions  $S_{i'}, \mu_{i'}(f_{i',j}) \leq \mu_i(f_{i,j})$ . An allocation of rates throughout the network is per-receiver-link-fair if each session’s allocation is per-receiver-link-fair.

**Fairness Property 4: (Per-Session-Link-Fairness)** An allocation is per-session-link-fair for a session  $S_i$  if  $a_{i,k} = \alpha_i$  for each receiver in  $S_i$  or there exists a fully utilized link  $l_j$  in  $S_i$ ’s route where for other sessions  $S_{i'}, \mu_{i'}(f_{i',j}) \leq \mu_i(f_{i,j})$ . An allocation of rates throughout the network is per-session-link-fair if each session’s allocation is per-session-link-fair.

Per-receiver-link-fairness requires that session  $S_i$  gets a “fair share” of link rate (with respect to the session’s utility function) along every path from sender  $x_i$  to its receivers. Per-session-link-fairness is a weaker version of this: a session must get a “fair share” of link rate (with respect to the session’s utility function) on at least one link in its route (i.e., along the route of at least one receiver). In Figure 2, session  $S_2$  is per-session-link-fair: on the route to receiver  $r_{2,2}$ , link  $l_3$  is fully utilized and session  $S_2$ ’s link rate on  $l_3$  is no less than the link rates of other sessions on  $l_3$ . It is also per-receiver-link-fair, because similar conditions hold on the route of its other receiver,  $r_{2,1}$ . Sessions  $S_1$  and  $S_3$  are also both per-receiver-link-fair and per-session-link-fair, making the network allocation both per-receiver-link-fair and per-session-link-fair.

## B. Multi-rate Session Impact on Fairness Properties

It is fairly easy to see that in a unicast network, Fairness Property 2 and Unicast Property 2 are identical, and the remaining multicast fairness properties are identical to Unicast Property 1. We now proceed to establish properties of util-max-min fair allocations in terms of the types of sessions (multi-rate or single-rate) within the network. All proofs appear in [18].

**Theorem 2:** A multi-rate util-max-min fair allocation satisfies Fairness Properties 1, 2, 3, and 4 also hold.

Theorem 2 tells us that if all sessions are multi-rate, then the util-max-min fair allocation satisfies all of our desired fairness

properties. We now introduce a mathematical ordering among allocations that allows us to comparatively examine the “util-max-min fairness” of an allocation within a network:

*Definition 2:* We say a vector  $(x_1, x_2, \dots, x_k)$  is *ordered* if for all  $i, 1 \leq i < k, x_i \leq x_{i+1}$ . Let  $X = (x_1, x_2, \dots, x_k)$  and  $Y = (y_1, y_2, \dots, y_k)$  be ordered vectors. We write  $X \leq_m Y$  ( $X$  is min-unfavorable to  $Y$ ) if no  $i$  exists such that  $x_i > y_i$ , or for any  $i$  where  $x_i > y_i$ , there is some  $j < i$  where  $x_j < y_j$ . We write  $X <_m Y$  to indicate  $(X \leq_m Y) \wedge (X \neq Y)$ .

Note that under the above definition,  $\leq_m$  is reflexive ( $X \leq_m X$ ), non-symmetric ( $X = Y \iff X \leq_m Y \wedge Y \leq_m X$ ), and transitive ( $W \leq_m X \wedge X \leq_m Y \Rightarrow W \leq_m Y$ ). Furthermore, for any pair,  $X$  and  $Y$ , of ordered vectors of identical length, either  $X \leq_m Y$  holds, or  $Y \leq_m X$  holds, or both. Min-unfavorability is a lexicographic ordering that is similar to alphabetizing two text strings of the same length. Let  $x_i$  represent the  $i$ th character of the first string, and  $y_i$  represent the  $i$ th character of the second string. Then  $X \leq_m Y$  if and only if  $X = Y$  or an alphabetization places  $X$  before  $Y$ . A version of this ordering has been applied specifically within unicast networks [5]. We now state the result regarding the uniqueness of a util-max-min-fair allocation. Unless stated otherwise, an allocation refers to an ordered vector of receiver utilities.

*Theorem 3 (Max-min fair uniqueness)* There is a unique util-max-min fair allocation for any network. Furthermore, if  $A$  is the util-max-min fair allocation (of receiver utilities), and  $B$  is some other feasible allocation (of receiver utilities), then  $A <_m B$ .

Note that Theorem 3 holds for a network containing any mix of single-rate and multi-rate sessions, as well as an arbitrary set of session utility functions,  $\{\mu_i\}$ . Theorem 3 along with the definition of min-unfavorability can be combined to show that the util-max-min fair allocation maximizes the minimum utilities allocated to sessions in a network: since all allocations are min-unfavorable to the util-max-min fair allocation, there exists a threshold utility  $x'$  such that for any receiver utility  $z < x'$ , the number of receivers that achieve a utility at or below  $z$  is minimal (smaller or equal) within the util-max-min fair allocation. Furthermore, the number of receivers that achieve a utility at or below  $x'$  is minimized (strictly smaller) within the util-max-min fair allocation. This result can be stated more formally as a general property of min-unfavorability:

*Lemma 1:*  $X <_m Y \iff \exists x'$  such that  $\forall z < x', |\{x_i \in X : x_i \leq z\}| \geq |\{y_i \in Y : y_i \leq z\}|$  and  $|\{x_i \in X : x_i \leq x'\}| > |\{y_i \in Y : y_i \leq x'\}|$ .

Because the min-unfavorable relation is transitive, it gives a strict ordering among the feasible allocations for a network, where the util-max-min fair allocation is the maximum under the ordering. Thus, one can quantitatively compare the util-max-min fairness of two allocations  $A$  and  $B$ .

### C. Fairness limitations of single-rate sessions

Theorem 2 states that a multi-rate util-max-min fair allocation satisfies our four desirable fairness properties. Let us now see where a single-rate util-max-min fair allocation fails to do so. The fact that a single-rate util-max-min allocation is per-session-link-fair is a simple extension of the results in [23] which demonstrate that the (non-utility-based) single-rate




Fig. 3. An example where a single-rate session would fail all but one of the fairness properties.

