

# Optimization-Based Congestion Control for Multicast Communications<sup>\*</sup>

Jonathan K. Shapiro, Don Towsley, and Jim Kurose

Department of Computer Science  
University of Massachusetts at Amherst

**Abstract.** Widespread deployment of multicast depends on the existence of congestion control protocols that are provably fair to unicast traffic. In this work, we present an optimization-based congestion control mechanism for single-rate multicast communication with provable fairness properties. The optimization-based approach attempts to find an allocation of rates that maximizes the aggregate utility of the network. We show that the utility of multicast sessions must be carefully defined if a widely accepted property of aggregate utility is to hold. Our definition of session utility amounts to maximizing a weighted sum of simple utility functions, with weights determined by the number of receivers. The fairness properties of the optimal rate allocation depend both on the weights and form of utility function used. We present analysis for idealized topologies showing that while our mechanism is not strictly fair to unicast, its unfairness can be controlled by appropriate choices of parameters.

## 1 Introduction

Widespread deployment of multicast communication in the Internet depends critically on the existence of practical congestion control mechanisms that allow multicast and unicast traffic to share network resources fairly. Most service providers recognize multicast as an essential service to support a range of emerging network applications including audio and video broadcasting, bulk data delivery, and teleconferencing. Nevertheless, these network operators have been reluctant to enable multicast delivery in their networks, often citing concerns about the congestion such traffic may introduce. There is a clear need for multicast congestion control algorithms that are provably fair to unicast traffic if these concerns are to be addressed. In this paper, we present a congestion control mechanism for single-rate multicast traffic based on an economic theory of resource allocation and show that although it is not strictly fair to unicast traffic, its unfairness is bounded and can be controlled.

We first formulate the multicast congestion control problem as a utility maximization problem, extending existing work for unicast. A naive, *sender oriented*,

---

<sup>\*</sup> This work was supported in part by the National Science Foundation (NSF) under grant number ANI-9980552

generalization of existing formulations for unicast treats single-source multicast sessions no differently from unicast sessions, modeling each by an unweighted utility function and maximizing the sum of session utilities. One problem with this naive approach is that it penalizes individual multicast sessions for using more network resources than unicast sessions without rewarding them for the bandwidth saved on links shared by multiple receivers. More serious than its unfairness to multicast sessions, the sender-oriented approach turns out to violate a generally accepted property of aggregate utility, namely, that the preference of the aggregate does not change if we simply measure utility on a different scale. This common-sense notion is why, for example, we reject as nonsense the statement that, as a group, residents of New York prefer a temperature of 70 degrees to 60 degrees Fahrenheit, but prefer a temperature of 15.5 to 21 degrees Celsius. If this invariance property is violated in the congestion control problem, the network will be controlled about an operating point determined by an arbitrary choice of utility scale. We introduce a *receiver-oriented* approach that uses session weights based on the number of receivers and preserves invariance under a change in utility scale. Moreover, we show that such an approach is necessary a necessary condition for satisfying this property.

A consequence of adding session weights based on the number of receivers is that the resulting rate allocations tend to favor sessions with more receivers over those with fewer. Since the weighted sum does not remove the original penalty against sessions that use more resources, it is not immediately clear whether multicast sessions fare better or worse than unicast sessions under our formulation. We show that while our formulation favors multicast sessions, the resulting unfairness can be controlled and remains bounded in the simple network topologies we have considered.

Our work is based on a promising economics-inspired approach called *optimization based congestion control*, which casts the congestion problem as one of utility maximization (alternately, cost minimization). This approach provides an elegant theoretical framework in which congestion signals are interpreted as prices, network users are modeled as utility maximizers, and the network sets prices in such a way to drive a set of self-interested users toward an operating point at which their aggregate utility is maximized. Specific link service disciplines and rate-control algorithms at end-hosts can be thought of as components of a distributed computation to solve the global optimization problem. Thus, improvements in congestion control can proceed in a principled fashion, driven by improvements in the underlying optimization algorithm. While the optimization-based approach has received much attention [1,2,3,4,5,6,7,8], it has only recently been applied to multicast congestion control [9,10]. Many existing multicast congestion control schemes [11,12] rely on heuristic techniques, such as adapting to a single receiver or a small group of representatives. In contrast, the optimization-based approach offers a formal foundation with which to develop congestion control mechanisms and understand their fairness properties and impact on the global behavior of the network.

The rest of this paper is structured as follows: In Section 2 we extend a unicast congestion control problem formulation to single-rate multicast. In Section 3 we consider multicast session utility functions in detail, presenting a axiomatic argument in favor of a particular definition. The fairness properties of our definition are analyzed in Sections 4-6 where we show that multicast sessions are favored over unicast sessions and present evidence that such unfairness can be controlled. We conclude by briefly discussing the development of practical control mechanisms based on our results and highlighting future work.

## 2 Problem Formulation

Optimization-based congestion control casts the problem of bandwidth sharing as one of utility maximization. Consider a network modeled as a set of directed links  $L$ , with capacity  $c_l$  for each link  $l \in L$ . Let  $C = (c_l, l \in L)$ . The workload for the network is generated by a set of sessions<sup>1</sup>  $S$ , which consume bandwidth. The set of links used by a particular session,  $s$ , is  $L(s) \subseteq L$ . The set of sessions using any particular link,  $l$ , is  $S(l) \subseteq S$ .

Each session is characterized by a utility function  $U_s$ , which is assumed to be increasing and concave in the session rate  $x_s$ . Session utility may also be a function of other parameters in addition to rate, such as number of receivers, but we will sometimes suppress these additional dependencies in the notation, writing  $U_s(x_s)$ . The network’s objective is to optimize social welfare, defined as the sum of session utilities.

$$\max_{x_s \geq 0} \sum_{s \in S} U_s(x_s) \tag{1}$$

$$\text{subject to } \sum_{s \in S(l)} x_s \leq c_l \quad l \in L \tag{2}$$

The problem (1-2) can be solved using convex optimization techniques [13]. Under a standard economic interpretation, the Lagrange multipliers of such techniques are referred to as *shadow prices* and can be shown to function as prices of network links [14]. The essential step in developing practical rate-control algorithms is to find a distributed algorithm for solving (1-2) in which each individual session need only compute a local optimization to set its own rates. There is a growing body of research devoted to finding such a distributed algorithm and using it as a basis for unicast rate-control in practical protocols [1,2,3,4,5,6,8,7, 15].

