

Bandwidth Allocation Policies for Unicast and Multicast Flows

Arnaud Legout
Institut EURECOM
Sophia Antipolis, FRANCE
legout@castify.net

Jörg Nonnenmacher
Castify Networks
Sophia Antipolis, FRANCE
nonnen@castify.net

Ernst W. Biersack
Institut EURECOM
Sophia Antipolis, FRANCE
erbi@eurecom.fr

April 13, 2001

Technical Report, Institut Eurecom, April 2001

Abstract

Using multicast delivery to multiple receivers reduces the aggregate bandwidth required from the network compared to using unicast delivery to each receiver. However multicast is not yet widely deployed in the Internet. One reason is the lack of incentive to use multicast delivery. To encourage the use of multicast delivery, we define a new bandwidth allocation policy, called *LogRD*, taking into account the number of downstream receivers. This policy gives more bandwidth to a multicast flow as compared to a unicast flow that shares the same bottleneck, however without starving the unicast flows. The *LogRD* policy provides also an answer to the question on how to treat a multicast flow compared to a unicast flow sharing the same bottleneck.

We investigate three bandwidth allocation policies for multicast flows and evaluate their impact on both receiver satisfaction and fairness using a simple analytical study and a comprehensive set of simulations. The policy that allocates the available bandwidth as a logarithmic function of the number of receivers downstream of the bottleneck achieves the best trade-off between receiver satisfaction and fairness.

Keywords: Unicast, Multicast, Bandwidth Allocation Policies.

1 Introduction

There is an increasing number of applications such as software distribution, audio/video conferences, and audio/video broadcasts where data is destined to multiple receivers. During the last decade, multicast routing and multicast delivery have evolved from being a pure research topic [7] to being experimentally deployed in the MBONE [11] to being supported by major router manufacturers and offered as a service by some ISPs. As a result, the Internet is becoming increasingly multicast capable. Multicast routing establishes a **tree** that connects the source with the receivers. The multicast tree is rooted at the sender and the leaves are the receivers. Multicast delivery sends data across this tree towards the receivers. As opposed to unicast delivery, data is not copied at the source, but is copied inside the network at branch points of the multicast distribution tree. The fact that only a *single copy* of data is sent over a link that leads to multiple receivers results in a bandwidth gain of multicast over unicast whenever a sender needs to send simultaneously to multiple receivers. Given R receivers, the **multicast gain** for the network is defined as the ratio of unicast bandwidth cost to multicast bandwidth cost, where bandwidth cost is the product of the delivery cost of one packet on one link

and the number of links the packet traverses from the sender to the R receivers for a particular transmission (unicast or multicast). In case of shortest path unicast and multicast routing between source and receivers, the multicast gain for the model of a full o -ary multicast tree is¹:

$$\log_o(R) \cdot \frac{R}{R-1} \cdot \frac{o-1}{o}$$

Even for random networks and multicast trees different from the idealized full o -ary tree, the multicast gain is largely determined by the logarithm of the number of receivers [21, 24].

Despite the widespread deployment of multicast capable networks, multicast is rarely provided as a service and network providers keep the multicast delivery option in their routers turned off. However, multicast results in bandwidth savings for the ISPs and allows the deployment of new services like audio/video broadcast. Several reasons contribute to the unavailability of multicast; multicast address allocation, security, network management, billing, lack of congestion control, lack of an incentive to use multicast are among the reasons that slow down the deployment of multicast (see [8] for a detailed discussion about the deployment issues for the IP multicast service). In this paper, we address how to increase the incentive to use multicast from the receiver's point of view. It could be argued that as multicast consumes less resources than unicast, a service using multicast should be charged less than the same service using unicast. However, as multicast is expensive to be deployed and probably more expensive to be managed (group management, pricing, security, etc.) than unicast, it is not clear whether a provider will charge less for multicast than for unicast. As discussed by Diot [8], multicast is only cost-effective for an ISP when it results in significant bandwidth savings. Indeed, as multicast is significantly more expensive than unicast, it is most of the time worthwhile to support small groups with unicast. We believe that the main incentive for a provider to use multicast is that multicast enables the deployment of new services that scale with a large number of receivers, for example audio and video broadcast.

The problem of providing receivers with an incentive to use multicast is very difficult. In general, users want high satisfaction, but do not care whether the provider uses unicast or multicast to deliver the content. The argument that multicast allows applications to scale with a large number of receivers is not a good argument for a user because it does not change the user's satisfaction, except if the service cannot be provided without multicast due to a very large number of receivers. If we give more bandwidth to multicast, a multicast user will experience a higher satisfaction than a unicast user which results in an incentive to use multicast.

We saw that it is not easy to establish precisely who benefits how much from multicast. However, we saw that multicast allows to deploy new services. Therefore, it is very important to give a receiver-incentive to use multicast in order to give to the receivers an indisputable benefit to use multicast. We want to give an incentive to use multicast by rewarding the multicast gain in the network to the receivers; at the same time we want to treat² unicast traffic fairly relative to multicast traffic. The two motivations for increasing the bandwidth share for multicast compared to unicast are: First, to give a receiver-incentive to use multicast; Second, to favor multicast due to its significant bandwidth saving. We believe that the second point can be highly controversial. It does not seem fair to give the same amount of bandwidth to a flow serving one receiver and to another one serving ten millions receivers. However, the notion of fairness is subjective and debatable.

We investigate bandwidth allocation policies that allocate the bandwidth locally at each single link to unicast and multicast traffic, and we evaluate globally the bandwidth perceived by the receivers. For three different bandwidth

¹See section 3.1 for some insights on the multicast gain and appendix A for a rigorous proof of the results.

²The problem of treating fairly unicast and multicast traffics is related to the more general question of how multicast flows should be treated in comparison to a unicast flow sharing the same bottleneck.

allocation policies, we examine the case where a unicast network is augmented with a multicast delivery service and evaluate the receiver satisfaction and the fairness among receivers.

The rest of the paper is organized as follows. In Section 2 we present the three bandwidth allocation strategies, and introduce the model and the assumptions for their comparison. In Section 3 we give some insights into the multicast gain, and we analytically study the strategies for simple network topologies. In Section 4 we show the effect of different bandwidth allocation policies on a hierarchical network topology. In Section 5 we discuss the practical issues of our strategies, and Section 6 concludes the paper.

2 Model

2.1 Assumptions

We examine, in this paper, how to best allocate the bandwidth of a link between competing unicast and multicast traffic. We consider scenarios with a given number k of unicast sources, a given number m of multicast sources, a different number M of receivers per multicast source, and a different bandwidth C for each network link to be allocated among the source-destination(s) pairs.

For this study, we make several assumptions and simplifications. The assumptions are: i) Knowledge in every network node about every flow S_i through an outgoing link l . ii) Knowledge in every network node about the number of receivers $R(S_i, l)$ for flow S_i reached via an outgoing link l . iii) Each node is making the bandwidth allocation independently. A particular receiver sees the bandwidth that is the minimum bandwidth of all the bandwidth allocations on the links from the source to this receiver. iv) The sources have the capability to send through different bottlenecks via a cumulative layered transmission [20, 19]. For receivers of the same multicast delivery, the (bottleneck) bandwidth seen by different receivers may be different. In fact, each receiver sees the maximum available bandwidth on the path between the source and the receiver.

These assumptions are not restrictive in the sense that they do not simplify or limit the model. Indeed, i) and ii) are mandatory for per-flow bandwidth allocation with respect to the number of receivers. Weakening assumption ii) to require only, for instance, the knowledge in some network nodes about roughly the number of receivers per flow reached via an outgoing link, is a area for future research. Assumption iii) simply considers independent nodes, and iv) guarantees that the sources are able to get all the available bandwidth.

However, in order to make the evaluation of the model more tractable, we make two simplifications concerning the traffic: i) A constant bit rate traffic for every flow. ii) No arriving or departing flows. Simplification i) means that we do not consider the throughput variations of a flow, for instance, due to congestion control. Therefore, the sources immediately get all the available bandwidth. Simplification ii) means that we do not consider the dynamics of the flows, for instance, in the case of Web traffic (multiple arriving and departing flows to get a web page). As we consider a static scenario, the sources remain stable at the optimal rate. These simplifications are useful to eliminate all side effects and interferences due to dynamic scenarios. We do not claim our model to take into account the dynamics of the real Internet, but to provide a snapshot. At a given moment in time, we evaluate the impact of different bandwidth allocation policies for a given scenario. Adding dynamics to our model would not improve our study, but simply adds complexity in the evaluation of the bandwidth allocation policies. Indeed, the dynamics is not related to the bandwidth allocation policies, but to the ability of the sources to get the available bandwidth. The impact of the dynamics of the flows on the bandwidth allocation policies is, however, an avenue for future research.

2.2 Bandwidth Allocation Strategies

We present three bandwidth allocation policies. It is important to us to employ the bandwidth-efficient multicast without starving unicast traffic and to give at the same time an incentive for receivers to connect via multicast, rather than via unicast. Our objective is twofold: On one hand we want to increase the average receiver satisfaction, on the other hand, we want to assure a fairness among different receivers.

We assume a network of nodes connected via links. At the beginning, we assume every network link l has a link bandwidth C_l . We compare three different strategies for allocating the link bandwidth C_l to the flows flowing across link l . Let n_l be the number of flows over a link l . Each of the flows originates at a source S_i , $i \in \{1, \dots, n_l\}$. We say that a receiver r is **downstream** of link l if the data sent from the source to receiver r is transmitted across link l . Then, for a flow originating at source S_i , $R(S_i, l)$ denotes the **number of receivers that are downstream** of link l . For an allocation policy p , $B_p(S_i, l)$ denotes the bandwidth shared of link l allocated to the receivers of S_i that are downstream of l .