max-min fair allocation is per-session-link-fair. However, the single-rate util-max-min fair allocation does not always satisfy the other fairness properties. Consider the simple example in Figure 3, whose labeling is performed in an identical manner to that of Figure 2. For simplicity, we again assume that each receiver’s utility equals its rate, i.e.,  $\forall i, \mu_i(x) = x$ . Here, we have a network with two sessions,  $S_1$  and  $S_2$ , whose respective senders,  $x_1$  and  $x_2$ , are located at the same point in the network. We assume that the maximum desired utilities are large,  $\alpha_1 = \alpha_2 = \infty$ , such that they do not bound receiving rates in this network. Session  $S_1$  is a single-rate session containing three receivers  $r_{1,1}, r_{1,2}, r_{1,3}$ , session  $S_2$  is a unicast session whose receiver  $r_{2,1}$  is located at the same point in the network as receiver  $r_{1,1}$ . In the util-max-min fair allocation, receivers in session  $S_1$  receive at a rate of 2 (since this fully utilizes link  $l_2$  and all receivers must receive at the same rate in a single-rate session), the receiver in  $S_2$  receives at a rate of 3. Receivers  $r_{1,1}$  and  $r_{2,1}$  fail to achieve same-path-receiver-fairness, since they have the same routes, but differing receiving rates - and hence differing utilities. Receiver  $r_{1,3}$ ’s rate does not satisfy fully-utilized-receiver-fairness, because there is no fully utilized link along its route on which its utility is the largest compared to other receivers whose routes cross the same link. It follows that fully-utilized-receiver-fairness does not hold for session  $S_1$ , nor does it hold for the network. Last, per-receiver-link-fairness fails to hold for session  $S_1$  (hence for the network as well) on the route to receiver  $r_{1,3}$ , since no link on this route is fully utilized. Per-receiver-link-fairness also fails to hold on the route to receiver  $r_{1,1}$ . This is because link  $l_1$  is the only fully utilized link on  $r_{1,1}$ ’s route, and the link rate (and hence utility) of session  $S_1$  on  $l_1$  is smaller than that of session  $S_2$ . This example demonstrates that three out of the four desirable properties can fail to hold for single-rate util-max-min fair allocations.

### D. Combining Multi-rate and Single-rate Sessions

We have examined the extent to which our four desirable properties hold for networks in which all sessions are the same type. Let us now consider these properties in the context of a network that contains a combination of multi-rate and single-rate sessions. Single-rate sessions are likely to always exist due to application constraints, such as a requirement that all receivers must complete receipt of data at approximately the same time.

*Theorem 4:* Consider a network  $N = \{G, \{S_1, \dots, S_m\}, \tau, \sigma\}$  in which session types can differ, i.e., there can exist a pair of

sessions,  $S_i, S_{i'} \in N$  such that  $\sigma(S_i) \neq \sigma(S_{i'})$ . Then, the following are properties of the util-max-min fair allocation of  $N$ :

(a) Fully-utilized-receiver-fairness holds for each receiver  $r_{i,k} \in S_i$  where  $\sigma(S_i) = \mathcal{M}$ .

(b) per-receiver-link-fairness holds for each session  $S_i$  where  $\sigma(S_i) = \mathcal{M}$ .

(c) Per-session-link-fairness holds for all sessions  $S_i$ .

(d) Same-path-receiver-fairness holds between any two receivers  $r_{i,k}$  and  $r_{i',k'}$  where  $\sigma(S_i) = \sigma(S_{i'}) = \mathcal{M}$ .

(e) If  $\sigma(S_i) = \mathcal{M}$  and  $\sigma(S_{i'}) = \mathcal{S}$ , and  $r_{i,k} \in S_i$  and  $r_{i',k'} \in S_{i'}$  have identical routes, then either  $a_{i,k} = \alpha_i$  or  $a_{i,k} \geq a_{i',k'}$ .

Theorem 4 states that, even with the presence of single-rate sessions, multirate sessions within the util-max-min fair allocation exhibit all four desirable fairness properties.

*Lemma 2:* Let  $N = (G, \{S_1, \dots, S_m\}, \tau, \sigma)$  and  $\bar{N} = (G, \{S_1, \dots, S_m\}, \tau, \bar{\sigma})$  be networks where the set of multi-rate sessions in  $\bar{N}$  is a subset of the set of multi-rate sessions in  $N$ , (i.e.,  $\forall i, \bar{\sigma}(S_i) = \mathcal{M} \Rightarrow \sigma(S_i) = \mathcal{M}$ ). If  $A$  is the ordered vector of util-max-min fair receiver utilities in  $N$ , and  $\bar{A}$  is the ordered vector of util-max-min fair receiver utilities in  $\bar{N}$ , then  $\bar{A} \leq_m A$ .

*Corollary 1:* Let  $N = (G, \{S_1, \dots, S_m\}, \tau, \sigma)$  be a multi-rate network ( $\forall i, \sigma(S_i) = \mathcal{M}$ ), and let  $\bar{N} = (G, \{S_1, \dots, S_m\}, \tau, \bar{\sigma})$  be identical to  $N$ , except that  $\bar{\sigma}(S_i) = \mathcal{S}$  for some sessions. Let  $A$  be the ordered vector of receiver utilities for a multi-rate util-max-min fair allocation within  $N$ , and let  $B$  be the ordered vector of receiver utilities in  $\bar{N}$ . Then  $B \leq_m A$ .

### E. Impact of Session Type on Receiver Rates

Now, let us consider how varying session types affects receiving utilities on a session-by-session basis. We can prove that if all sessions' types are fixed except for session  $S_i$ , then if  $S_i$  is multi-rate, all of its receivers will achieve utilities that are no less than what they would achieve if  $S_i$  is single-rate (see Lemma 9 in [18]). Unfortunately, this result does not extend to the case when several sessions can switch types. In fact, it is rather difficult to say what happens to receiver utilities due to changes in the session type or the network topology. For example, one might conjecture that removing a receiver  $r_{i,k}$  from a session would only increase other receivers' fair utilities. Our intuition was that this would be the case since the removal frees up bandwidth that can then be used by other receivers whose route crosses  $r_{i,k}$ 's route. However, the util-max-min fair allocation of bandwidth after the receiver is removed can cause receiver rates and thus receiver utilities (both in session  $S_i$  and in other sessions) to vary in either direction.

To see this, consider the examples in Figure 4. Again, we consider sessions where each receiver's rate equals its utility (i.e.,  $\mu_i(x) = x$ ). Both networks contain three multi-rate sessions,  $S_1, S_2$ , and  $S_3$ .  $S_1$  and  $S_2$  each contain a single receiver,  $S_3$  contains two receivers, the second ( $r_{3,2}$ ) is subsequently removed. The util-max-min fair utilities for receivers are indicated before and after this removal. Note that in Figure 4(a),  $r_{3,1}$ 's util-max-min fair utility decreases and  $r_{1,1}$ 's utility increases as a result of the removal. In Figure 4(b),  $r_{3,1}$ 's utility increases and  $r_{1,1}$ 's utility decreases. This demonstrates that removing receivers from sessions can have a non-obvious impact on the

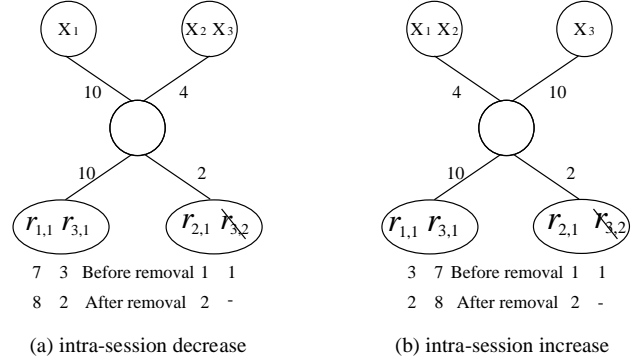


Fig. 4. The change in util-max-min fair rates due to a removal of a receiver from a session.