Observe that the topologies of sessions are not explicit in the formulation. For a unicast session, the links of  $L(s)$  are arranged end-to-end, forming an acyclic path between a source and a receiver. However,  $L(s)$  can be any subset of links—for example, a tree in the case of multicast. There is a requirement that the session employs a single rate  $x_s$  on all of its links. Thus, this formulation is readily generalized to single-rate multicast sessions.

<sup>1</sup> The terms ‘session’ and ‘user’ are synonymous in this paper.

The case of multi-rate multicast is somewhat more complicated since the single rate requirement is replaced with a constraint that a session's rate on link  $l$ , be the maximum of its rate on any downstream link. Since a session can now have different rates on different links, it makes little sense to endow the session with a utility that is a function of scalar rate. Instead, recent treatments of the multi-rate problem [9,10] have altered the model by associating a utility function with each receiver. It is worth observing that this change to the model, while arising naturally from the multi-rate constraint, has been introduced without consideration of its effect on the global operating point. While a complete solution to the multi-rate problem is beyond the scope of this paper<sup>2</sup>, our work provides a formal justification for the use of receiver utility functions even in the case of single-rate multicast where no such modelling pressure exists.

## 2.1 Application to Multicast

Single-rate multicast represents an important class of multicast applications. Many important applications, such as bulk data transfer [16] typically operate at a single rate. Even applications such as streaming video, for which multi-rate multicast is often considered well suited, single-rate multicast is used in current practice. It remains unclear whether multi-rate multicast for video is viable on the Internet, where it must be implemented using layered multicast schemes that have substantial overhead [17]. Furthermore, even if layered multicast is used, single-rate congestion control techniques may be appropriate to adapt the rates of individual layers.

It would appear that congestion control for single-rate multicast is a trivial extension of the unicast problem and can take advantage of existing approaches. It is important, however, to evaluate this claim carefully, given the importance of single rate multicast in practice. Certainly there are implementation issues in multicast that complicate the extension of unicast optimization-based rate control protocols based on packet marking schemes [5,6,4]. Equally serious, are the conceptual difficulties that arise in an uncritical application of the unicast solution of the underlying optimization problem. It is not immediately clear what the fairness properties of the resulting rate allocations are and, more fundamentally, what it means to define a utility function for a multicast session.

To develop our intuitions about the conceptual problems mentioned above, consider the approach of Low and Lapsley [5]. This approach finds a solution to problem (1-2) by solving its dual to obtain the following optimality condition for session rate  $x_s$ :

$$x_s(\lambda^s) = U_s'^{-1}(\lambda^s) \quad (3)$$

$$\lambda^s = \sum_{l \in L(s)} \lambda_l \quad (4)$$

---

<sup>2</sup> Solving the resulting multi-rate optimization problem is further complicated by the coupling of problem variables due to the multi-rate constraint and because the max function is nondifferentiable. See recent works by Kar, Sarkar, and Tassioulas [9] and Deb and Srikant [10] for treatments of the multi-rate problem.

where  $\lambda_l$  is the shadow price of link  $l$  and  $U'_s{}^{-1}$  is the inverse marginal utility function for session  $s$ . It can be shown that  $U'_s{}^{-1}(\lambda^s)$ , and, hence, the rate allocated to session  $s$  is a decreasing function of the total session price  $\lambda^s$ . A large multicast session typically uses many more links than would a unicast session between the source and any single receiver and can therefore expect to see a higher session price than the unicast session. since  $U'_s{}^{-1}(\lambda^s)$  is decreasing, multicast sessions will receive lower rates than unicast sessions along the same end-to-end paths, casting doubt on whether individual receivers have any incentive to join multicast groups. It may be reasonable to adopt a new definition of session utility with a bias in favor of multicast sessions to encourage bandwidth sharing. However, we must be careful not to overcompensate for the high session prices seen by multicast sessions, yielding allocations one would not consider fair to unicast sessions.

In the following sections we analyze the impact of such definitions on the fairness properties of the resulting congestion control mechanism. It will turn out that the class of receiver-oriented session utility functions, while not absolutely fair to unicast sessions, does not starve them in the presence of larger sessions. Moreover, we will see that utility functions in this class make sense in a way that other functions do not.

### 3 Multicast Utility Functions

In Section 2.1, we generalized the unicast optimization problem formulation to accommodate single rate multicast sessions. However, there is a subtle problem with this model that makes it difficult to apply to single-rate multicast. The problem concerns the definition of utility for an individual multicast session. An unweighted utility function is used to characterize the benefit of a higher rate to the session. For a unicast session, it makes little difference whether we consider this benefit to belong to the sender or receiver. For the purpose of unicast congestion control, we can treat the sender's and receiver's objectives as being one and the same. For a multicast session with multiple receivers it is unclear whether session utility belongs to the sender or should be split in some way among the receivers.

One approach towards defining multicast session utility ignores the multiplicity of receivers and defines it only with respect to the sender.<sup>3</sup> An alternative approach would be to define session utility as a function of the utilities of the receivers in the session. We informally refer to these two approaches as *sender-oriented* and *receiver-oriented*, respectively. While a receiver-oriented approach emerges naturally from the model in the case of multi-rate multicast, it is not immediately clear which approach is most appropriate for single-rate multicast. Later in this section we will formalize these definitions and argue in favor of a receiver-oriented approach. Before doing so, however, we will digress briefly to provide some background about the use of utility functions in economics and the theory of social choice.

<sup>3</sup> We are assuming that multicast sessions have a single source.

### 3.1 Utility Functions and Social Welfare

The use of concave increasing utility functions to represent session utility has a natural and intuitive interpretation. Utility is a monotonically increasing function of its input when individuals prefer having as much of the input as possible. The concavity of the utility function captures the idea of diminishing marginal utility<sup>4</sup>. Both concavity and monotonicity are appropriate assumptions in the case of bandwidth for elastic traffic [18], where the input to the utility function is the session rate.<sup>5</sup>

Utility can be difficult to quantify precisely; there is no clear unit of utility and no agreed upon scale. Comparing the utility of two individuals can be tricky, particularly when they do not share the same utility function. Because of the difficulty in performing interpersonal comparisons of utility, economists customarily think of utility as an *ordinal magnitude*, meaning that the absolute magnitude of utility is meaningless, but that the relative magnitudes of utilities at various rates for an individual session define preferences among rates and the relative differences in magnitude indicate the strength of the preferences [19]. A consequence of considering only ordinal magnitudes is that utility functions are unique only up to a linear transformation. That is, the utility maximizing behavior of an individual with utility function  $u(x)$  is indistinguishable from one whose utility function is a linear transformation of  $u(x)$ . This restriction makes intuitive sense because a linear transformation simply represents a change in scale and a translation of the zero point of the utility function.