The three bandwidth allocation strategies for the bandwidth of a single link l are:

- **Receiver Independent (RI):** Bandwidth is allocated in equal shares among all flows through a link – independent of the number of receivers downstream. At a link l , each flow is allocated the share:

$$B_{RI}(S_i, l) = \frac{1}{n_l} C_l$$

The motivation for this strategy is: the *RI* strategy does not represent any changes in the current bandwidth allocation policy. This allocation policy weighs multicast and unicast traffic equally. We consider this policy as the benchmark against which we compare the other two policies.

- **Linear Receiver Dependent (LinRD):** The share of bandwidth of link l allocated to a particular flow S_i depends linearly on the number of receivers $R(S_i, l)$ that are downstream of link l :

$$B_{LinRD}(S_i, l) = \frac{R(S_i, l)}{\sum_{j=1}^{n_l} R(S_j, l)} C_l$$

The motivation for this strategy is: given R receivers for S_i downstream of link l , the absence of multicast forces the separate delivery to each of those R receivers via a separate unicast flow³. For a multicast flow, we allocate a share that corresponds to the aggregate bandwidth of R separate unicast flows.

- **Logarithmic Receiver Dependent (LogRD):** The share of bandwidth of link l allocated to a particular stream S_i depends logarithmically on the number of receivers $R(S_i, l)$ that are downstream of link l :

$$B_{LogRD}(S_i, l) = \frac{1 + \ln R(S_i, l)}{\sum_{j=1}^{n_l} (1 + \ln R(S_j, l))} C_l$$

The motivation for this strategy is: multicast receivers are rewarded with the multicast gain from the network. The bandwidth of link l allocated to a particular flow is, just like the multicast gain, logarithmic in the number of receivers that are downstream of link l :

³We assume shortest path routing in the case of unicast and multicast.

Our three strategies are representatives of *classes* of strategies. We do not claim that the strategies we pick are the best representatives of each class. It is not the aim of this paper to find the best representative of a class, but to study the trends between the classes. One can define numerous classes of strategies. We do not claim that one of the three classes of strategies is optimal. However, we restrict ourselves to these three strategies as we believe these policies shed light on the fundamental issues that come with the introduction of the number of receivers in the bandwidth allocation. The following example illustrates the bandwidth allocation for the case of the *Linear Receiver Dependent* policy. We have two multicast flows originating at S_1 and S_2 with three receivers each (see Fig. 1).




Figure 1: Bandwidth allocation for linear receiver-dependent policy.

For link 1, the available bandwidth C_1 is allocated as follows: Since $R(S_1, 1) = 3$ and $R(S_2, 1) = 3$, we get $B_{LinRD}(S_1, 1) = B_{LinRD}(S_2, 1) = \frac{3}{3+3}C_1 = 0.5C_1$. For link 4, we have $R(S_1, 4) = 2$ and $R(S_2, 4) = 1$. Therefore we get $B_{LinRD}(S_1, 4) = \frac{2}{3}C_4$ and $B_{LinRD}(S_2, 4) = \frac{1}{3}C_4$. Given these bandwidth allocations, the bandwidth seen by a particular receiver r is the bandwidth of the bottleneck link on the path from the source to r . For example, the bandwidth seen by receiver R_1^3 is $\min(\frac{1}{2}C_1, \frac{2}{3}C_4, \frac{1}{2}C_6)$.

The way we allocate bandwidth could lead to scenarios where bandwidth needs to be reallocated, we call this the *bandwidth reallocation problem*. Imagine three flows F_1 , F_2 , and F_3 with only one receiver each. The flows F_1 and F_2 share a link l_C of bandwidth C , and flows F_2 and F_3 share a link $l_{\frac{C}{2}}$ of bandwidth $\frac{C}{2}$. With any of the

three policies, the bandwidth allocated on link l_C is $\frac{C}{2}$ for F_1 and F_2 and the bandwidth allocated on link $l_{\frac{C}{2}}$ is $\frac{C}{4}$ for F_2 and F_3 . Therefore, F_2 cannot use its allocated bandwidth $\frac{C}{2}$ on link l_C . However, as we consider static scenarios with constant bit rate flows, the bandwidth that is not used by F_2 cannot be reallocated to F_1 . This is the *bandwidth reallocation problem*. This problem could adversely impact the results of the simulation. One way to solve this problem is to consider dynamic flows which grab the available bandwidth in case of unused bandwidth. This is contrary to the simplifications required by our model. Another way to solve this problem is to statically reallocate the unused bandwidth. However, in case of a complex topology, this leads to convergence problems that are beyond the scope of this paper. In fact, we decided to evaluate, for each simulation, the amount of unused bandwidth, and we found that there is very little unused bandwidth. Therefore, we do not expect the bandwidth reallocation problem to adversely impact the results of our simulations.

2.3 Criteria for Comparing the Strategies

Our goal is to increase the mean *receiver satisfaction*, however, not at the detriment of *fairness*. In order to evaluate *receiver satisfaction* and *fairness*, we define two basic measures, one describing the average user satisfaction, the other one describing the fairness among users.

Receiver Satisfaction

There are many ways to define receiver satisfaction and the most accurate is receiver utility. Unfortunately, utility is a theoretical notion that does not allow to compare the utility of two different receivers and does not give an absolute (i.e. for all receivers) scale of utility [12]. We measure **receiver satisfaction** as the bandwidth an average receiver sees⁴. Let r be a receiver of a source S and let (l_1, l_2, \dots, l_L) be the path of L links from the source to r , then, the bandwidth seen by the receiver r is:

$$B_p^r = \min_{i=1, \dots, L} \{B_p(S, l_i)\} \quad , \quad p \in \{RI, LinRD, LogRD\}$$

With the total number of receivers R of all sources we define the **mean bandwidth** \bar{B}_p as:

$$\bar{B}_p = \frac{1}{R} \sum_{r=1}^R B_p^r \quad , \quad p \in \{RI, LinRD, LogRD\} \quad (1)$$

Jiang et al. [17] introduced a global measure for the throughput delivered via the whole network that is defined as the sum of the mean throughput over all the flows. For the global throughput measure, it is possible to weight multicast flows with a factor R^α , where R is the number of receivers and $0 < \alpha < 1$. To the best of the authors knowledge, the approach of Jiang et al. [17] is the only one taking into account the number of receivers of a multicast flow. While their approach takes into account the number of receivers to measure the global network throughput, our approach is different in two aspects: First, we take the number of receivers into account for the *allocation* of the bandwidth on links and use a policy (*LogRD*) that weights multicast flows in the allocation with the logarithm of the number of receivers. Second, we measure receiver satisfaction with respect to all receivers, not just the ones of a single group.

⁴While there are other criteria to measure satisfaction such as delay or jitter, bandwidth is a measure of interest to the largest number of applications.

Fairness

For inter-receiver fairness, several measures exist, including the product measure [2] and the fairness index [16]. For a discussion of the different measures see [13].

Jiang et al. [17] defined inter-receiver fairness for a single multicast flow as the sum of the receiver's utilities, where utility is highest around the fair share. Due to the intricacies coming with the utility function, we do not consider a utility function and use a fairness measure that takes into account all receivers of all flows.

We use the standard deviation of the bandwidth among receivers to be the measure of choice for inter-receiver fairness.

$$\sigma = \sqrt{\frac{1}{R} \sum_{r=1}^R (\bar{B}_p - B_p^r)^2} \quad , \quad p \in \{RI, LinRD, LogRD\} \quad (2)$$

The key point with this fairness measure is that we consider a notion of fairness independent of the network and of the localization of the bottlenecks. Indeed, each receiver has a given satisfaction. The feeling of fairness for each receiver only depends on the satisfaction of the other receivers, but is independent of any network parameters. For instance, if a receiver has a satisfaction lower than all the other receivers, he will feel a high unfairness even if his low satisfaction is due to a slow modem.

We define **ideal fairness** as the case where all receivers receive the same bandwidth. For **ideal fairness** our measure $\sigma = 0$ has its lowest value. In all other cases, the bandwidth sharing among receivers is unfair and $\sigma > 0$.

Optimality

The question now is how to optimize both *receiver satisfaction* and *fairness*. For the strategy p and the scenario s , let $\sigma(p, s)$ be the function that defines our fairness criteria and $\bar{B}(p, s)$ be the function that defines our receiver satisfaction. An accurate definition of s is: $s + p$ defines the full knowledge of all parameters that have an influence on receiver satisfaction and fairness. So s defines all the parameters without the strategy p . We define

$$\sigma_{max}(s) = \min_p \sigma(p, s)$$

and

$$\bar{B}_{max}(s) = \max_p \bar{B}(p, s)$$

We want to find a function $F(s)$ such as $\forall s: \sigma(F(s), s) = \sigma_{max}(s)$ and $\forall s: \bar{B}(F(s), s) = \bar{B}_{max}(s)$. If such a function $F(s)$ exists for all s , it means that there exists a pair $(F(s), s)$ that defines for all s an optimal point for both *receiver satisfaction* and *fairness*. Feldman [12] shows that *receiver satisfaction* is inconsistent with *fairness*⁵, which means it is impossible to find such a function $F(s)$ that defines an optimal point for both *receiver satisfaction* and *fairness* for all s . So we cannot give a general mathematical criteria to decide which bandwidth allocation strategy is the best. Moreover, in most of the cases it is impossible to find an optimal point for both \bar{B} and σ .