Fig. 5. A session whose receivers have differing utility functions that does not exhibit same-path-receiver-fairness or fully-utilized-receiver-fairness.

util-max-min fair utilities of the remaining receivers in the network.

### F. Varying Utility Functions Within a Session

We briefly consider what happens if we allow receivers in the same session to utilize different utility functions. Define  $\mu_{i,k}$  to be the utility function for receiver  $r_{i,k}$  such that  $a_{i,k} = \mu_{i,k}(d_{i,k})$ . The proof of existence and uniqueness of the util-max-min-fair requires small modifications to the algorithm in Figure 1 and proofs in [18] of Lemma 4 and Theorem 3. Definition 1 remains well-defined in this context: the definition requires a comparison of receiver utilities, and does not require that receivers in the same session utilize the same utility function. However, Fairness properties per-receiver-link-fairness and per-session-link-fairness are not well-defined in this context: these definitions require the existence of the session's unique utility function,  $\mu_i$ .

We now demonstrate that fairness properties same-path-receiver-fairness and fully-utilized-receiver-fairness need not hold when different receivers within a session can employ different utility functions. Consider the multi-rate util-max-min fair allocation shown in Figure 5. Session  $S_1$  is a multi-rate multicast session in which  $\mu_{1,1}(x) = .5x$ , and  $\mu_{1,2}(x) = x$ . Session  $S_2$  is a unicast sessions in which  $\mu_{2,1}(x) = x$ . Links  $l_1, l_2$ , and  $l_3$  have respective capacities of 5, 6, and 4. The receiving rate and utility within the util-max-min fair allocation is given as  $d_{i,k}(a_{i,k})$  immediately to the right of each receiver  $r_{i,k}$ . Utilization of link bandwidth by the two sessions is indicated adjacent to each link  $l_j$  in the form  $(f_{1,j}, f_{2,j})$ . In this example, even though  $r_{1,1}$  and  $r_{2,1}$  have identical data paths, it

is the case that  $a_{1,1} \neq a_{2,1}$ . Hence, the allocation is not same-path-receiver-fair. Since the routes to these two receivers share all network links, there is no link on which  $a_{2,1} \geq a_{1,1}$ , such that the allocation is not fully-utilized-receiver-fair.

### G. Section Summary

We now summarize the main results of this section. We have shown that if multicast sessions are multi-rate, then the util-max-min fair allocation satisfies additional desirable fairness properties that do not necessarily hold in a single-rate util-max-min fair allocation. We also examined networks in which some of the sessions are single-rate, while the remaining are multi-rate. By examining fairness properties on a per-session basis, we find that all of the fairness properties hold in general only in multi-rate sessions. Next, we used the min-unfavorable relation to comparatively examine any two max-min fair allocations for a network that contains a mix of single-rate and multi-rate sessions. We find that max-min fair allocation of a set of sessions is min-unfavorable to the same set of sessions where a subset of single-rate sessions is “replaced” by multi-rate sessions. Finally, we showed that if receivers in the same session have differing utility functions, then two of the desirable fairness properties need not hold, and the other two are not well-defined.

## III. ACHIEVING MULTI-RATE MAX-MIN FAIRNESS WITH LAYERING

In the previous section, we motivated the use of multi-rate sessions by showing that in theory they yield more desirable util-max-min fair allocations. One way to then obtain these rates in practice is to have the sender configure layers so that each receiver can obtain its fair rate by joining some subset of layers. However, the number of layers may need to be as large as the number of receivers in the session, making such an approach infeasible for large multicast sessions. Furthermore, the number of layers and the rate per layer is often beyond the control of the session itself, due to application-specific requirements, a limitation in the availability of multicast groups, or because it is too difficult for the sender to obtain the feedback needed to appropriately configure the rates of each of the layers. In this section, we examine how receivers can obtain their long term average max-min fair rates by repeated joins and leaves from multicast groups on which data is sent at a restricted set of rates. We will see that such a mechanism will force us to reconsider our previous assumption of how receiver rates impact link rates in the network. For simplicity, we assume in the remainder of the paper that each receiver’s utility matches its receiving rate, i.e.,  $\forall i, \mu_i(x) = x$ .

Let us first discuss the implementation of a layered multicast approach. Data to be transferred is split into  $M$  layers by the sender, where layers are transmitted on separate multicast groups, each at some rate. The layers are ordered  $L_1, \dots, L_M$ , such that all receivers desiring transmission join the group containing layer  $L_1$ , and any receiver that joins the group containing layer  $L_j$  must also join or already be joined to layer  $L_i$  for all  $1 \leq i < j$  (henceforth, this is implied when we say that the receiver joins the layer or joins *up to* the layer). A receiver joined up to layer  $L_i$  receives data from the sender at an aggregate rate equal to the sum of the rates of layers  $L_1$  through  $L_i$ .

Joining layers increases the aggregate rate, while leaving layers decreases the aggregate rate.<sup>5</sup>

Let us examine why receivers must join and leave layers to obtain max-min fair rates. An obvious alternative is to require receivers to choose rates that can be obtained by joining up to a given layer and remaining at that rate for the duration of the session. This makes a finite set of rates available to the receiver. However, if these layers cannot be configured to the needs of receivers for reasons described above, the max-min fair allocation might not even exist! As an example, consider a simple network that consists of a single link with capacity  $c$ , and let there be two layered multicast sessions,  $S_1$ , and  $S_2$  that traverse this link. Each session contains a single receiver, respectively denoted  $r_1$  and  $r_2$ . The sender for session  $S_1$  provides three layers, and sends at a rate of  $c/3$  per layer. The sender for session  $S_2$  provides two layers, and sends at rate  $c/2$  per layer. The set of feasible allocations is  $\{(0, 0), (0, c/2), (0, c), (c/3, 0), (c/3, c/2), (2c/3, 0), (c, 0)\}$ , where  $(d_1, d_2)$  implies receiver  $r_i$  receives at a rate of  $d_i$ . None of these allocations are max-min fair. For instance,  $(d_1, d_2) = (c/3, c/2)$  is not max-min fair since  $(\bar{d}_1, \bar{d}_2) = (2c/3, 0)$  is feasible, and  $d_1 < \bar{d}_1$ , but  $d_2 > d_1$ , hence there is no  $j$  where  $\bar{d}_j < d_j \leq d_1$ <sup>6</sup> (contradicting the defined requirement for max-min fairness). The reader can easily verify that none of the other feasible allocations is max-min fair.