The notion of an aggregate utility function is a compelling extension of the concept of individual utility. Aggregate utility is defined by a *social welfare function* (SWF) that maps the vector of all session utilities to a scalar utility value representing the social desirability of the corresponding vector of rates. Since the SWF is not one-to-one, it induces a partial ordering over allocations of rates, known as the *social preference relation* (SPR). As with individual utility functions, we are primarily interested in this preference relation rather than the absolute magnitude of the SWF.

In optimization-based congestion control, we adopt the sum of individual utilities as the SWF. In general, there are many ways to define the SWF, each carrying with it some subjective judgment about how individual preferences should be combined to determine a social preference. It is possible to specify desirable properties for a SWF axiomatically. Perhaps the most important result of social choice theory is Arrow's Impossibility Theorem, which states that it is not always possible to satisfy all desiderata [20].

For example, a commonly cited property of SWFs is *independence of irrelevant alternatives*, which states that the socially preferred allocation should be

<sup>4</sup> The term 'marginal utility' is used in economics to refer to the first derivative of the utility function.

<sup>5</sup> In this section, utility will be assumed to be a function of session rate; we do so for the sake of concreteness and continuity with the rest of the paper. It should be understood, however, that the discussion presented here applies to any utility function.

invariant under a change in individual *utility functions* that leaves individuals' *preferences* unaffected. It is straightforward to show that the sum of individual utility functions violates this property. Indeed, it is precisely this violation that allows Kelly to associate optimal rates under different functional forms of utility with different formal definitions of fairness [1].

Although independence of irrelevant alternatives is neither required nor (in light of Arrow's Impossibility Theorem) worth pursuing for the congestion control application, a related but weaker property is still worthy of consideration.

- **Invariance Under Linear Transformation (ILT):** Let  $u$  be a vector of utility functions and  $v$  be a transformed vector such that  $v_i(x) = \alpha u_i(x) + \beta$ . Let  $U(u(x))$  be a SWF, where  $u(x) = (u_i(x_i))$  is the vector of session utilities for rate vector  $x$ . We say that a SWF is *invariant under a linear transformation* if, for any two rate vectors  $x$  and  $y$ ,

$$U(u(x)) \geq U(u(y)) \Rightarrow U(v(x)) \geq U(v(y))$$

for any values of  $\alpha, \beta$ . In words, the SWF induces the same preference relation for  $u$  and  $v$ .

The ILT property builds on the assertion that individual utility functions are unique up to a linear transformation, saying that aggregate preferences, too, should be invariant under such a transformation. We will see shortly that under some definitions of multicast session utility the ILT property is satisfied, while under others it is not.

### 3.2 Sender- and Receiver-Oriented Utility Functions

We now formally define sender- and receiver-oriented concepts of session utility. Consider a single-rate multicast session  $s$  with rate  $x$  and receiver set  $\mathcal{R}$  with size  $R$ . In the sender-oriented approach, session utility function is a concave increasing function  $u_s$  of the session rate.

$$U_{snd} = u_s(x) \tag{5}$$

In the receiver-oriented approach, each receiver  $i \in \mathcal{R}$  has a utility function  $u_i(x)$ , which is concave and increasing. The session utility function is the sum of receiver utilities.

$$U_{rcv} = \sum_{i \in \mathcal{R}} u_i(x) \tag{6}$$

We can convert these definitions into an alternate form by introducing two requirements. First, we require that all receivers in a session have identical utility functions.

$$u_i(x) = u_r(s) \quad \forall i \in \mathcal{R}$$

We typically think of utility functions as representing application characteristics and sometimes as being imposed by network mechanisms. For example, following

the example of Kunniyur and Srikant [4], we use  $u(x) = -1/x$  to model TCP-style congestion control.<sup>6</sup> To the extent that receivers within a session share the same application requirements, it is also reasonable to assume they share a utility function. We feel that this is a natural assumption in the case of single-rate multicast. The second requirement is that both sender- and receiver-oriented utility functions should reduce to the same standard unicast utility function up to a linear transformation when  $R = 1$ . These two restrictions allow us to express both types of session utility functions as the product of a base utility function  $u(\cdot)$  and a scaling function  $f(\cdot)$ . The base utility function,  $u(\cdot)$  depends only on the session rate and is concave and increasing. It can be thought of as the utility function of a session with a single receiver. The scaling function  $f(\cdot)$  depends on the number of receivers in the session. It must be monotonic in its argument, although it need not be strictly increasing.

For a sender-oriented definition of session utility,  $f(R) = \kappa$ , where  $\kappa$  is a constant.

$$U_{snd}(x, R) = \kappa u(x) \quad (7)$$

For a receiver-oriented definition,  $f(R) = \kappa R$ , where  $\kappa$  is a constant.

$$U_{rcv}(x, R) = \kappa R u(x) \quad (8)$$

It is possible to entertain other definitions of session utility. We choose these because they are commonplace and mathematically tractable. One obtains equation (7) by treating all sessions equivalently, regardless of the number of receivers. Equation (8) reflects the idea that multicast session utility is itself a social welfare function, representing the aggregate utility of the receiver set. Under our assumptions, this equation is equivalent to the sum of receiver utilities—a simple and commonly used social welfare function

### 3.3 The Session-Splitting Problem

In Section 3.2, we identified two alternative definitions of multicast session utility based on sender- or receiver orientation. Now we consider these two definitions in more detail and determine which makes sense in the context of congestion control. We begin by attempting to capture the effect of flexible group membership using an optimization-based approach. Golestani and Sabnani [23] observe that if receivers in a session can be dropped and reassigned to a different session in response to congestion, it is often desirable to split a multicast group into subgroups with different rates. One can think of this approach as an approximation of multi-rate multicast that does not violate the constraint of having a single rate per session and requires less overhead at the receiver [17].

The presence of additional sessions in the network after splitting may increase contention on existing bottlenecks or even create new bottlenecks. Thus not all

<sup>6</sup> A more accurate TCP utility function is  $u(x) = (\sqrt{2}/T) \tan^{-1}(Tx/\sqrt{2})$ , where  $T$  is the round trip time [21,22]. Kunniyur and Srikant's approximate function  $u(x) = -1/x$  is valid for small end-to-end loss probabilities.



ways of splitting a session lead to an overall improvement in received rates. Ideally, one would like to find a way to split the session that offers a higher rate to some receivers without reducing the rates of any others. A less ideal, but perhaps still tolerable split might reduce some receivers' rates but improve the utilization of the network and allow many more receivers to receive at a higher rate. In this section, we consider the use of sender- and receiver-oriented social welfare functions to determine whether splitting a session will improve aggregate utility.