Therefore, we evaluate the allocation policies with respect to the tradeoff between *receiver satisfaction* and *fairness*. Of course, we can define criteria that can apply in our scenarios, for instance, strategy A is better than strategy B if $\frac{\sigma_A}{\sigma_B} \leq L_f$ and $\frac{\bar{B}_A}{\bar{B}_B} \geq I_s$ where L_f is the maximum loss of *fairness* accepted for strategy A and I_s is the minimum

⁵In terms of mathematical economics we can say that Pareto optimality is inconsistent with fairness criteria [12].

increase of *receiver satisfaction* for strategy *A*. But, the choice of L_f and I_s needs a fine tuning and seems pretty artificial to us.

Receiver satisfaction and fairness are criteria for comparison that are meaningful only in the same experiment. It does not make sense to compare the satisfaction and the fairness among different sets of users. Moreover, it is impossible to define an absolute level in satisfaction and fairness. In particular, it is not trivial to decide whether a certain increase in satisfaction is worthwhile when it comes at the price of a decrease in fairness. Hopefully, for our study the behavior of the three strategies will be different enough to define distinct operating points. Therefore, the evaluation of the tradeoff between *receiver satisfaction* and *fairness* does not pose any problem.

3 Analytical Study

We first give some insights into the multicast gain and the global impact of a local bandwidth allocation policy. A rigorous discussion of both points is given in appendix A and appendix B. Then, we compare the three bandwidth allocation policies from Section 2 for basic network topologies in order to gain some insights in their behavior. In Section 4 we study the policies for a hierarchical network topology.

3.1 Insights on Multicast Gain

We can define the multicast gain in multiple ways and each definition may capture very different elements. We restrict ourselves to the case of a full o -ary distribution tree with either receivers at the leaves – in this case we model a point-to-point network – or with broadcast LANs at the leaves. We consider one case where the unicast and the multicast cost only depends on the number of links (the unlimited bandwidth case) and another case where the unicast and the multicast cost depends on the bandwidth used (the limited bandwidth case).

We define the **bandwidth cost** as the sum of all the bandwidths consumed on all the links of the tree. We define the **link cost** as the sum of all the links used on the tree; we count the same link n times when the same data are sent n times on this link. Let C_U be the unicast bandwidth/link cost from the sender to all of the receivers and C_M the multicast bandwidth/link cost from the same sender to the same receivers.

For the bandwidth-unlimited case, every link of the tree has unlimited bandwidth. Let C_U and C_M be the link cost for unicast and multicast, respectively. We define the multicast gain as the ratio $\frac{C_U}{C_M}$. If we consider one receiver on each leaf of the tree, the multicast gain depends logarithmically on the number of receivers. If we consider one LAN on each leaf of the tree, the multicast gain depends logarithmically on the number of LANs and linearly on the number of receivers per LAN (see appendix A.1 for more details).

For the bandwidth-limited case, every link of the tree has a capacity C . Let C_U and C_M be the bandwidth cost for unicast and multicast, respectively. Unfortunately, for the bandwidth-limited case, the multicast gain defined as $\frac{C_U}{C_M}$ makes no sense because it is smaller than 1 for a large number of multicast receivers (see appendix A.2 for more details). We define another measure that combines the satisfaction and the cost that we call cost per satisfaction $GB = \frac{\text{global cost}}{\text{global satisfaction}}$, that tells us how much bandwidth we invest to get a unit of satisfaction. Now, we define the multicast gain as $\frac{GB_U}{GB_M}$ where GB_U and GB_M are the unicast and multicast cost per satisfaction, respectively. If we consider one receiver on each leaf of the tree, the gain depends logarithmically on the number of receivers. If we consider one LAN on each leaf of the multicast tree, the gain depends logarithmically on the number of LANs and linearly on the number of receivers per LAN (see appendix A.2 for more details).

In conclusion, for both the bandwidth unlimited and limited case, the multicast gain has a logarithmic trend with the number of receivers in case of point-to-point networks. The multicast gain has also a logarithmic trend with the number of LANs, but a linear trend with the number of receivers per LAN. Therefore, with a small number of receivers per LANs the multicast gain is logarithmic but with a large number of receivers per LANs the multicast gain is linear. Appendix A gives an analytical proof of these results.

3.2 Insights on the Global Impact of a Local Bandwidth Allocation Policy

In section 2.2, we suggest the *LogRD* policy because we want to reward the multicast receivers with the multicast gain. However, it is not clear whether allocating locally the bandwidth as a logarithmic function of the number of downstream receivers achieves to reward the multicast receivers with the multicast gain, which is a global notion.

To clarify this point, we consider a full o -ary tree for the bandwidth-unlimited case when there is one receiver per leaf. We find (see appendix B for a proof) that the policy that rewards multicast with its gain is the *LinRD* policy and not the *LogRD* policy as expected. If we reward multicast with its real gain using the *LinRD* policy, we will give to multicast the bandwidth that corresponds to the aggregate bandwidth of R separate unicast flows (see section 2.2). However, we have to consider that we use multicast in order to save bandwidth. If we allocate to a multicast flow the same bandwidth than the bandwidth used by R separate unicast flows, the use of multicast makes no sense as it does not save bandwidth compared to unicast. Therefore, rewarding a multicast flow with its gain (as defined in appendix A) makes no sense.

In the following, we will see that the *LinRD* is a very aggressive policy for unicast flows while the *LogRD* policy gives very good results for both the unicast and multicast flows.

3.3 Comparison of the Bandwidth Allocation Policies

3.3.1 Star Topology

We consider the case where k unicast flows need to share the link bandwidth C with a single multicast flow with m downstream receivers, see Fig. 2.

With the *RI* strategy, the bandwidth share of a link is $\frac{1}{k+1}C$ for both a unicast and a multicast flow. The *LinRD* strategy gives a share of $\frac{1}{m+k}C$ to each unicast flow and a share of $\frac{m}{m+k}C$ to the multicast flow. The *LogRD* strategy results in a bandwidth of $\frac{1}{k+(1+\ln m)}C$ for a unicast flow and $\frac{1+\ln m}{k+(1+\ln m)}C$ for the multicast flow.

The mean receiver bandwidths over all receivers (unicast and multicast) for the three policies are:

$$\begin{aligned}\bar{B}_{RI} &= \frac{1}{k+m} \sum_{i=1}^{k+m} \frac{C}{k+1} = \frac{C}{k+1} \\ \bar{B}_{LinRD} &= \frac{1}{k+m} \left(\sum_{i=1}^k \frac{C}{m+k} + \sum_{i=1}^m \frac{mC}{m+k} \right) = \frac{k+m^2}{(k+m)^2} C \\ \bar{B}_{LogRD} &= \frac{1}{k+m} \left(\sum_{i=1}^k \frac{C}{k+(1+\ln m)} + \sum_{i=1}^m \frac{C(1+\ln m)}{k+(1+\ln m)} \right) = \frac{k+m(1+\ln m)}{(k+m)(k+1+\ln m)} C\end{aligned}$$

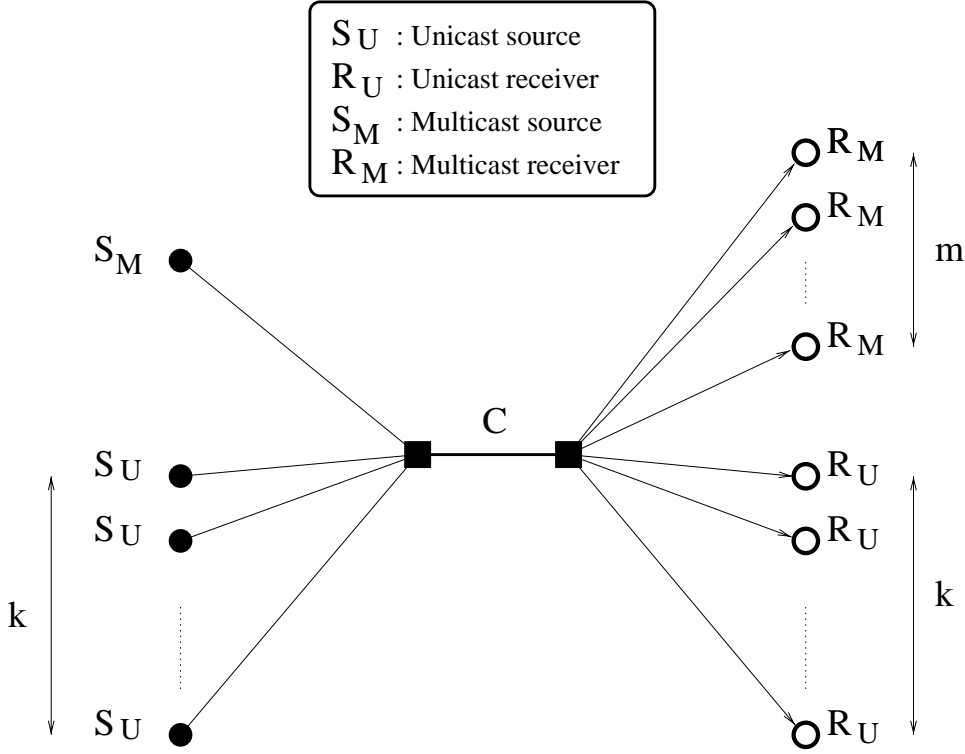


Figure 2: One multicast flow and k unicast flows over a single link.


By comparing the equations for any number of multicast receivers, $m > 1$, and any number of unicast flows $k > 1$ we obtain:

$$\bar{B}_{LinRD} > \bar{B}_{LogRD} > \bar{B}_{RI} \quad (3)$$


The receiver-dependent bandwidth allocation strategies, $LinRD$ and $LogRD$, outperform the receiver-independent strategy RI by providing a higher bandwidth to an average receiver. This is shown in Fig. 3, where the mean bandwidths are normalized by \bar{B}_{RI} , in which case the values depicted express the bandwidth gain of any policy over RI .