Although it is not possible to achieve a max-min fair rate allocation when receivers are restricted to joining some arbitrarily chosen fixed set of layers for the entire length of a session, it is possible to achieve long-term average max-min fair rates through joins and leaves. The idea of using long term average rates also appears in current definitions of TCP-fairness [3], [12], [15], [24]. We define the *quantum*,  $\Delta t$ , to be the minimum amount of time over which a receiver’s average rate is computed. We say that a rate of  $r$  is obtained through a link during the  $i$ th quantum if  $r\Delta t$  bytes pass through the link between times  $i\Delta t$  and  $(i+1)\Delta t$ . We say that a link  $l_j$  can support a capacity of  $c_j$  if it is able to forward  $c_j\Delta t$  bytes within each time quantum. We note that rapid changes in transmission rate are undesirable for certain applications like the streaming of live multimedia. For such applications, it would be most useful to use a large quantum such that the period of rate adjustment would be less frequent. Selecting an appropriate quantum to meet the needs of specific applications is beyond the scope of this paper.

Let us now consider an idealized network where a receiver can use joins and leaves to obtain its fair rate. The network is ideal in that we assume that network propagation delays and leave latencies are negligible compared to  $\Delta t$  and to packet inter-arrival times for each session. In this model, a packet traverses a link  $l_j$  only if it is received by some receiver  $r_{i,k} \in R_j$ . We also assume that all packets are of equal size, and for any receiver  $r_{i,k}$ , let  $d_{i,k} \leq m_i$  be its fair packet rate (in packets/sec) within the network. Consider a single layer (multicast group), where the transmission rate on the layer,  $\rho$ , satis-

<sup>5</sup>We make the assumption that there is some utility in receiving at a faster rate, e.g., audio and video transmissions increase in clarity, reliable data transmissions take less time.

<sup>6</sup>Or less formally,  $r_1$ ’s increase in rate does not result in a decrease in any receiver’s rate whose original rate was less than  $r_1$ ’s.



fies  $\rho \geq \max\{d_{i,k} : r_{i,k} \in S_i\}$ . Receiver  $r_{i,k}$  joins the single layer so that it receives the first  $d_{i,k}\Delta t$  packets within the quantum,<sup>7</sup> then leaves the group. This is clearly possible, since  $d_{i,k} \leq m_i \leq \rho$ , and  $\rho\Delta t$  packets are transmitted on the layer during the quantum.

An alternative to achieving max-min fair allocations in a with timed joins and leaves in a discrete-rate setting is to instead achieve what is called a *maximally fair* allocation [20]. Maximally fair allocations are equivalent to max-min fair allocations when receiver rates are not restricted to fixed, discrete values. However, when such a restriction holds, the rates need not be unique, and we do not consider this alternative definition of fairness here.

In the join-leave scenario discussed above, for any link  $l_j$  and session  $S_i$  where  $|R_{i,j}| > 0$ , there is some receiver  $r_{i,k'}$  that receives  $d_{i,k'} = \max\{d_{i,k} | r_{i,k} \in R_{i,j}\}$  packets per time quantum. Hence, this is the minimum number of packets that traverse link  $l_j$  for session  $S_i$  per quantum. Transmitting exactly this number of packets requires that all other receivers  $r_{i,k} \in R_{i,j}$  receive a subset of the packets that are received by  $r_{i,k'}$  per quantum. When this is not the case,  $f_{i,j} > d_{i,k'}$ .

**Definition 3:** We define the **link-overuse** of a link  $l_j$  for a session  $S_i$  to be  $f_{i,j} / \max\{d_{i,k} | r_{i,k} \in R_{i,j}\}$ , where  $f_{i,j}$  is the long-term average link rate  $l_j$  by session  $S_i$ , and  $d_{i,k}$  is the long-term average rate for receiver  $r_{i,k}$ . We say a session's bandwidth utilization of a link is **efficient** for session  $S_i$  if the link's link-overuse for that session is 1.0, and define a session  $S_i$ 's **efficient link rate** to be  $\max\{d_{i,k} | r_{i,k} \in R_{i,j}\}$ .

To understand the impact on link-overuse of coordination between receiver joins and leaves, let us examine what happens on a shared link when there is no implicit join/leave coordination. Our demonstration is performed using a simple protocol that is unlikely to be used in practice, but is used here for ease of analysis. In this protocol, we assume that each receiver  $r_{i,k}$  within session  $S_i$  randomly chooses the  $d_{i,k}\Delta t$  packets it should receive within the quantum, with each packet having an equally likely chance of being chosen as any other in that quantum. In this case, the expected utilization of link  $j$  by  $S_i$ ,  $E[U_{i,j}]$ , is  $\rho(1 - \prod_{t=1}^s (1 - d_{i,k_t}/\rho))$ , where  $\{d_{i,k_1}, \dots, d_{i,k_s}\}$  are the rates of receivers that are members of the set  $R_{i,j}$  (derivation in Appendix A).

Figure 6 shows how the number of receivers within a session that utilize a link (i.e.,  $|R_{i,j}|$ ) impacts the link-overuse of a layer in this scenario. The number of receivers is shown on the  $x$ -axis, while the session's link-overuse is indicated on the  $y$ -axis. The curves represent various configurations of  $\{d_{i,1}, \dots, d_{i,s}\}$ . For curves labeled **All**  $z$ , ( $z = 0.1, 0.5$ , or  $0.9$ ),  $d_{i,k}$  is set respectively to  $\rho/10, \rho/2$ , and  $9\rho/10$  for all receivers. For curves labeled **1st  $w$  rest  $z$** ,  $d_{i,1} = w\rho$ , and  $d_{i,s} = z\rho$  for  $1 < s \leq |R_{i,j}|$ . Note that in each plot, the efficient link rate remains constant as the number of receivers is varied.

We find that for link-overuse to be high, the ratio of the efficient link rate to the transmission rate (i.e.,  $\max_{r_{i,k} \in R_{i,j}} \{d_{i,k}\} / \rho$ ) must be small. In fact, the link-overuse can only be as large as the multiplicative inverse of this value

<sup>7</sup>If  $d_{i,k}\Delta t$  is not an integer, then it can elect to receive  $\lfloor d_{i,k}\Delta t \rfloor$  packets in each quantum, and periodically receive  $\lceil d_{i,k}\Delta t \rceil$  to come arbitrarily close to  $d_{i,k}\Delta t$ .