In general, the choice of sender- or receiver-oriented utility as well as the form of the base utility function will affect the social welfare function. However, for a fixed choice of these factors, we expect the SWF to be well-defined for all possible ways of splitting the session. Additionally, the optimal way of splitting a session should be invariant under a linear transformation of the base utility function. If this were not the case, an arbitrary rescaling of utility could determine whether splitting a session is preferred over not splitting. We will observe that this invariance holds in the case of a receiver-oriented SWF but not in the case of a sender-oriented one.

We begin by formalizing the session splitting problem in terms of utility maximization. In the session-splitting problem, we have a network  $(N, L)$  with link capacities  $C = (c_l, l \in L)$ . A set of receivers  $R \subset N$  could be served by one or more multicast sessions with source  $s \in N - R$ . We assume that the number and rates of all other sessions in the network are fixed. Capacities in  $C$  thus represent the available capacity for multicast sessions serving receiver set  $R$ . Each session's rate is limited by its most constrained receiver, that is, by the receiver with the lowest link capacity along the path between it and the source. If this bottleneck link is not shared by all of the receivers, then it may be possible to split the session into two or more sessions yielding a higher rate to some receivers.

Splitting the session is equivalent to partitioning the receiver set into disjoint subsets  $P = \{P_1, P_2, \dots, P_N\}$ . We will use  $\mathcal{P}$  to denote the set of all possible partitions of  $R$ . Each partition in  $\mathcal{P}$  represents one possible way to divide the receiver set into sessions. Each element of a partition represents a subset of  $R$  to be served by a different session. Rates may vary among sessions, but all receivers within a session must receive at a single rate. Computing the rates for each session is, itself, a non-trivial problem since some links will be shared by more than one session. There are many possible mechanisms for determining session rates. One example is the greedy algorithm suggested by Rubenstein, Kurose and Towsley [24] to achieve max-min fairness among the sessions.

For our purposes, it is sufficient to assume that we have some deterministic mechanism to perform this rate assignment, which we model as a rate allocation function  $X : \mathcal{P} \times \mathbb{Z}^+ \rightarrow \mathbb{R}$ . Given a partition  $P$  and an index  $i$ , the rate allocation function returns the rate of the session serving  $P_i$ .

The session-splitting problem requires us to find a partition that maximizes the aggregate utility of the network. Recall that the optimization-based approach defines aggregate utility as the (possibly weighted) sum of all session utilities.

Thus the optimal splitting is a partition that solves

$$\max_{P \in \mathcal{P}} U(P; f, u, X)$$

where

$$U(P; f, u, X) = \sum_{i=1}^{|P|} f(|P_i|) u(X(P, i))$$

is the aggregate utility function. We can choose the scaling function  $f$  from equations (7) and (8) to solve this problem for sender- and receiver-oriented definitions of session utility.

The aggregate utility function defines a partial ordering over  $\mathcal{P}$ . In economic terms, this ordering is the social preference relation over all possible partitions of the receiver set. As explained in Section 3.1, it is customary to regard utility functions as unique up to a linear transformation. A reasonable restriction, therefore, is only to allow social preference relations that remain invariant under a linear transformation of the base utility function, as captured by the following axiom, similar to the ILT property in Section 3.1:

**Axiom 1** *Let  $f(\cdot)$  be a fixed scaling function and  $X(\cdot, \cdot)$  be a fixed rate allocation function. For any base utility function  $u(\cdot)$ , let  $v(\cdot)$  be another base utility function such that*

$$v(x) = \alpha u(x) + \beta$$

where  $\alpha$  and  $\beta$  are constants. Then for all  $P, Q \in \mathcal{P}$ ,

$$\begin{aligned} U(P; f, u, X) \geq U(Q; f, u, X) \\ \iff U(P; f, v, X) \geq U(Q; f, v, X) \end{aligned}$$

**Theorem 1.** *Let  $f_{snd}(R) = \kappa$  and  $f_{rcv}(R) = \kappa R$  be sender- and receiver-oriented scaling functions. For any base utility function  $u$  and rate allocation function  $X(\cdot, \cdot)$ , the aggregate utility function  $U(\cdot; f_{rcv}, u, X)$  satisfies Axiom 1, while  $U(\cdot; f_{snd}, u, X)$  does not. Furthermore, Axiom 1 can only be satisfied using the scaling function  $f(R) = f_{rcv}(R)$ .*

The proof of Theorem 1 is quite straightforward and is omitted here due to space limitations. Interested readers can find it in [25]. One immediate consequence of this theorem is that if one accepts that Axiom 1 is indeed an appropriate requirement for any “reasonable” definition of aggregate utility, then our sender-oriented utility definition is not “reasonable”. In fact, the only reasonable definition of session utility is a receiver-oriented one.

## 4 Consequences of Receiver-Oriented Utility

In Section 3.3 we argued that receiver oriented session utility functions are an appropriate model for multicast session utility in the session splitting problem.

In this section, we return to the original congestion control problem and determine whether using receiver-oriented utility functions leads to fair sharing of bandwidth between unicast and multicast sessions. We rewrite the network optimization problem (1-2) as

$$\max_{x=(x_s, s \in S)} \sum_s \kappa R_s u(x_s) \tag{9}$$

$$\text{subject to } \sum_{s \in S(l)} x_s \leq c_l, \quad \forall l \in L \tag{10}$$

The Kuhn-Tucker conditions for optimality are

$$\kappa R_s du/dx = \sum_{l \in L(s)} \lambda_l \tag{11}$$

$$\lambda_l(x^l - c_l) = 0, \quad (x^l - c_l) \leq 0 \tag{12}$$

where the  $\lambda_l$  are Lagrange multipliers or link prices and  $x^l = \sum_{s \in S(l)} x_s$  is the aggregate rate seen at link  $l$ . As before, we also write  $\lambda^s = \sum_{l \in L(s)} \lambda_l$  as the total session price seen by session  $s$ .