Fig. 3(a) shows the mean bandwidth for $m = 60$ multicast receivers and an increasing number of unicasts $k = 1, \dots, 200$. The receiver-dependent policies $LinRD$ and $LogRD$ show an increase in the mean bandwidth when the number of unicasts is small compared to the number of multicast receivers. The increase with the $LogRD$ policy is less significant than the increase with the $LinRD$ policy since the $LogRD$ policy gives less bandwidth to the multicast flow than the $LinRD$ policy for the same number of receivers. Additionally, more link bandwidth is allocated to the multicast flow than in the case of a higher number of unicasts, which result in a lower share for multicast. With an increasing number of unicasts, the gain of $LinRD$ and $LogRD$ decreases.

After assessing the bandwidth gain of $LinRD$ and $LogRD$ for a number of unicast receivers higher than the number of multicast receivers, we turn our attention to the case where the number of multicast receivers is increasing $m = 1, \dots, 200$ and becomes much higher than the number of unicasts ($k = 60$). Fig. 3(b) shows that the mean bandwidth for $LinRD$ and $LogRD$ is increasing to multiples of the bandwidth of RI .



(a) Increasing the number k of unicasts; 60 multicast receivers.



(b) Increasing the size m of the multicast group; 60 unicasts.

Figure 3: Normalized mean bandwidth for the Star topology.

We saw that the receiver-dependent policies significantly reward multicast receivers and that the *LinRD* policy is better than the *LogRD* policy with respect to the receiver satisfaction. Now, we have to study the impact of the receiver-dependent policies on the fairness.




Figure 4: Standard deviation for the Star topology. Increasing the size $m = 1, \dots, 200$ of the multicast group; $k = 60$ unicasts.

The following equations give the standard deviation over all receivers for the three policies:

$$\begin{aligned}\sigma_{RI} &= 0 \\ \sigma_{LinRD} &= C(m-1)\sqrt{\frac{k \cdot m}{(k+m)^3(k+m-1)}} \\ \sigma_{LogRD} &= \frac{C \cdot \ln m}{k+1+\ln m}\sqrt{\frac{k \cdot m}{(k+m)(k+m-1)}}\end{aligned}$$

By comparing the equations for any number of multicast receivers, $m > 1$, and any number of unicast flows $k > 1$ we obtain:

$$\sigma_{LinRD} > \sigma_{LogRD} > \sigma_{RI} \quad (4)$$

While the *LinRD* is the best policy among our three policies with respect to the receiver satisfaction, it is the worst policy in terms of fairness. Fig. 4 shows the standard deviation for $k = 60$ unicast flows and an increasing multicast group $m = 1, \dots, 200$. With the Star topology, all unicast receivers see the same bandwidth and all multicast receivers see the same bandwidth. Between unicast receivers and multicast receivers no difference exists for the *RI* strategy. For the *LinRD* strategy a multicast receiver receives m times more bandwidth than a unicast receiver and for the *LogRD* strategy a multicast receiver receives $(1 + \ln m)$ times more bandwidth than a unicast receiver. The standard deviation for all the receivers is slightly increased with the *LogRD* policy compared to the *RI* policy, and is more significantly increased with the *LinRD* policy compared to the *RI* policy (see Fig. 4).

The high bandwidth gains of the *LinRD* strategy result in a high unfairness for the unicast receivers. For *LogRD*, the repartitioning of the link bandwidth between unicast and multicast receivers is less unequal than in the case of *LinRD*. In summary, the *LogRD* policy leads to a significant increase in receiver satisfaction, while it introduces only a small decrease in fairness. We can conclude that among the three strategies *LogRD* makes the best tradeoff between receiver satisfaction and fairness.

Surprisingly we will obtain nearly the same results in Section 4.3 when we examine the three policies on a large random network. The similarity of the Fig. 3(b), and 4, with the figures of Section 4.3 indicates that the simple Star topology with a single shared link can serve as a model for large networks.

3.3.2 Chain Topology

We now study bandwidth allocation for the case where a multicast flow is traversing a unicast environment of several links. We use a chain topology, as shown in Fig. 5, where k unicast flows need to share the bandwidth with a single multicast flow leading to m receivers. However, the unicast flows do not share bandwidth among each other, as opposed to the previous single shared link case for the star topology.

At each link, the *RI* strategy allocates in $\frac{1}{2}C$ for both the unicast flow and the multicast flow. The *LinRD* strategy results in a share of $\frac{1}{m+1}C$ for the unicast flow and $\frac{m}{m+1}C$ for the multicast flow. The *LogRD* strategy results in a share of $\frac{1}{2+\ln m}C$ for the unicast flow and a share of $\frac{1+\ln m}{2+\ln m}C$ for the multicast flow.

The mean receiver bandwidth for the three cases is:

$$\bar{B}_{RI} = \frac{1}{k+m} \sum_{i=1}^{k+m} \frac{C}{2} = \frac{C}{2}$$

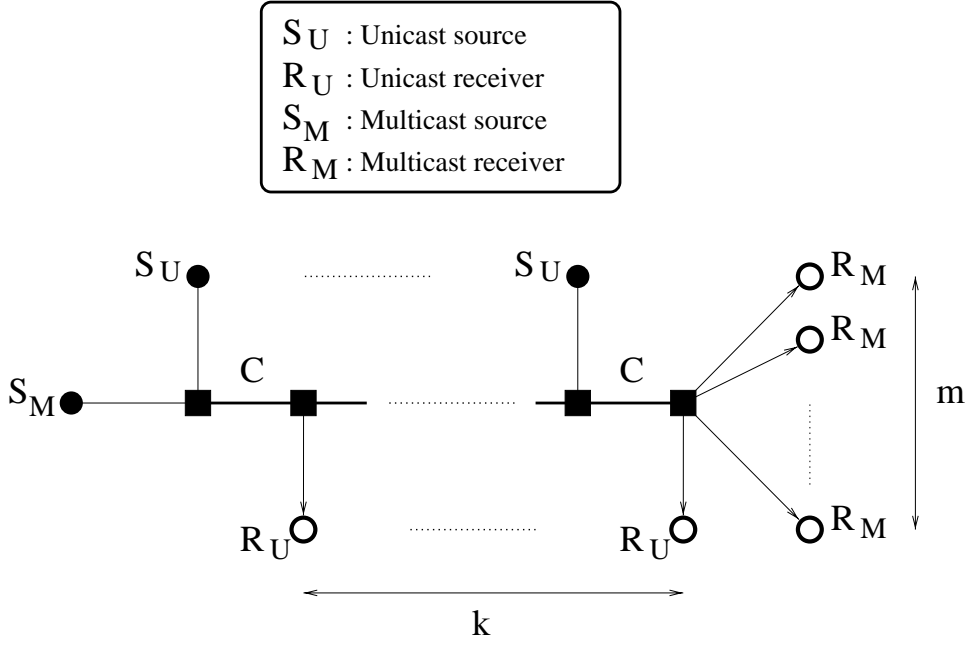


Figure 5: One multicast flow and k unicast flows over a chain of links.

$$\bar{B}_{LinRD} = \frac{1}{k+m} \left(\sum_{i=1}^k \frac{C}{m+1} + \sum_{i=1}^m \frac{m \cdot C}{m+1} \right) = \frac{k+m^2}{(k+m)(m+1)} C$$

$$\bar{B}_{LogRD} = \frac{1}{k+m} \left(\sum_{i=1}^k \frac{C}{2+\ln m} + \sum_{i=1}^m \frac{C(1+\ln m)}{2+\ln m} \right) = \frac{k+m+m \cdot \ln m}{(k+m)(2+\ln m)} C$$

The strategy with the highest mean bandwidth depends on the relation between the number of multicast receivers and the number of unicast flows. If the number of unicasts equals the number of multicast receivers, $k = m$, then all policies result in the same average receiver bandwidth of $C/2$. For all other cases, with $k > 1$ and $m > 1$ we have:

$$\begin{aligned} \bar{B}_{RI} &> \bar{B}_{LogRD} > \bar{B}_{LinRD} \quad , \quad k > m \\ \bar{B}_{LinRD} &> \bar{B}_{LogRD} > \bar{B}_{RI} \quad , \quad k < m \end{aligned} \quad (5)$$

The receiver-dependent policies *LinRD* and *LogRD* perform better than the *RI* policy when the size of the multicast group is larger than the number of unicast sessions. While the number of multicast receivers can increase to large numbers and is only limited by the number of hosts in the network, the number of unicast crossing traffic is limited by the length of the path source-receiver. This is shown in Fig. 6, where the mean bandwidths are normalized by \bar{B}_{RI} , in which case the values depicted express the bandwidth gain of any policy over *RI*.