Fig. 6. Link overuse of a single layer with random joins




Fig. 7. An example where a network fails to achieve session-perspective fairness properties due to link-overuse.

(e.g.,  $\max\{d_{i,k} | r_{i,k} \in R_{i,j}\} / \rho = .1$  bounds link-overuse from above by 10), and asymptotically reaches this value with an increase in the number of receivers that share the link. In other words, for link-overuse to be high, all receivers must require only a small percentage of packets per quantum from a layer.

A second result is that for a fixed efficient link rate, link-overuse increases most rapidly as a function of the number of receivers when all receivers receive at the same rate. In other words, an upper bound on how additional receivers impact link-overuse is obtained by considering a network in which all receivers within a session have identical fair rates.

These results gives a preliminary indication as to what impacts the magnitude of link-overuse within a network. We find that having additional layers often leads to a reduction in link-overuse that is sometimes substantial, and that it never increases link-overuse beyond that exhibited for the single-layer case. Details of these results can be found in Appendix E of [17].

Note that our assumption in Section II that  $f_{i,j} = \max\{d_{i,k} : r_{i,k} \in R_{i,j}\}$  amounts to an assumption that multi-rate sessions are efficient (i.e., on all links in the network, a multi-rate session's link rate equals its efficient link rate). When there are multi-rate sessions that are not efficient, a multi-rate max-min fair allocation might not satisfy per-session-link-fairness (and hence might not satisfy per-receiver-link-fairness). To show this, we consider the network shown in Figure 7, whose labeling is similar to that of Figures 2 and 3. We again assume that the maximum desired rates are large so as not to bound receiving rates, e.g., let  $m_1 = m_2 = \infty$ . Here, session  $S_1$  is multi-rate with a link-overuse of 2 over the shared link,  $l_4$ . This could




Fig. 8. The impact of link-overuse on fair rates.

occur, for instance, if receivers  $r_{1,3}$  and  $r_{1,2}$  would receive at rate 3 on each even-numbered quantum and at rate 1 on each odd-numbered quantum while receiver  $r_{1,1}$  receives at rate 3 on each odd-numbered quantum at rate 1 on each even-numbered quantum. Since the maximum receiving rate for receivers in  $S_1$  (all of whose data-paths traverse  $l_4$ ) is 2,  $f_{1,4} = 4$ . Since this is the only link that is fully utilized, and  $f_{1,4} > f_{2,4}$ , per-session-link-fairness fails to hold for session  $S_2$ . It follows that per-receiver-link-fairness fails to hold for session  $S_2$  as well.

It is trivial to show that Fairness Properties 2 and 1 hold even when sessions are not efficient.

#### A. The impact of link-overuse on fair rates

Let us now examine the impact that link-overuse has on fairness within a network. We now demonstrate why sessions with lower link-overuse are preferable to ones with high link-overuse. We begin by relaxing our assumption made in Section II that  $f_{i,j} = \max\{d_{i,k} : r_{i,k} \in R_{i,j}\}$ . We extend our definition of a session to be a tuple  $S_i = (X_i, \{r_{i,1}, \dots, r_{i,k_i}\}, v_i)$  that now includes a link-overuse function  $v_i$ . Here,  $v_i$  maps a set (of arbitrary size) of receiver rates to a link rate. Given an allocation of receiver rates,  $A$ , session  $S_i$ 's link rate for link  $l_j$  is computed as  $f_{i,j} = v_i(\{d_{i,k} : r_{i,k} \in R_{i,j}\})$ . In Section II,  $v_i$  is simply the max operation. Since  $f_{i,j} \geq d_{i,k}$  must hold whenever  $r_{i,k} \in R_j$  (for reasons discussed in Section II), it is necessary that  $v_i(\{d_{i,k} : r_{i,k} \in R_{i,j}\}) \geq \max\{d_{i,k} : r_{i,k} \in R_{i,j}\}$ .

**Lemma 3:** Let  $N = (G, \{S_1, \dots, S_m\}, \tau, \sigma)$ , and  $\bar{N} = (G, \{\bar{S}_1, \dots, \bar{S}_m\}, \tau, \sigma)$  be identical networks, where each session  $S_i$  in  $N$  is identical to  $\bar{S}_i$  in  $\bar{N}$ , except for their respective link-overuse functions,  $v_i$  and  $\bar{v}_i$ . Assume sessions in  $\bar{N}$  exhibit higher link-overuse than those in  $N$ , (i.e., for each session  $S_i$  and any set of real numbers,  $X$ ,  $v_i(X) \leq \bar{v}_i(X)$ ). Let  $A$  be the max-min fair allocation in  $N$  and  $\bar{A}$  the max-min fair allocation in  $\bar{N}$ . Then  $\bar{A} \leq_m A$ .

Lemma 3 states the following: assume that sessions are “replaced” by sessions that are identical, except that the session link rates required to support a given set of receiver rates are higher (e.g., the amount of coordination of joins and leaves between receivers within a session is reduced). It follows that the resulting max-min fair allocation is “less max-min fair” than the max-min fair allocation for the network with the sessions prior to the “replacement”.

We know that a link-overuse greater than 1.0 produces max-min fair rate allocations within the network that might not exhibit the session-perspective fairness properties, per-receiver-link-fairness and per-session-link-fairness. Also, using the min-unfavorable relation, we have shown that increased link-overuse might reduce the “max-min fairness” of a max-min fair allocation. Let us now quantitatively examine how link-overuse may impact fair rates. Consider a set of  $n$  sessions whose receiver rates are constrained by the same link,  $l$  with capacity  $c$ . Let  $m$  of these sessions be multi-rate and have a link-overuse  $v$  on link  $l$ , and the remaining  $n - m$  sessions be of an arbitrary session type and have link-overuse 1. Since we assume that all receivers' rates are constrained by link  $l$ , their max-min fair rates are all equal to  $\frac{c}{(n-m)+mv}$ . Figure 8 shows the receivers' rates as a function of the link-overuse,  $v$ . The  $x$ -axis indicates  $v$ , the various curves represent various values of the ratio of sessions,  $m/n$ , that exhibit link-overuse  $v$ . The  $y$ -axis presents the fair rate normalized by  $c/n$ , the fair rate for all the receivers in the network when all sessions are efficient.

Figure 8 indicates that even modest levels of link-overuse can substantially reduce the fair receiver rates for all sessions in the network. From this we draw two conclusions: first, it is important to maintain low link-overuse on network links to keep fair rates high. Second, when multi-rate sessions make up a small percentage of the sessions in the network, they have less of an impact on the fair rates of sessions. Due to the current proliferation of unicast traffic within the network, we expect that less than 5% of the sessions within the network will be multi-rate. This means that low levels of link-overuse greater than 1.0 can be tolerated.