From the first Kuhn-Tucker condition (11), we observe that the use of receiver-oriented utility functions creates a bias in favor of sessions with large numbers of receivers. To see this, note that

$$du/dx = \lambda^s (\kappa R_s)^{-1} \tag{13}$$

The optimal rate for session  $s$ ,  $x_s^*$  is given by

$$x_s^* = u'^{-1} (\lambda^s (\kappa R_s)^{-1}) \tag{14}$$

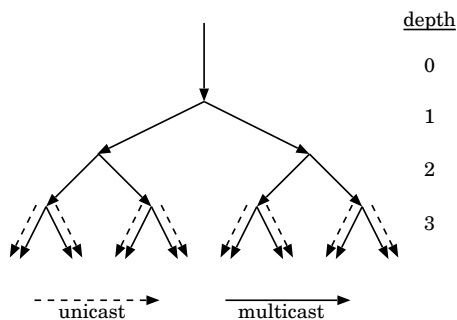
Equation (13) states that, at optimality, the a session’s marginal utility should be proportional to its price divided by the number of receivers. We refer to the ratio  $\lambda^s/(\kappa R_s)$  as the *effective session price*. The optimal rate can therefore be obtained by taking the inverse of the marginal utility function as shown in (14). Since  $U_s$  is concave,  $u'$  is a strictly decreasing function of  $x$  and its inverse is also a decreasing function. For a fixed session price, a session with a larger number of receivers has a lower effective session price and thus receives a higher rate. We refer to this effect as “tyranny of the majority” (ToM).

ToM is a source of unfairness against unicast flows since multicast flows with the same total session price will receive a higher rate. However, the fact that multicast sessions tend to use more links than unicast sessions, particularly as the number of receivers becomes large, means that the session price  $\lambda^s$  for a multicast flow is likely to be higher than that of a unicast session. To understand the fairness properties of rate allocations under receiver oriented utility functions we must determine whether the price increase associated with the scaling of multicast trees is sufficient to limit the effect of ToM as more receivers are added.<sup>7</sup>

<sup>7</sup> If one holds that improving the rate of many receivers at the expense of a few is reasonable, giving a larger share of bandwidth to larger groups may not seem unfair.

## 5 Effect of Multiple Points of Congestion

In the previous section, we saw that ToM and the scaling of multicast trees have opposite effects. As we will see shortly, these effects are not necessarily equal in strength. The effect of ToM is likely to be the stronger of the two, allowing sessions with more receivers to receive a greater share of bandwidth. Whether we choose to accept this form of controlled unfairness or introduce a correction, we require a more precise understanding of the interaction of the two effects. In this section, we show that the functional form of the base utility function can be chosen to limit the strength of the ToM bias.



**Fig. 1.** A binary multicast tree of depth 3 with a sharing depth of 3.

Consider a multicast session in the form of a complete tree of degree  $k$  and depth  $D$ , such as the one shown in Fig. 1. Each link of the tree has capacity  $c$ . We will use a receiver-oriented definition of session utility, but allow an arbitrary base utility function  $u(x)$ . The tree has a receiver at each leaf, giving  $R = k^D$  receivers in total. The multicast session shares the network with a set of  $k^D$  one-hop unicast sessions, which are evenly distributed across the links at depth  $d$ . There are  $k^{D-d}$  unicast sessions on each of  $k^d$  links at level  $d$ . We refer to  $d$ , the depth in the multicast tree at which unicast sessions share links, as the *sharing depth*. Let  $x = (x_s)$  be the vector of session rates, where  $x_0$  is the multicast rate and  $x_1, \dots, x_R$  are the rates of the unicast sessions. Shadow prices are represented by a vector of multipliers  $\lambda = (\lambda_1, \dots, \lambda_L)$ , where  $L = k^d$ .<sup>8</sup> For a particular choice of sharing depth  $d$ , we can now form the Lagrangian for the basic optimization problem (1).

---

We take the position that a bounded bias in favor of large groups is a defensible form of “controlled unfairness” but that there must be a mechanism to prevent starvation of unicast flows.

<sup>8</sup> The vector  $\lambda$  contains elements for only those links with nonzero price, namely, the  $k^d$  links at depth  $d$ .

$$L_d(x, \lambda) = k^D u(x_0) + \sum_{i=1}^R u(x_i) - \sum_{j=1}^L \lambda_j g_j(x) \tag{15}$$

where  $\{g_j\}$  is the set of capacity constraints for the shared links.

$$g_j(x) = x_0 + \sum_{l=L(j-1, k, d, D)}^{L(j, k, d, D)-1} x_l - c \leq 0 \tag{16}$$

$$L(j, k, d, D) = j k^{D-d}$$

We use the symmetry of the tree topology to reduce the problem to three variables: the multicast session rate  $x_m$ , the unicast session rate  $x_i$ , and the shadow price of a congested link  $\lambda$ . We rewrite the link capacity constraint

$$g_j(x) = g(x) = x_m + k^{D-d} x_i - c \tag{17}$$

Solving the first-derivative conditions of the reduced problem for the logarithmic base utility function  $u(x) = \log(x)$  gives

$$x_m = c/2, \quad x_i = c/(2 k^{D-d}), \quad \lambda = 2 k^{D-d}/c \tag{18}$$

We observe the following facts about this result:

- At the system optimum, the multicast session receives rate  $x_m = c/2$ . This result is independent of the tree depth  $D$ , the sharing depth  $d$ , and the tree degree  $k$ .
- The invariance of the optimal multicast rate is a direct result of the choice of a logarithmic base utility function. As we will see, this property does not hold for other utility functions.
- The remaining capacity on the shared links is split evenly among the sharing unicast sessions. Since the number of sharing sessions is  $k^{D-d}$ , the optimal unicast rate depends on  $D$ ,  $d$  and  $k$ .
- The total price charged to the multicast session is

$$\lambda k^d = 2 k^D/c,$$

which is independent of the sharing depth. Under a receiver-oriented definition of session utility, this price is divided by the number of receivers to obtain the effective session price. Thus, effective session price is independent of  $d$ ,  $D$  and  $k$  under a logarithmic utility function.

Since the invariance of the multicast session rate appears to derive from a special choice of utility function, it is interesting to explore the behavior as we modify the functional form. We can derive the following optimality condition:

$$u'(x_m) = \frac{1}{k^{D-d}} u'\left(\frac{c - x_m}{k^{D-d}}\right), \tag{19}$$

Equation (19) relates the marginal utility function  $u'(x)$  to the function  $u'^*(x) = a u'(a(c-x))$  obtained when we rotate  $u'$  about the line  $x = c$  and scale both the argument and the result by the same factor  $a$ . Any point at which these two functions intersect satisfies the optimality condition. Note that  $u'(x)$  is the derivative of a concave and strictly increasing utility function, and therefore must be strictly decreasing. Thus,  $u'(x)$  and  $u'^*(x)$  intersect in exactly one point, establishing the uniqueness of the solution. Observe also that the scaling factor  $a = 1/k^{D-d} \leq 1$ . Scaling the argument of  $u'(c-x)$  compresses the function along the horizontal axis and moves the point of intersection to the left, while scaling its result compresses the function along the vertical axis and moves the point of intersection to the right.