Fig. 6(a) shows the mean bandwidth for $m = 30$ multicast receivers and an increasing number of unicast sessions $k = 1, \dots, 200$. As the number of unicasts increases, receiver-dependent policies become worse than *RI* policy. Fig. 6(b) shows the mean bandwidth for $k = 30$ unicast receivers and an increasing number of multicast receivers. The receiver-dependent policies perform worse than the *RI* policy for small multicast group sizes, but as the size of

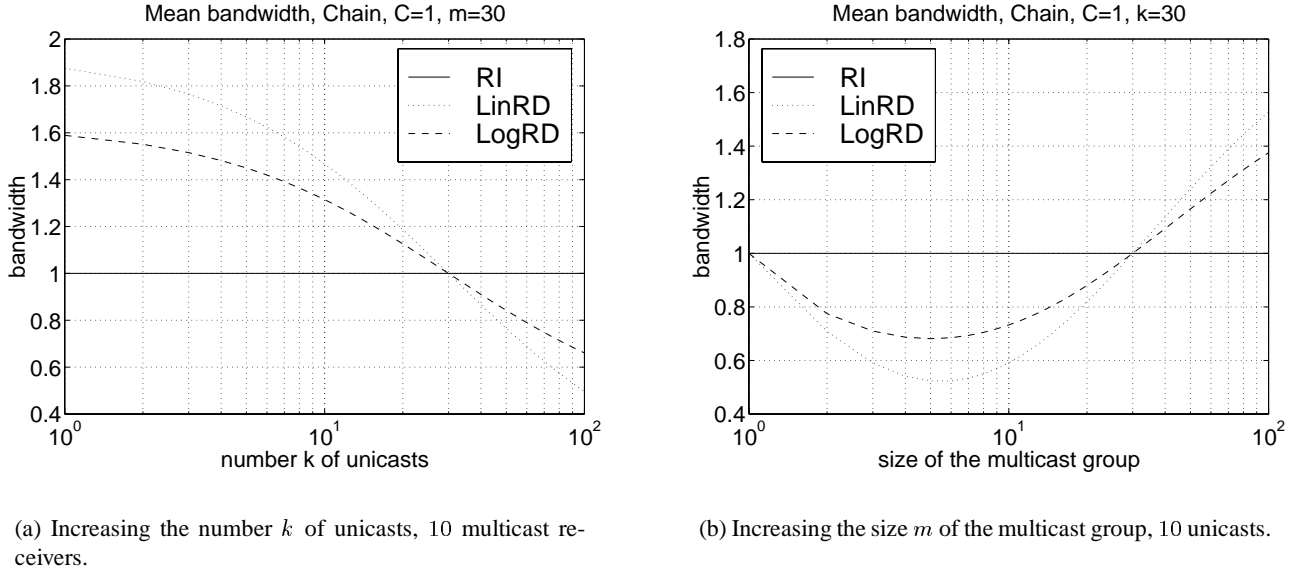


Figure 6: Normalized mean bandwidth for the Chain topology.

the multicast group increases the bandwidth gain for receiver dependent policies increases rapidly. In Fig. 6(b), for the multicast group size $m = 30$, the three policies lead to the same mean bandwidth, for the multicast group size $m = 50$, the *LinRD* policy yields to more than 20% gain over the *RI* policy and the *LogRD* policy yields to more than 15% gain over the *RI* policy.

We see that, concerning the receiver satisfaction, the receiver-dependent policies have a more complex behavior with a chain topology than with a star topology. To complete the study of the chain topology, we look at the fairness.

The standard deviation over all the receivers for the three policies is:

$$\begin{aligned}\sigma_{RI} &= 0 \\ \sigma_{LinRD} &= \frac{C(m-1)}{m+1} \sqrt{\frac{k \cdot m}{(k+m)(k+m-1)}} \\ \sigma_{LogRD} &= \frac{C \cdot \ln m}{2 + \ln m} \sqrt{\frac{k \cdot m}{k+m} k + m - 1}\end{aligned}$$

By comparing the equations for any number of multicast receivers, $m > 1$, and any number of unicast flows $k > 1$ we obtain:

$$\sigma_{LinRD} > \sigma_{LogRD} > \sigma_{RI} \quad (6)$$

The *LinRD* policy, as for the star topology, has to the worst fairness. Fig. 7 shows the standard deviation for $k = 30$ unicast flows and an increasing multicast group $m = 1, \dots, 200$. For *RI*, unicast receivers and multicast receivers obtain the same share, for *LinRD* a multicast receiver receives m times more bandwidth than a unicast

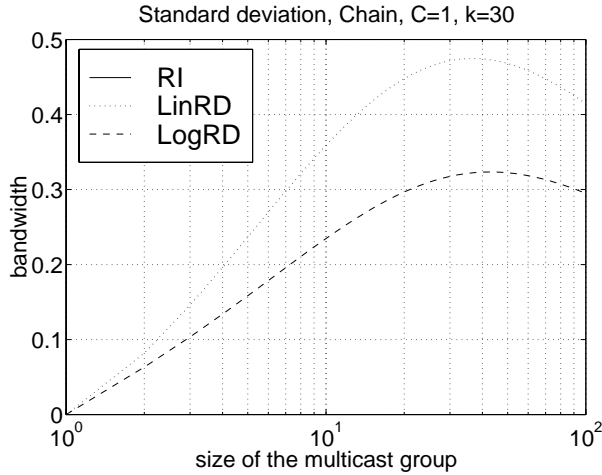


Figure 7: Standard deviation for the Chain topology as a function of the size m of the multicast group for $k = 30$ unicasts.

receiver and for *LogRD* a multicast receiver receives $(1 + \ln m)$ times more bandwidth than a unicast receiver. As the multicast session size m increases, the unicast flows get less bandwidth under the *LinRD* and the *LogRD* strategy, while the *RI* strategy gives the same bandwidth to unicast and multicast receivers. The *LinRD* policy leads to a worse fairness than the *LogRD* policy, however, the gap between the two policies is smaller than with the *Star* topology (compare Fig. 7 and Fig. 4).

We conclude that among the three strategies the *LogRD* strategy achieves for large group sizes the best compromise between receiver satisfaction and fairness. However, for the *Chain* topology the superiority of the *LogRD* policy is not as obvious as for the *Star* topology.

This simple analytical study allowed to identify some principal trends in the allocation behavior of the three strategies studied. The *LogRD* policy seems to be the best compromise between receiver satisfaction and fairness. To deepen the insight gained with our analytical study, we will study the three strategies via simulation on a large hierarchical topology.

4 Simulation

We now examine the allocation strategies on network topologies that are richer in connectivity. The generation of realistic network topologies is subject of active research [3, 10, 25, 26]. It is commonly agreed that hierarchical topologies better represent a real Internet network than do flat topologies. We use `tiers` [10] to create hierarchical topologies consisting of three levels: WAN, MAN, and LAN that aim to model the structure of the Internet topology [10]. For details about the network generation with `tiers` and the used parameters the reader is referred to Appendix C.

4.1 Unicast Flows Only

Our first simulation aims to determine the right number of unicast flows to define a meaningful unicast environment. We start with our random topology *RT* and add at random locations of the LAN-leaves unicast senders and unicast

receivers. The number of unicast flows ranges from 50 to 4000. Each simulation is repeated five times and averages are taken over the five repetitions. We compute for each plot 95% confidence intervals.

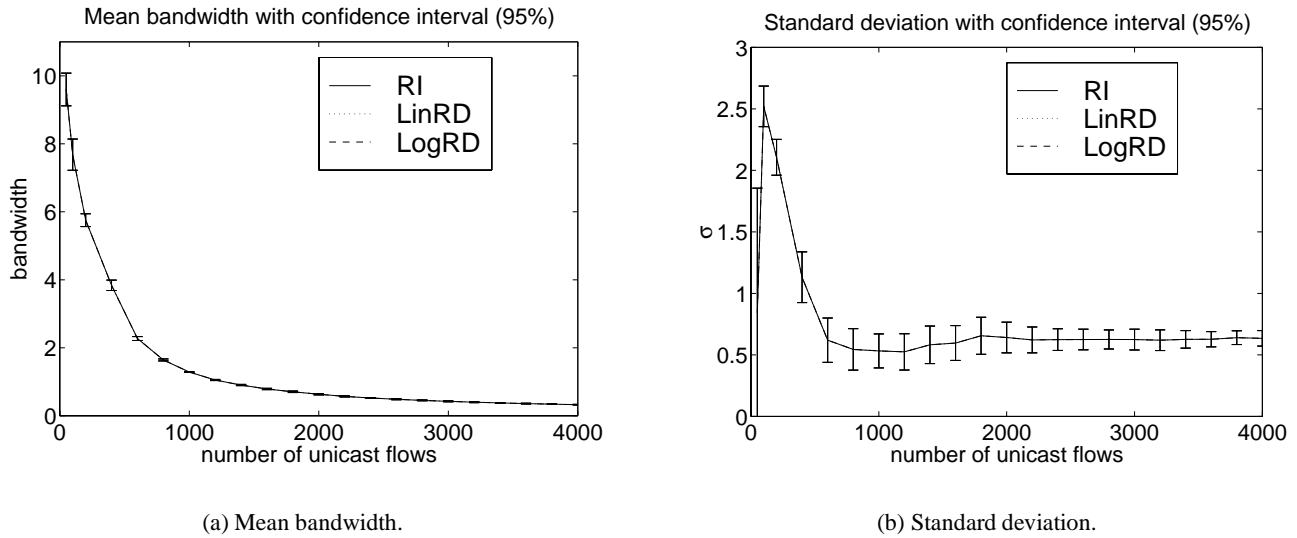


Figure 8: Mean bandwidth (Mbit/s) and standard deviation of all receivers for an increasing number of unicast flows, $k = [50, \dots, 4000]$.

First of all, we see in Fig. 8 that the 3 allocation policies give the same allocation. Indeed, there are only unicast flows and the differences of behavior between the policies depend only on the number of receivers downstream a link for a flow, which is always one in this example.

Secondly, the mean bandwidth (Fig. 8(a)) decreases as the number of unicast flows increases. An added unicast flows decreases the average share. For instance, if we take one link of capacity C shared by all unicast flows, k unicast flows on that link obtain a bandwidth of $\frac{C}{k}$ each.

We plot the standard deviation in Fig. 8(b). For a small number of unicast flows, we have high standard deviation. Since there are few unicast flows with respect to the network size, the random locations of the unicast hosts have a great impact on the bandwidth allocated. The number of LANs in our topology is 180. So, 180 unicast flows lead on average to one receiver per LAN. A number of unicast flows chosen too small for a large network results in links shared only by a small number of flows. Hence, the statistical measure becomes meaningless. When the network is lightly loaded adding one flow can heavily change the bandwidth allocated to other flows, and we observe a large heterogeneity in the bandwidth allocated to the different receivers. On the other hand, for 1800 unicast flows, the mean number of receivers per LAN is 10, so the heterogeneity due to the random distribution of the pairs sender-receiver does not lead to high standard deviation. According to Fig. 8(b), we chose our unicast environment with 2000 unicast flows to obtain a low bias due to the random location of the sender-receiver pairs.