These results raise an interesting dilemma: should multi-rate protocols be used to achieve fairness from the receivers' perspectives, even if it means failing to achieve per-session-link-fairness (a fairness property that holds when all network sessions are single-rate and unicast)? We argue that yes, multi-rate protocols should still be used, because the “unfair” additional usage of link bandwidth due to link-overuse can be justified in that the session is transmitting data to multiple receivers. A similar argument is used in [9] to allocate link bandwidth to sessions in a manner that is proportional to the number of receivers within the session.

The reduction in rate due to link-overuse can occur whenever a multi-rate session tries to achieve some form of fairness using joins and leaves of layers. For example, in [24], receiver join experiments are coordinated within a network where TCP-fairness is the fairness criteria. The coordination prevents “bottleneck bandwidth allocated to [the] protocol instance [from] not being fully exploited.” This lack of “exploitation” is, in effect, an artifact of link-overuse.

#### IV. OVERUSE IN PRACTICAL CONGESTION CONTROL PROTOCOLS

In Section III, we showed that a lack of join and leave coordination within a session increases the session's link-overuse on links shared by that session's receivers. This in turn is likely to reduce their fair receiving rates. Our final contribution is to show that link-overuse can easily be kept quite low in practice. We show this by measuring the link-overuse of several Internet

layered congestion control protocols that vary in the degree to which joins are coordinated among receivers. In these protocols, receivers react to congestion by leaving layers, and probe for available bandwidth by joining layers. We compute each protocol's link-overuse using analysis and simulation of simple network models. Because of the simplicity of the models, there may be some differences between what we observe and what will actually occur in practice. However, we do not expect these differences to affect our conclusions.

In each protocol, a receiver leaves the highest layer joined (unless only joined to one layer) whenever it observes a *congestion event*: an indication that some part of its data-path is being overutilized. In practice, a congestion event may be the loss of a packet by the receiver, or a bit set within a packet by the network used to indicate that the receiving rate should be lowered [14]. If no congestion events are observed by a receiver within a sequence of packet arrivals, it joins an additional layer (unless already joined to all layers). Using these protocols, a receiver repeatedly adjusts the set of layers to which it is joined for the duration of the session. The protocols differ in the degree to which joins are coordinated within a session.

- In the *Uncoordinated* protocol, there is no inherent coordination: upon receiving a packet, a receiver randomly decides whether to join an additional layer.

- In the *Deterministic* protocol, there is also no inherent coordination; a receiver joins an additional layer after receiving a fixed number of packets without loss since its last join or leave event.

- In the *Coordinated* protocol, the sender indicates (e.g., through a field within its transmitted packet) when receivers should join an additional layer. This is done in such a way so that when the field indicates that receivers joined up to layer  $i$  should join layer  $i + 1$ , it also indicates that receivers joined up to layer  $j < i$  should join layer  $j + 1$ .

The additional details of the protocols (layer rates, join-period) are based on the choices made in [24]. For instance, we require that the aggregate rate of layers 1 through  $i$  equals  $2^{i-1}$ , and that the expected number of packets received by a receiver between a previous join/leave event to its join to layer  $i + 1$  equals  $2^{2(i-1)}$ .<sup>8</sup> Because of these protocols' similarities to the protocol in [24], we anticipate these protocols are suitable for the same set of continuous stream and reliable bulk data transfer applications described in [24]. Due to a lack of round-trip-time dependence, these protocols come closer to achieving max-min fair rates than TCP-fair rates. A more precise description of these protocols and how they differ from the protocol in [24] is found in [17].

We model packet loss (or equivalently, congestion marking of packets) as a Bernoulli loss process. The reader can consider the loss process to be fairly accurate for a network where the number of flows across links is large, so that there is little correlation between the rate of an individual flow and the link loss rate [26]. Our model also assumes that receivers' reactions to coordinated events (shared loss, coordinated joins) take effect at the same time: two receivers that see identical loss patterns would be joined to the same set of layers. Under these con-

<sup>8</sup>In [24], the number of packets received equals  $2^{2(i-1)}$  (i.e., it is a deterministic value).





Fig. 9. Network models for coordination experiments

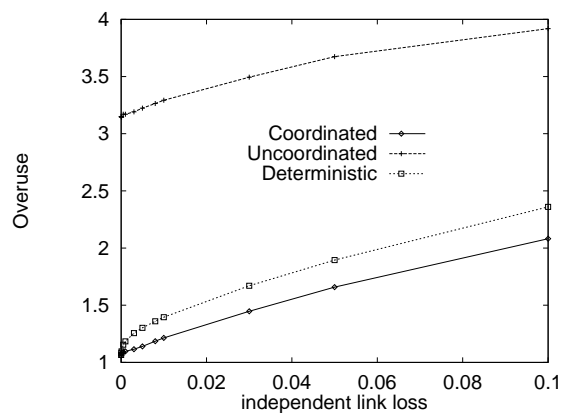
ditions, it can be argued that these protocols come “close” to achieving the max-min fair rates, i.e., the expected rate does not exactly equal the max-min fair rate, but the difference is fairly small.

Our experiments use modified star networks, as shown Figure 9, to examine how shared loss (i.e., loss on the shared link abutting the sender) and independent loss (i.e., loss on the fanout links) impact link-overuse. The initial set of experiments uses the topology in Figure 9(a). Using Markov models of the protocols over this network, we examine how different values of shared and independent loss impact the link-overuse of a session on the shared link. The details of these models appear in [17]. We summarize the most important finding: link-overuse is highest when receivers experience the same end-to-end loss rates. This result follows intuitively from our observation in Section III that link-overuse is highest when all receivers' receiving rates are equal.

Our Markov models are too computation-intensive to allow us to examine sessions with large sets of receivers. Instead, we turn to simulation. Figure 10 shows simulations of the protocols using 8 layers with 100 receivers in the session that have identical end-to-end loss rates, configured in the modified-star topology of Figure 9(b). In Figure 10(a), the shared loss rate is fixed to 0.0001 (i.e., very low shared loss), and the loss rate on each of the fanout links is given on the  $x$ -axis. Each curve shows the link-overuse for the three protocols we consider. Each point plotted is the mean of 30 experiments where the sender transmits 100,000 packets, the variance is less than 1% with 95% confidence. Figure 10(b) plots similar results, but where the shared loss rate is .05. We see that for all protocols, link-overuse remains below 5 for reasonable loss rates. By having the sender coordinate joins as in the Coordinated protocol, link-overuse remains below 2.5, even when there are 100 receivers within the session, each of whose data-path contains the shared link. We observed negligible changes in the results when we increased the number of receivers beyond 100. Since our previous results indicate that link-overuse is highest when all receivers have identical end-to-end loss rates, we can conclude that sender-coordinated congestion control protocols can keep link-overuse below 2.5. This is low enough so that, in networks where multi-rate protocols make up a small percentage of sessions, multi-rate protocols will yield fair allocations with sufficiently desirable fairness properties.