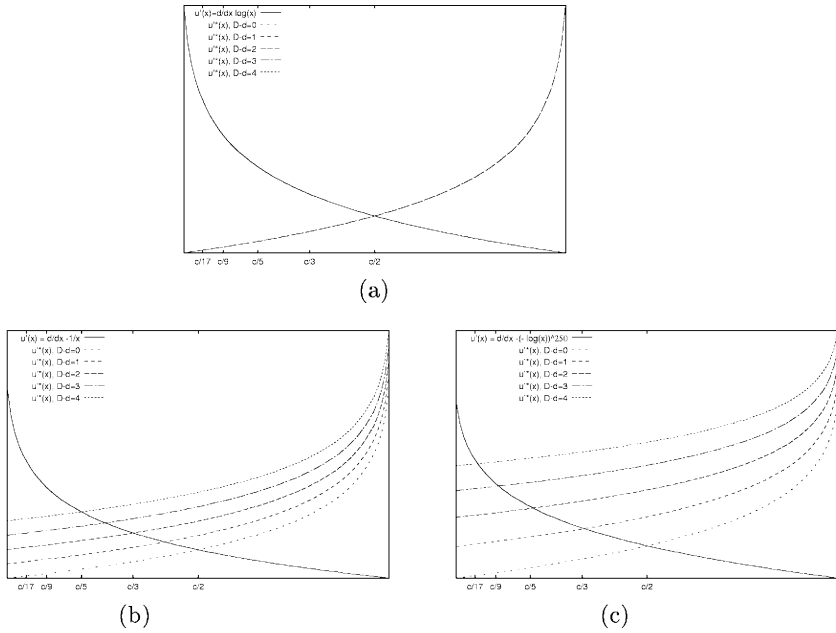
Figure 2 shows how the allocated multicast rate changes as we vary the sharing depth  $d$  (hence, the number of congested links) in a binary tree for three choices of base utility function:  $u(x) = \log(x)$ ,  $u(x) = -1/x$  and  $u(x) = -(-\log(x))^\alpha$ . The first function is the now familiar logarithmic utility function. The second is the *minimum potential delay* (MPD) utility function introduced by Massoulié and Roberts [26] and shown by Kunniyur and Srikant to model the utility of TCP traffic [4]. The third function is shown by Kelly to yield max-min fairness in the limit as  $\alpha \rightarrow \infty$  [1].<sup>9</sup> In all three graphs, the single decreasing function is  $u'(x)$ , the first derivative of the base utility function, and the family of increasing functions are  $u'^*(x)$  for decreasing  $a$  (increasing  $D-d$ ). The points where  $u'^*(x)$  intersects  $u'(x)$  give the optimal rates for the multicast session as a fraction of available capacity.

As established above, the intersection point is invariant and equal to  $c/2$  for logarithmic utility. The intersection point is also fixed at  $c/2$  when  $a = 1$  for all three functions, corresponding to a sharing depth equal to the maximum tree depth. In both the MPD and max-min fair utility functions, however, the intersection point moves to the left as  $a$  decreases. That is, as the sharing depth moves closer to the top of the tree, the *number* of bottleneck links decreases. However, as more unicast sessions share each bottleneck, the *price* on each congested link increases and the multicast session receives a smaller fraction of the available bandwidth.

Under the definition of max-min fairness for single-rate multicast [24], the multicast session must share bandwidth equally with all sessions on its most congested link. Thus, in the max-min fair allocation for our  $k$ -ary tree example,  $x_m = c/(k^{D-d} + 1)$ . In the case of Kelly's max-min fair utility function, we see that the optimal rates approach the max-min fair allocations, indicated by the tick-marks along the x-axis in Fig. 2. It appears that the points of intersection converge to these values as we transform the logarithmic utility function into the max-min fair utility function by increasing the exponent  $\alpha$ . Demonstrating this convergence formally is somewhat difficult.

We can establish a similar result for a family of utility functions that includes both the logarithmic and MPD utility functions and also yields max-min fairness as a limiting case. Consider the family of utility functions  $u(x)$  with first deriva-

<sup>9</sup> In our numerical analysis, we take  $\alpha$  to a reasonably high power. ( $\alpha = 250$ )



**Fig. 2.** Figure showing the effect on the optimal allocation for a binary multicast tree as we vary the sharing depth. These graphs show three different marginal utility functions,  $u'(x)$  along with their transformations  $u'^*(x)$  for various choices of  $D - d$  with the y-axis shown in log scale. The x-coordinate of the points of intersection give the optimal session rates as a fraction of available capacity. Max-min fair allocations for different values of  $D - d$  are indicated along the x-axis. The figure shows that the logarithmic utility function (top) gives the multicast session half the available bandwidth regardless of the number of sharing unicast sessions, whereas the max-min fair utility function (bottom) splits bandwidth evenly among all sessions on the shared link regardless of the number of receivers. The MPD utility function (center) represents a compromise between these two extremes.

tives  $u'(x) = 1/x^{\alpha+1}$ . Such functions include  $u(x) = \log(x)$ ,  $u(x) = -x^{-\alpha}/\alpha$ . This family of functions was originally identified by Mo and Walrand [7]. Members of this family are mathematically tractable since the functions  $u'(x)$  are homogeneous, satisfying

$$u'(tx) = t^{-r} u'(x) \tag{20}$$

where  $r = \alpha + 1$ . We can simplify the optimality condition of (19).

$$x_m / (c - x_m) = a^{(1-r)/r} \tag{21}$$

As a further simplification, we can express the multicast rate as a fraction,  $p$ , of available capacity,  $x_m = pc$ . Solving for  $p$ , we get

$$p = \left(1 + a^{(1-r)/r}\right)^{-1}. \quad (22)$$

In the limit of large  $\alpha$ ,  $p$  converges to the max-min fair allocation.

$$\lim_{r \rightarrow \infty} \left(1 + a^{(1-r)/r}\right)^{-1} = a/(1+a) = 1/(k^{D-d} + 1). \quad (23)$$

Following Kelly's example in [1], we can prove that  $u(x) = -x^{-\alpha}/\alpha$  always gives max-min fairness in the limit  $\alpha \rightarrow \infty$ , by providing an absolute priority to smaller flows.<sup>10</sup> For two rates such that  $x_{s^*} < x_s$ ,

$$u'(x_{s^*})/u'(x) = (x_s/x_{s^*})^{\alpha+1} \rightarrow \infty \quad \text{as } \alpha \rightarrow \infty$$

## 6 Fairness to Unicast Sessions


In Section 5, we observed that a multicast session was able to obtain a higher rate than unicast sessions sharing the same bottleneck links. We showed that this unfairness is bounded in the presence of multiple points of congestion. However, this result exploited features of an idealized multicast session topology. Adopting a somewhat more realistic model in this section, we investigate whether the same type of bounded unfairness is possible in a more general setting with receiver-oriented utility functions. We also consider whether there is any multicast utility function that allows a strictly equal split of shared bottleneck bandwidth between a multicast and a unicast session.