4.2 Simulation Setup

For our simulations we proceed as follows.

- 2000 unicast sources and 2000 unicast receivers are chosen at random locations among the hosts.
- One multicast source and $1, \dots, 6000$ receivers are chosen at random locations. Depending on the experiment, this may be repeated several times to obtain several multicast trees, each with a single source and the same number of receivers.
- We use shortest path routing [6] through the network to connect the 2000 unicast source-receiver pairs and to build the source-receivers multicast tree [9]. As routing metric, the length of the link as generated by `tiers` is used.
- For every network link, the number of flows across that link is calculated. By tracing back the paths from the receivers to the source, the number of receivers downstream is determined for each flow on every link.
- At each link using the information about the number of flows and the number of receivers downstream, the bandwidth for each flow traversing that link is allocated via one of the three strategies: *RI*, *LinRD*, and *LogRD*.
- In order to determine the bandwidth seen by a receiver r , the minimum bandwidth allocated to a flow on all the links along the path from source to receiver is taken as the bandwidth B_p^r seen by r for strategy p (see section 2.3).

The result of the simulation gives the mean bandwidth \bar{B}_p for the three bandwidth allocation strategies. We conduct different experiments with a single and with multiple multicast groups.

4.3 Single Multicast Group

For this experiment, we add one multicast group to the 2000 unicast flows. The size of the multicast group varies from 1 up to 6000 receivers. There are 70 hosts on each LAN and the number of potential senders/receivers is therefore 12600. This experiment shows the impact of the group size on the bandwidth allocated to the receivers under the three allocation strategies. This simulation is repeated five times and averages are taken over the five repetitions.

We simulate small groups sizes ($m = [1, \dots, 100]$), then large groups sizes ($m = [100, \dots, 3000]$), and finally evaluate the asymptotic behavior of our policies ($m = [3000, \dots, 6000]$). The asymptotic case does not aim to model a real scenario, but gives an indication about the behavior of our policies in extreme cases. While 6000 multicast receivers seems a lot compared to the 2000 unicast flows, this case gives a good indication about the robustness of the policies. We display the results with a logarithmic x-axis.

Fig. 9(a) shows that the average user receives more bandwidth when the allocation depends on the number of receivers. A significant difference between the allocation strategies appears for a group size m greater than 100. For small group sizes, unicast flows determine the mean bandwidth due to the high amount of unicast receivers compared to multicast receivers. We claim that receiver-dependent policies increase receiver satisfaction.

A more accurate analysis needs to distinguish between unicast and multicast receivers. Multicast receivers are rewarded with a higher bandwidth than unicast receivers for using multicast as the comparison between Fig. 10(a) and Fig. 10(b) shows. This is not surprising as our policies reward using multicast. Moreover, the increase in bandwidth allocated to multicast receivers leads to a significant decrease of bandwidth available for unicast receivers for the *LinRD* policy, while the decrease of bandwidth is negligible for the *LogRD* policy (Fig. 10(a)) even in the asymptotic case. In conclusion, the *LogRD* policy is the only policy among the three policies that leads to a significant increase




Figure 9: Mean bandwidth (Mbit/s) and standard deviation of all receivers for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

of receiver satisfaction for the average multicast receiver without affecting the receiver satisfaction for the average unicast receiver.

The standard deviation for the average user increases with the size of the multicast group for the receiver-dependent policies (Fig. 9(b)). This unfairness is caused by the difference of the lower bandwidth allocated to the unicast flows compared to the higher bandwidth given to the a multicast flow (Fig. 10(a) and 10(b)). For *LinRD* and *LogRD*, σ tends to flatten for large group sizes, since the multicast receivers determine, due to their large number, the standard deviation. The standard deviation for unicast receivers (Fig. 11(a)) is independent of the multicast group size and of the policies. For a small increasing group size, fairness first becomes worse among multicast receivers, as indicated by the increasing standard deviation in Fig. 11(b), since the sparse multicast receiver setting results in a high heterogeneity of the allocated bandwidth. As the group size increases further, multicast flows are allocated more bandwidth due to an increasing number of receivers downstream. Therefore, the standard deviation decreases with the number of receivers. In the asymptotic part, the standard deviation for the *LinRD* policy decreases faster than for the *LogRD* policy since as the number of receivers increases, the amount of bandwidth allocated to the multicast flow approaches the maximum bandwidth (the bandwidth of a LAN), see Fig. 10(b). Therefore, all the receivers see a high bandwidth near the maximum, which leads to low standard deviation. Another interesting observation is that the multicast receivers among each other have a higher heterogeneity in the bandwidth received than have the unicast receivers, compare Fig. 11(a) and Fig. 11(b). A few bottlenecks are sufficient to split the multicast receivers in large subgroups with significant differences in bandwidth allocation that subsequently result in a higher standard deviation. For the 2000 unicast receivers, the same bottlenecks affect only a few receivers.

The standard deviation taken over all the receivers hides the worst case performance experienced by any individual receiver. To complete our study, we measure the minimum bandwidth, which gives an indication about the worst

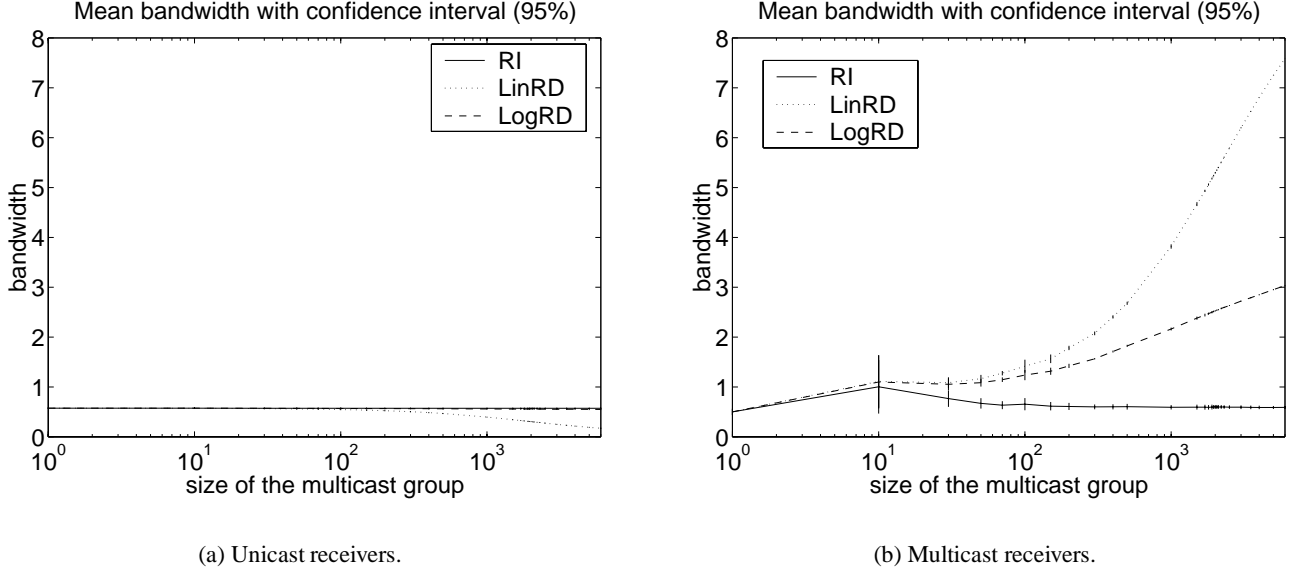


Figure 10: Mean bandwidth (Mbit/s) of unicast and multicast receivers with confidence interval (95%) for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

case behavior seen by any receiver. The minimum bandwidth over all the receivers is dictated by the minimum bandwidth over the *unicast* receivers (we give only one plot, Fig. 12(a)). As the size of the multicast group increases, the minimum bandwidth seen by the unicast receivers dramatically decreases for the *LinRD* policy, whereas the minimum bandwidth for the *LogRD* policy remains close to the one for the *RI* policy even in the asymptotic part of the curve. We can point out another interesting result: the minimum bandwidth for the *RI* policy stays constant even for very large group sizes; the *LinRD* policy that simulates the bandwidth that would be allocated if we replace the multicast flow by an equivalent number of unicast flows, results in a minimum bandwidth that rapidly decreases toward zero. Therefore, we note the positive impact of multicast on the bandwidth allocated, and multicast *greatly improves the worst case bandwidth allocation*. We see in Fig. 12(b) that the minimum bandwidth increases for *multicast* receivers with the size of the multicast group for the receiver-dependent policies. In conclusion, the *LinRD* policy leads to important degradation of the fairness when the multicast group size increases, whereas the *LogRD* policy always remains close to the *RI* policy.

For the *RI* policy, we see that the increase in the multicast group size does not influence the average user satisfaction (Fig. 9(a)), nor the fairness among different receivers (Fig. 9(b)). Also, the difference between unicast and multicast receivers is minor concerning the bandwidth both received (Fig. 10(a) and 10(b)), and the unfairness (Fig. 11(a) and 11(b)). The *LogRD* policy is the only policy among our policies that significantly increases receiver satisfaction (Fig. 9(a)), keeps fairness close to the one of the *RI* policy (Fig. 9(b)), and does not starve unicast flows, even in asymptotic cases (Fig. 12(a)).