(a) Low shared loss .0001



(b) High shared loss .05

Fig. 10. 100 receivers, 8 layers

## V. RELATED / FUTURE WORK

The application of layering, in the context of video transmission, to maximize usage of available bandwidth and the benefits of coordination of receiver join events within a session is discussed in [13], and further explored in [10]. Clever use of parity coding techniques extend layering's applicability to reliable multicast [4], [16], [24]. Preliminary experiments and definitions of various forms of fairness for layered approaches are explored in [10], [24], as well as in [11], which discusses at a high level how using a layered approach can change the max-min fair allocation. An examination that uses fairness metrics to compare various allocation strategies for layered multicast protocols is presented in [9]. There, the authors argue that link bandwidth should be allocated to sessions in some manner that is proportional to the number of receivers in the session because doing so increases the average "receiver satisfaction". However, none of these works look consider how multi-rate approaches affect fairness properties (in comparison to single-rate approaches) throughout a large-scale network.

Much of the remaining work that deals with multicast fairness assumes that sessions are single-rate [3], [15], [23], [25], and therefore compromise fairness from the receiver perspective, due to tight binding of receiver rates within a session. There has been some work that discusses how one might choose a single-rate session's rate in order to maximize a measure of fairness on a per-receiver basis [6].

There are numerous issues that remain open with regard to using layering to achieve multi-rate max-min fairness. The effects of layering on desirable fairness properties for other definitions of fairness is one possible avenue for examination. We believe that many of our results can be directly applied to TCP-fairness by constructing a definition of max-min fairness where receiver rates are assigned weights (i.e., a receiver's rate is weighted by the inverse of round trip time). It would also be interesting and useful to extend definitions of fairness to multicast sessions with multiple senders. There are also many issues that deal with the practicality of using layering to achieve fairness. One question that comes to mind is whether priority dropping schemes for layered approaches [1] might aid in reducing link-overuse by

increasing coordination among receivers. Also, multicast routing technology must be improved to make layered approaches practical for congestion control and fairness purposes. For instance, join and leave latencies complicate coordination among various receivers within a session, which is likely to increase link-overuse. We believe that long leave latencies will also increase link-overuse (a link continues to receive at the rate prior to the leave, until the leave takes effect, while the receiver's rate reduces immediately). We expect that many such problems are solvable, perhaps with the aid of active routing technology [22]. For instance, placing the decision to add and drop layers at the active nodes, rather than at receivers, should increase the coordination of the joins and leaves of layers by downstream receivers, thereby reducing link-overuse. Such an approach would make a link-overuse of 1.0 feasible for a layered multi-rate session.

It is also unclear whether bandwidth can be shared fairly by sessions that measure fairness on different timescales (i.e., use different quanta), especially in networks like the Internet, where a session's fair allocation may vary due to startup and/or termination of other sessions within the network. Finally, our models contain numerous simplifications of what exists in practice; they are merely used to illustrate concepts, identify challenges, and provide a basic understanding of what can be expected in practice. Extensive development and testing is still necessary to verify that our hypotheses presented here do in fact occur in practice.

## VI. CONCLUSION

We have explored how multi-rate multicast, achievable using layered multicast approaches, can impact fairness within a network. In particular, we showed that in theory, multi-rate sessions can achieve several desirable fairness properties that cannot be achieved in general networks using single-rate sessions. In a practical environment, we demonstrate how receivers can join and leave layers so that their rates are max-min fair over a long term average. Unfortunately, this join-leave process has several practical difficulties. One difficulty that we address is link-overuse: an excessive use of bandwidth by a session over a link shared by multiple receivers in the session. High link-

overuse not only leads to failure of several fairness properties from a session perspective (i.e., fairness of session link rates), but is also likely to reduce most receivers' fair rates. However, our subsequent analysis shows that based on the portion of network sessions that are expected to be multi-rate, practical solutions can keep the amount of link-overuse low enough such that layering can be used to improve fairness within multicast networks.

## APPENDIX

### I. EXPECTED BANDWIDTH WITH RANDOM JOINS

We compute the expected bandwidth for session  $S_i$  on a link  $l_j$ . For simplicity, we write  $R = |R_{i,j}|$ , and denote the set of receivers from session  $S_i$  whose data-path utilizes this link (i.e.,  $R_{i,j}$ ) as  $\{r_1, \dots, r_R\}$ , and let  $a_t$  be the number of packets that receiver  $r_t$  must receive per quantum.

Let  $\rho$  packets be transmitted in a time quantum, and let  $X_i$  be a random variable that equals 1 if any receiver is joined when packet  $i$  is transmitted, and 0 otherwise ( $1 \leq i \leq \rho$ ). Let  $Y_{i,t}$  be a random variable that equals 1 if receiver  $r_t$  joins to receive packet  $i$ , and 0 otherwise. Since we assume a receiver chooses the packets it is to receive from a uniform distribution, we have  $\Pr(Y_{i,t} = 1) = a_t/\rho$ .

$$\begin{aligned} E[X_i] &= 1 - \prod_{t=1}^R \Pr(Y_{i,t} = 0) = 1 - \prod_{t=1}^R (1 - a_t/\rho) \\ E[U_{i,j}] &= E\left[\sum_{i=1}^{\rho} X_i\right] = \sum_{i=1}^{\rho} E[X_i] \\ &= \rho \left(1 - \prod_{t=1}^R (1 - a_t/\rho)\right) \end{aligned}$$