Adopting the fairness objective proposed by Handley, Floyd and Whetten [28]—that the algorithm be provably fair relative to TCP in the steady state, we define a generalized notion of TCP fairness. We say that a multicast session utility function  $U(x; r) = f(R)u(x)$  is *strictly unicast-fair* if the optimal rate for the multicast session is the same as would be obtained by a unicast session with utility function  $u(x)$  along the most congested source-to-receiver path in the multicast tree. This definition is equivalent to TCP-fairness in the case where  $u(x) = -1/x$ , the MPD utility function.

We first show that neither sender nor receiver oriented multicast utility functions lead to strict unicast-fair allocations and derive a result suggesting that strict fairness is difficult to achieve under any definition of session utility. Consider the modified star network topology shown in Figure 3. A single multicast session with source node  $s$  and receivers  $\{1, \dots, R\}$  shares the network with  $R$  unicast sessions, one from  $s$  to each receiver. Link  $l_0$  from the source to the central node is shared by all sessions and has effectively infinite capacity. Each link  $l_i$  from the center to receiver  $i$  is shared by the multicast session and one unicast session. Link  $l_1$  is the bottleneck link, with capacity  $\beta c$ , where  $\beta < 1$  and  $c$  is the capacity of all other links  $l_i$ ,  $i > 1$ . Receiver 1 is the most congested receiver in the multicast session.

<sup>10</sup> A similar result is reported in recent work by Bonald and Massoulié [27].





**Fig. 3.** A multicast tree with a modified star topology. Receiver 1 is most congested.

We give the unicast sessions the MPD utility function  $u(x) = -1/x$ . The multicast function has utility function  $u(x; R) = f(R)u(x)$ . Let  $x_m$  be the rate of the multicast session and  $x_i$  be the rate of the unicast session to receiver  $i$ . A strictly tcp-fair allocation would split the bandwidth on  $l_1$  equally between  $x_m$  and  $x_1$ ,  $x_m = x_1 = \beta c/2$ . We can substitute this rate into the optimality conditions of the optimization problem (9-10) to determine the appropriate scaling function  $f(R)$  that will lead to the tcp-fair allocation, obtaining

$$f(R) = 1 + \frac{(R - 1) \beta^2}{2\sqrt{2} (2 - \beta)^2} \tag{24}$$

This result shows that tcp-fairness can be achieved in the optimization-based framework by maximizing a weighted sum of utilities with weights given by a scaling function  $f(R)$ . However, the presence of  $\beta$ , a topological parameter, in the scaling function suggests that the correct scaling function depends on topological properties of the network.

We now consider a generalized version of the previous example with no explicitly defined network topology. Consider a network containing a set of links  $\mathcal{L}$ . The network is shared by two sessions  $v$  and  $w$ , which have rates  $x_v$  and  $x_w$ , respectively. Each session uses a subset of links in the network and session  $w$  only uses a proper subset of links that are also used by  $v$ . Formally,  $L(w) \subset L(v) \subseteq \mathcal{L}$ . The sessions have  $R_v$  and  $R_w$  receivers with  $R_v > R_w$ . We assume that the path to the most constrained receiver in both  $v$  and  $w$  is the same and is therefore entirely contained in  $L(w)$ . The Lagrangian for the optimization problem is.

$$L(x; \lambda) = f(R_v) u'(x_v) + f(R_w) u'(x_w) + \sum_{l \in L(w)} \lambda_l (x_v + x_w - c_l) + \sum_{l \in L(v) - L(w)} \lambda_l (x_v - c_l) \tag{25}$$

From the Kuhn-Tucker conditions, we derive an optimality condition on the ratio of marginal utilities.

$$u'(x_v)/u'(x_w) = (f(R_w) \lambda^v)/(f(R_v) \lambda^w) \tag{26}$$

where

$$\lambda^v = \sum_{l \in L(v)} \lambda_l, \quad \lambda^w = \sum_{l \in L(w)} \lambda_l \quad (27)$$

Consider the family of base utility functions satisfying  $u'(x) = -1/x^\alpha$ ,  $\alpha \geq 1$ . Note that this family includes both the MPD and logarithmic utility functions. The ratio of session rates is

$$x_w/x_v = ((f(R_w)\lambda^v)/(f(R_v)\lambda^w))^{1/\alpha} \quad (28)$$

In a strictly tcp-fair allocation, the ratio  $x_w/x_v = 1$ . From equation (28), it is clear that the actual value of this ratio depends on both the choice of scaling function and the ratio  $\lambda^v/\lambda^w$ . It is also apparent that the ratio  $x_w/x_v$  approaches 1 in the limit as  $\alpha \rightarrow \infty$ . Thus, the exponent  $\alpha$  offers one way to control unfairness for any choice of scaling function; increasing it moves the resulting rate allocation closer to max-min fairness. However, only the max-min utility function can guarantee strict unicast fairness for an arbitrary choice of scaling function.

If a utility function other than max-min is used, providing strict unicast fairness requires careful selection of the scaling function. For example, strict unicast fairness could be achieved by exploiting a scaling law relating the total price of a multicast session to its number of receivers. Chuang and Sirbu propose such a law for static multicast costs [29] with the form

$$\lambda^s \propto R_s^k \quad (29)$$

The authors empirically evaluate the scaling exponent  $k$ , finding its value to be constant over a wide range of network topologies. This law assumes, however, that link costs in the network are static. To be applicable for the purposes of congestion control, such a scaling law would have to be established for dynamically changing prices that reflect link congestion. If such a scaling law can be found, then strict unicast fairness would result from a multicast session utility function

$$U_s(x) \propto R_s^{-k} u(x).$$

We leave the search for such a scaling law as direction for future research, but note here that, as presented in Section 3.3 the sum of session utilities under such a multicast utility function would not be invariant under a linear transformation of  $u(x)$ .