Finally, one also should note the similarity between Fig. 9(a), 9(b) obtained by simulation for a large network and Fig. 3(b), 4 obtained by analysis of the star topology. This suggests that the star topology is a good model to study the impact of the three different bandwidth allocation policies.





Figure 11: Standard deviation of unicast and multicast receivers with confidence interval (95%) for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

4.4 Multiple Multicast Groups


We now consider the case of multiple multicast groups and 2000 unicast sessions. We add to the 2000 unicast sessions multicast sessions of 100 receivers each. The number of multicast sessions ranges from 2 to 100. There are 100 hosts on each LAN, the number of potential receivers/senders is therefore 18000. The simulations were repeated five times and average are taken over the five repetitions.

In this section, each plot can be partitioned into three parts: the first part shows the results for a small number of multicast receivers with respect to the number of unicast receivers ($M = [1, \dots, 10]$ groups), the second part shows the results for a large number of multicast receivers compared to the number of unicast receivers ($M = [10, \dots, 50]$ groups), and the third part indicates the asymptotic behavior of our policies ($M = [50, \dots, 100]$ groups). Again, the asymptotic case gives an indication about the behavior of our policies in extreme cases and about the robustness of our policies.

The three policies give nearly the same mean bandwidth over all the receivers (Fig. 13(a)). The *LogRD* policy is the best policy for the mean bandwidth over all the receivers. We can explain this behavior with our simple analytical study. We see for the chain topology that there are some cases where the *LinRD* strategy gives worse results than the *LogRD* and the *RI* strategy. We can consider a real network as a composition of stars and chains, therefore, it is not surprising to observe, for a large topology, a composition of the behavior of both the star and chain topology. We see that the bandwidth allocation of the *LogRD* policy over the *RI* policy first slightly increases as the number of multicast groups increases (until $M = 10$), and then decreases with the number of multicast groups. For $M \leq 10$, the number of multicast receivers that benefits from the receiver-dependent policies increases and so the differences between receiver-dependent and receiver-independent policies increase. However, in the second part of the curves ($M > 10$), the number of multicast sessions tends to have more impact than the number of multicast receivers. Indeed,



(a) Minimum bandwidth of unicast receivers.



(b) Minimum bandwidth of multicast receivers.

Figure 12: Minimum bandwidth (Mbit/s) with confidence interval (95%) of the unicast receivers and of the multicast receivers for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

when the number of multicast session increases we have two behaviors: i) As the number of sessions (unicast or multicast) increases the bandwidth available for each session decreases, and therefore, the benefits due to receiver-dependent policies decreases; ii) The receiver-dependent policies reward multicast flows as a function of the number of receivers. But, if all the flows have the same number of receivers, receiver-dependent policies do not make any significant difference. Fig. 14(a) shows that the *LogRD* policy gives roughly the same bandwidth than the *RI* policy for unicast receivers whereas the *LinRD* policy leads to a lower bandwidth. Fig. 14(b) shows a very important result, the receiver-dependent policies significantly reward the multicast receivers compared to the *RI* policy. As the number of multicast groups increases, the differences between the policies decrease, since the number of multicast sessions tends to have more impact on the mean bandwidth than the number of multicast receivers. Fig. 14(b) shows that the receiver dependent policies achieve their objective, which is to reward multicast flows.

Fig. 13(b) shows that standard deviation is roughly the same for the three bandwidth allocation policies. Fig. 15(b) shows that the multicast receivers have higher standard deviation with the receiver-dependent policies than with *RI*. The standard deviation is roughly the same for the three bandwidth allocation policies for the unicast receivers (Fig. 15(a)). As the number of multicast sessions increases, multicast flows dominate due to the high amount of multicast receivers compared to unicast receivers, and therefore, the standard deviation of multicast receivers for the three bandwidth allocation strategies becomes close due to the high homogeneity of the sessions. The minimum bandwidth is dictated by the unicast receivers, so the plots for all the receivers and for the unicast receivers are the same. Fig. 16(a) shows an interesting result. The *LinRD* policy gives very little bandwidth to unicast receivers, whereas the *LogRD* policy allocates roughly the same minimum bandwidth than the *RI* policy. Fig. 16(b) shows the minimum bandwidth for multicast receivers is slightly better for the receiver-dependent policy than for *RI* for a small number of multicast sessions, and the minimum bandwidth is slightly worse for a large number of multicast sessions. Indeed, for a small

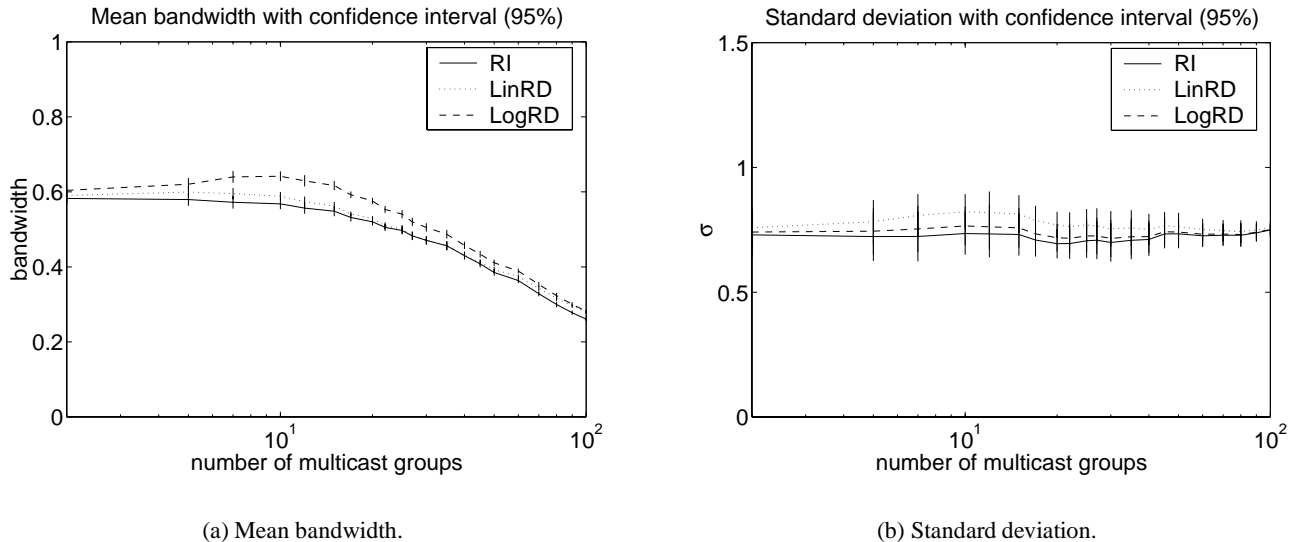


Figure 13: Mean bandwidth (Mbit/s) and standard deviation of all the receivers for an increasing number of multicast sessions, $k = 2000$, $M = [2, \dots, 100]$, $m = 100$.

number of multicast sessions the interaction between sessions is low, therefore the probability that a multicast session decreases the bandwidth seen by a multicast receiver of another session is low. But, for a large number of multicast sessions, the interaction between multicast sessions is high, and the probability that a multicast session decreases the bandwidth seen by a multicast receiver of another session is higher.

We did another experiment that aims to model small conferencing groups where multicast groups of a size 20 are added. But the results of this experiment do not differ from the results of the experiment with multicast group sizes of 100 receivers and we do not present these results.

In conclusion, the receiver satisfaction and fairness of all the receivers are roughly the same for the three bandwidth allocation strategies (Fig. 13), but the *LogRD* policy is the only policy that greatly improves the average bandwidth allocated to multicast receivers (Fig. 14(b)) without starving unicast flows (Fig. 16(a)).

5 Practical Aspects

5.1 Estimating the Number of Downstream Receivers

Up to now, we quantified the advantages of using bandwidth allocation strategies based on the number of downstream receivers. Estimating the number of receivers downstream of a network node has a certain cost but has other benefits that largely outweigh this cost. Two examples of these benefits are feedback accumulation and multicast charging.

One of the important points of the feedback accumulation process is the estimation of the number of downstream receivers. Given the number of receivers is known in the network nodes, the distributed process of feedback accumulation [23], or feedback filtering in network nodes becomes possible and has a condition to terminate upon.

While multicast saves bandwidth, it is currently not widely offered by network operators due to the lack of a valid

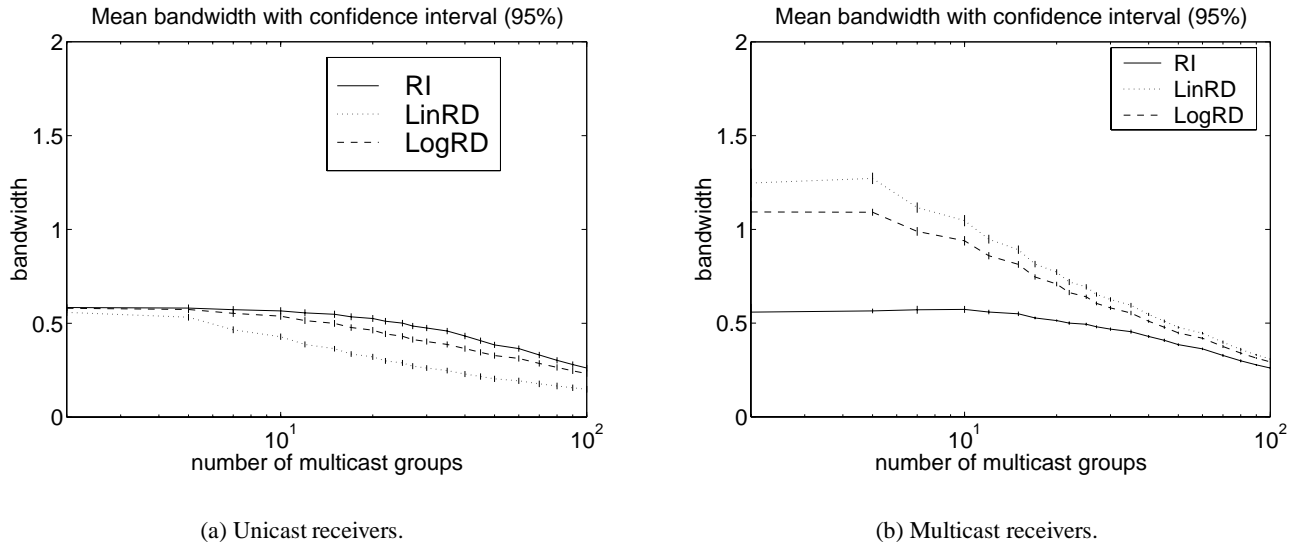


Figure 14: Mean bandwidth (Mbit/s) of unicast and multicast receivers with confidence interval (95%) for an increasing number of multicast sessions, $k = 2000$, $M = [2, \dots, 100]$, $m = 100$.