## REFERENCES


- [1] S. Bajaj, L. Breslau, and S. Shenker. Uniform Versus Priority Dropping for Layered Video. In *Proceedings of SIGCOMM'98*, Vancouver, Canada, September 1998.
- [2] D. Bertsekas and R. Gallager. *Data Networks*. Prentice-Hall, Englewood Cliffs, NJ, 1992.
- [3] S. Bhattacharyya, D. Towsley, and J. Kurose. The Loss Path Multiplicity Problem for Multicast Congestion Control. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [4] J. Byers, M. Luby, M. Mitzenmacher, and A. Rege. A Digital Fountain Approach to Reliable Distribution of Bulk Data. In *Proceedings of SIGCOMM'98*, Vancouver, Canada, September 1998.
- [5] Z. Cao and E. Zegura. Utility Max-Min: An Application-Oriented Bandwidth Allocation Scheme. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [6] T. Jiang, M. Ammar, and E. Zegura. Inter-Receiver Fairness: A Novel Performance Measure for Multicast ABR Sessions. In *Proceedings of ACM SIGMETRICS'98*, Madison, WI, June 1998.
- [7] F. Kelly. Charging and Rate Control for Elastic Traffic. *European Transactions on Telecommunications*, 8, 1997.
- [8] F. Kelly and D. Tan. Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability. *Journal of the Operational Research Society*, 49, 1998.
- [9] A. Legout, J. Nonnenmacher, and E. Biersack. Bandwidth Allocation Policies for Unicast and Multicast Streams. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [10] X. Li, S. Paul, and M. Ammar. Layered Video Multicast with Retransmissions (LVMR): Evaluation of Hierarchical Rate Control. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [11] X. Li, S. Paul, and M. Ammar. Multi-Session Rate Control for Layered Video Multicast. In *Proceedings of Symposium on Multimedia Computing and Networking*, San Jose, CA, January 1999.
- [12] J. Madhavi and S. Floyd. TCP-Friendly Unicast Rate-Based Flow Control. Note sent to e2e mailing list, January 1997.
- [13] S. McCanne, V. Jacobson, and M. Vetterli. Receiver Driven Layered Multicast. In *Proceedings of SIGCOMM'96*, Stanford, CA, August 1996.
- [14] K.K. Ramakrishnan and S. Floyd. A Proposal to Add Explicit Congestion Notification (ECN) to IP. RFC 2481, January 1999.
- [15] I. Rhee, N. Balaguru, and G. Rouskas. MTCP: Scalable TCP-like Congestion Control for Reliable Multicast. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.
- [16] L. Rizzo and L. Vicisano. RMDP: An FEC-based Reliable Multicast Protocol for Wireless Environments. *Mobile Computing and Communications Review*, 2(2), April 1998.
- [17] D. Rubenstein, J. Kurose, and D. Towsley. The Impact of Multicast Layering on Network Fairness. Technical report, University of Massachusetts, CMPSCI 99-08, February 1999. A shorter version of this report appeared in *Proceedings of ACM SIGCOMM'99*, September 1999, Cambridge, MA.
- [18] D. Rubenstein, J. Kurose, and D. Towsley. The Impact of Multicast Layering on Network Fairness. Technical report, Columbia University, NETWORK 02-01, January 2002. This report is an extension of University of Massachusetts technical report CMPSCI 99-08.
- [19] S. Sarkar and L. Tassiulas. Distributed Algorithms for Computation of Fair Rates in Multirate Multicast Trees. In *Proceedings of IEEE INFOCOM'00*, Tel-Aviv, Israel, March 2000.
- [20] S. Sarkar and L. Tassiulas. Fair Allocation of Discrete Bandwidth Layers in Multicast Networks. In *Proceedings of IEEE INFOCOM'00*, Tel-Aviv, Israel, March 2000.
- [21] S. Shenker. Making Greed Work in Networks: A Game-Theoretic Analysis of Switch Service Disciplines. In *Proceedings of ACM SIGCOMM'94*, London, UK, August 1994.
- [22] D. Tennenhouse, J. Smith, D. Sincoskie, D. Wetherall, and G. Minden. A Survey of Active Network Research. *IEEE Communications Magazine*, January 1997.
- [23] H. Tzeng and K. Siu. On Max-Min Fair Congestion Control for Multicast ABR Service in ATM. *IEEE JSAC*, 15(3), April 1997.
- [24] L. Vicisano, L. Rizzo, and J. Crowcroft. TCP-like Congestion Control for Layered Multicast Data Transfer. In *Proceedings of INFOCOM'98*, San Francisco, CA, March 1998.
- [25] H. Wang and M. Schwartz. Achieving Bounded Fairness for Multicast and TCP Traffic in the Internet. In *Proceedings of SIGCOMM'98*, Vancouver, Canada, September 1998.
- [26] M. Yajnik, S.B. Moon, J. Kurose, and D. Towsley. Measurement and Modeling of the Temporal Dependence in Packet Loss. In *Proceedings of IEEE INFOCOM'99*, New York, NY, March 1999.

**Dan Rubenstein** received the B.S. degree in mathematics from MIT, the M.A. degree in mathematics from UCLA and a Ph.D in computer science from the University of Massachusetts Amherst in 1992, 1994 and 2000. Since 2000, he has been a member of the Department of Electrical Engineering and has a courtesy appointment in the Department of Computer Science at Columbia University in New York City. Dr. Rubenstein is a recipient of the ACM SIGMETRICS Best Student Paper Award in 2000.

**Jim Kurose** Jim Kurose received a Ph.D in computer science from Columbia University. He is currently a Professor in the Department of Computer Science at the University of Massachusetts, where he is also co-director of the Networking Research Laboratory of the Multimedia Systems Laboratory. Professor Kurose has been a Visiting Scientist at IBM Research, at INRIA Sophia Antipolis, and at Institut Eurecom. His research interests include real-time and multimedia communication, network and operating system support for servers, and modeling and performance evaluation.

Dr. Kurose is the past Editor-in-Chief of the IEEE Transactions on Communications and of the IEEE/ACM Transactions on Networking. He has been active in the program committees for IEEE Infocom, ACM SIGCOMM, and ACM SIGMETRICS conferences for a number of years. He has received a number

of awards for his teaching. With Keith Ross, he is the author of an introductory networking textbook, "Computer Networking, a top down approach featuring the Internet," published by Addison Wesley Longman. He is a Fellow of the IEEE and ACM, and a member of Phi Beta Kappa, Eta Kappa Nu, and Sigma Xi.



**Don Towsley** Don Towsley holds a B.A. in Physics (1971) and a Ph.D. in Computer Science (1975) from University of Texas. From 1976 to 1985 he was a member of the faculty of the Department of Electrical and Computer Engineering at the University of Massachusetts, Amherst. He is currently a Distinguished Professor at the University of Massachusetts in the Department of Computer Science. He has held visiting positions at IBM T.J. Watson Research Center, Yorktown Heights, NY (1982-1983); Laboratoire MASI, Paris, France (1989-1990); INRIA, Sophia-

Antipolis, France (1996); and AT&T Labs - Research, Florham Park, NJ (1997). His research interests include networks, multimedia systems, and performance evaluation.

He currently serves as an Area Editor of the Journal of the ACM and on the Editorial board of Performance Evaluation, and has previously served on several editorial boards including those of the IEEE Transactions on Communications and IEEE/ACM Transactions on Networking. He was a Program Co-chair of the joint ACM SIGMETRICS and PERFORMANCE '92 conference. He is a member of ACM, ORSA and the IFIP Working Groups 6.3 and 7.3.

He received the 1998 IEEE Communications Society William Bennett Paper Award and several conference best paper awards. Last, he has been elected Fellow of both the ACM and IEEE.