## 7 Conclusion

This paper presented an optimization based scheme for multicast congestion control based on utility maximization. Appealing to the economic theory underlying this approach, we proposed the use of a receiver oriented definition of session utility. By considering the incentive to split multicast sessions into smaller sessions, we showed that only receiver oriented utility functions ensure that the optimal solution of the session-splitting problem remains invariant under a linear transformation of the utility scale.

We identified two sources of unfairness that arise when maximizing the sum of receiver oriented utility functions, one favoring unicast sessions and one favoring multicast. Unicast sessions are favored because they tend to use fewer links than multicast sessions and thus are charged a lower price for bandwidth. Multicast sessions are favored by the tyranny of the majority effect because the the sum of link prices in the session session is divided by the number of receivers and this reduced price is used to compute the session rate. When these two effects are combined, a net unfairness results that favors sessions with many receivers over sessions with few, with unicast sessions faring worst of all. This unfairness can be controlled, however, by choosing the form of the base utility function. While it is difficult to achieve strict fairness between unicast and multicast traffic, we argue that controlled unfairness is a reasonable goal, particularly as it provides an incentive to use multicast by rewarding larger groups.

Much work still remains to be done in this area. In the work presented here, we have focused on the economic interpretation of the optimization-based approach to reason about the fairness properties resource allocation at system equilibrium. A complimentary line of enquiry concerns the convergence and stability properties of the equilibria in the multicast case. A promising direction of future work is to extend the growing body of relevant research for unicast [30, 31,32,33] to the multicast case.

## References

1. Kelly, F.: Charging and rate control for elastic traffic. *European Transactions on Telecommunications* **volume 8** (1997) 33–37
2. Kelly, F.P., Maulloo, A., Tan, D.: Rate control in communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society* **49** (1998) 237–252
3. Gibbens, R., Kelly, F.: Resource pricing and the evolution of congestion control. *Automatica* **35** (1999) 1969–1985
4. Kunniyur, S., Srikant, R.: End-to-end congestion control: utility functions, random losses and ecn marks. In: *Proc. INFOCOM*. (2000)
5. Low, S.H., Lapsley, D.E.: Optimization flow control, i: Basic algorithm and convergence. *IEEE/ACM Transactions on Networking* (1999)
6. Athuraliya, S., Lapsley, D., Low, S.: An enhanced random early marking algorithm for internet flow control. In: *Proc. INFOCOM*. (2000)
7. Mo, J., Walrand, J.: Fair end-to-end window-based congestion control. *IEEE/ACM Transactions on Networking* (1999)
8. Golestani, S., Bhattacharyya, S.: A class of end-to-end congestion control algorithms for the internet. In: *Proc. ICNP'98*. (1998)
9. Kar, K., Sarcar, S., Tassiulas, L.: Optimization based rate control for multirate multicast sessions. In: *Proceedings of Infocom 2001*. (2001)
10. Deb, S., Srikant, R.: Congestion control for fair resource allocation in networks with multicast flows. In: *Proc. of the IEEE Conference on Decision and Control*. (2001)
11. Kasera, S., Bhattacharyya, S., Keaton, M., Kiwior, D., Kurose, J., Towsley, D., Zabele, S.: Scalable fair reliable multicast using active services. *IEEE Networks Magazine* (2000)

12. Rizzo, L.: pgmcc: A TCP-friendly single-rate multicast. In: Proceedings of SIGCOMM. (2000) 17–28
13. Madden, P.: Concavity and Optimization in Microeconomics. Basil Blackwell (1986)
14. Hillier, F.S., Lieberman, G.J.: Introduction to Mathematical Programming. 2 edn. McGraw-Hill (1995)
15. P.Key, McAuley, D., Barham, P., Laevens, K.: Congestion pricing for congestion avoidance. Technical Report MSR-TR-99-15, Microsoft Research (1999)
16. Byers, J.W., Luby, M., Mitzenmacher, M., Rege, A.: A digital fountain approach to reliable distribution of bulk data. In: Proc. Sigcomm '98. (1998)
17. Li, X., Ammar, M.H., Paul, S.: Video multicast over the internet. IEEE Networks Magazine (1999)
18. Shenker, S.: Fundamental design issues for the future internet. IEEE J. Selected Areas Comm. **13** (1995) 1176–1188
19. Hirshleifer, J., Hirshleifer, D.: Price Theory and Applications. 6 edn. Prentice Hall (1997)
20. Arrow, K.J.: Social Choice and Individual Values. 2 edn. Yale Univ. Press (1963)
21. Kelly, F.P.: Mathematical modelling of the internet. In: Proc. 4th International Congress on Industrial and Applied Mathematics. (1999)
22. Low, S.H.: A duality model of tcp flow controls. In: Proc. ITC Specialist Seminar on IP Traffic Measurement, Modeling, and Management. (2000)
23. Golestani, S.J., Sabnani, K.K.: Fundamental observations on multicast congestion control in the internet. In: Proc. INFOCOM. (1999)
24. Rubenstein, D., Kurose, J., Towsley, D.: The impact of multicast layering on network fairness. In: Proc. SIGCOMM 99. (1999)
25. Shapiro, J.K., Towsley, D., Kurose, J.: Optimization-based congestion control for multicast communications. Technical Report UM-CS-2000-033, University of Massachusetts at Amherst (2000)
26. Massoulié, L., Roberts, J.: Bandwidth sharing: Objectives and algorithms. In: Proc. INFOCOM. (1999)
27. Bonald, T., Massoulié, L.: Impact of fairness on internet performance. In: Proc. ACM SIGMETRICS. (2001)
28. Handley, M., Floyd, S., Whetten, B.: Strawman specification for tcp friendly (reliable) multicast congestion control. Technical report, Reliable Multicast Research Group (1998)
29. Chuang, J., Sirbu, M.: Pricing multicast communications: A cost-based approach. In: Proc. INET'98. (1998)
30. Low, S.H., Paganini, F., Doyle, J.C.: Internet congestion control. IEEE Control Systems Magazine (2002) February
31. Johari, R., Tan, D.: End-to-end congestion control for the internet: Delays and stability. To appear in IEEE/ACM Transactions on Networking (2001)
32. Massoulié, L.: Stability of distributed congestion control with heterogeneous feedback delays. Technical report, Microsoft Research (2000)
33. Hollot, C., Misra, V., Towseley, D., Gong, W.: On designing improved controllers for aqm routers supporting tcp flows. In: Proc. IEEE Infocom 2001. (2001)