charging model [5, 15]. By knowing the number of receivers at the network nodes, different charging models for multicast can be applied, including charging models that use the number of receivers. In the case of a single source and multiple receivers, the amount of resources used with multicast depends on the number of receivers. For an ISP, in order to charge the source according to the resources consumed, the number of receivers is needed. The bandwidth allocation policy used impacts the charging in the sense that the allocation policy changes the number of resources consumed by a multicast flow, and changes the cost of a multicast flow for the ISP. However, in appendix B, we see that a simple local bandwidth allocation policy leads to a global cost that is a complex function of the number of receivers. It is not clear to us whether an ISP can charge a multicast flow with a simple linear or logarithmic function of the number of receivers. Moreover, several ISPs (see [8]) use flat rate pricing for multicast due to the lack of valid charging model. Even in the case of flat rate pricing, the number of downstream receivers is useful when a multicast tree spans multiple ISPs. In this case, we have a means to identify the number of receivers in each ISP. The charging issue is orthogonal to our paper and is an important area for future research.

The estimation of the number of downstream receivers is feasible, for instance, with the Express multicast routing protocol [15]. The cost of estimating the number of downstream receivers is highly dependent on the method used and the accuracy of the estimate required. As our policy is based on a logarithmic function, we only need a coarse estimate of the number of downstream receivers. Holbrook [15] describes a low overhead method for the estimation of the number of downstream receivers.

5.2 Introduction of the LogRD Policy

Another important question is how to introduce the *LogRD* policy in a real network without starving unicast flows. In section 4, we show that even in asymptotic cases the *LogRD* strategy does not starve unicast flows, but we do not have

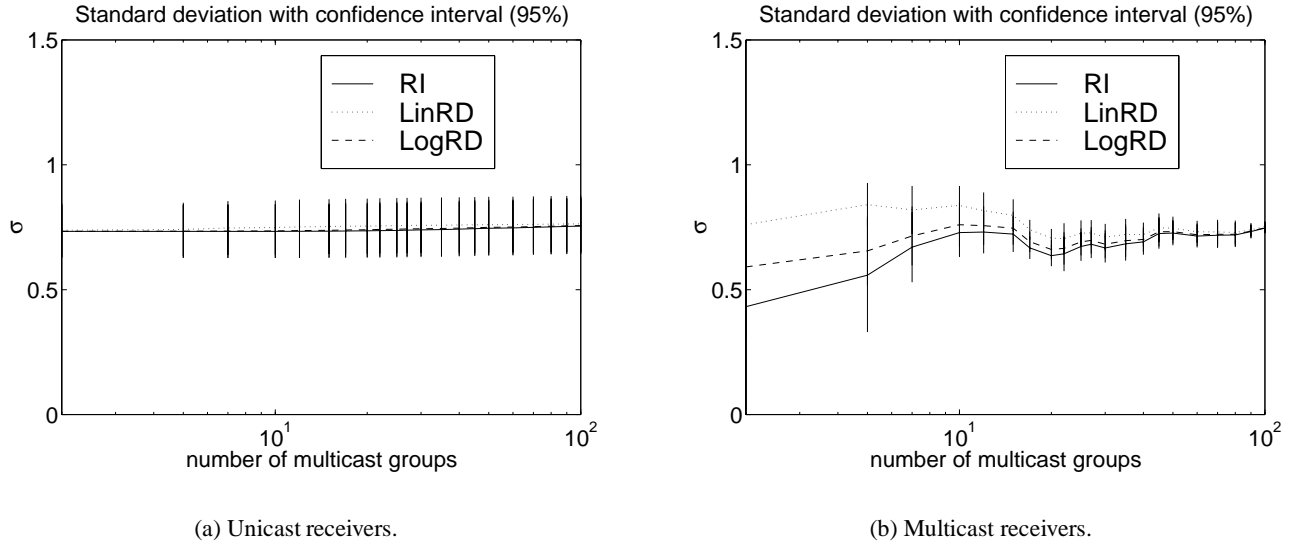


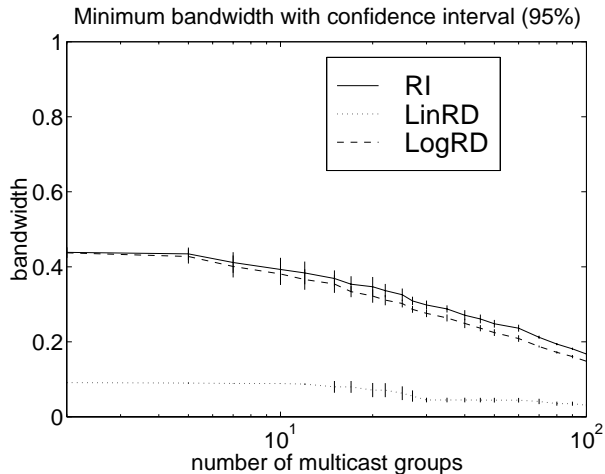
Figure 15: Standard deviation of unicast and multicast receivers with confidence interval (95%) for an increasing number of multicast sessions, $k = 2000$, $M = [2, \dots, 100]$, $m = 100$.

a hard guarantee about the bandwidth allocated to unicast receivers. For instance, one multicast flow with 1 million downstream receivers sharing the same bottleneck than a unicast flow will grab 93% of the available bandwidth. This is a large amount of the bandwidth, but that does not lead to a starvation of the unicast flow.


The *LogRD* policy will asymptotically – when the number of multicast receivers tends toward infinity – lead to an optimal receiver satisfaction (limited by the capacity of the network) and to a low fairness. In particular, the multicast flow will grab all the available bandwidth of the bottleneck link and starve all the unicast flows sharing this bottleneck link. It is possible to devise a strategy based on the *LogRD* policy that allocates to the multicast flows never more than K times the bandwidth allocated to the unicast flows sharing the same bottleneck. We can imagine the *LogRD* strategy to be used in a hierarchical link sharing scheme (see [14, 1] for hierarchical link sharing models). The idea is to introduce our policy in the general scheduler [14] (for instance we can configure the weight of a PGPS [22] scheduler with the *LogRD* policy to achieve our goal), and to add an administrative constraint in the link sharing scheduler (for instance we guarantee that unicast traffic receives at least $x\%$ of the link bandwidth). This is a simple way to allocate the bandwidth with respect to the *LogRD* policy, and to guarantee a minimum bandwidth for the unicast flows. Moreover, Kumar et al. [18] show that it is possible to integrate efficiently a mechanism like HWFQ [1] in a Gigabit router, and WFQ is already available in the recent routers [4].

5.3 Incremental Deployment

An important practical aspect is whether it is possible to incrementally deploy the *LogRD* policy. To answer this question we make the following experiment. We consider the random topology used in section 4 and a unicast environment consisting of 2000 unicast flows. We add to this unicast environment 20 multicast flows with a uniform group size of 50 multicast receivers randomly distributed. The simulation consists in varying the percentage of LANs,



(a) Minimum bandwidth of unicast receivers.



(b) Minimum bandwidth of multicast receivers.

Figure 16: Minimum bandwidth (Mbit/s) with confidence interval (95%) of the unicast receivers and of the multicast receivers for an increasing number of multicast sessions, $k = 2000$, $M = [2, \dots, 100]$, $m = 100$.

MANs, and WANs that use the *LogRD* policy compared to the *RI* policy. We make the assumption that each LAN, MAN, and WAN is an autonomous system managed by a single organization. So when an organization decides to use the *LogRD* policy, it changes the policy in all the routers of the LAN, MAN, or WAN it is responsible for. We say that a LAN, MAN or WAN is *LogRD* if all the routers use the *LogRD* policy. The simulation consists in varying the number of *LogRD* LANs and MANs from 0% to 100%, for the WAN we only look at a full support (all routers are *LogRD*) or no support (all routers are *RI*). We call these percentages respectively *perLAN*, *perMAN*, and *perWAN*. This simulation is repeated five times and averages are taken over the five repetitions. The results are given with a confidence interval of $95\% \pm 20\text{Kbit/s}$ around the mean bandwidth.

The main behavior we see in Fig. 17 is the *interdependency* of the parameters *perLAN*, *perMAN*, and *perWAN* on the mean bandwidth for the multicast receivers. An isolated deployment of the *LogRD* in just the LANs, MANs, or WANs does not allow to achieve a mean bandwidth close to the mean bandwidth obtained when the whole network is *LogRD*. For instance, the *perMAN* parameter does not have a significant influence on the mean bandwidth when *perLAN* = 0. However, when *perLAN* = 100 and *perWAN* = 100, the *perMAN* parameter has a significant influence on the mean bandwidth. The results obtained depend on the network configuration (number of LANs, MANs, and WANs, link bandwidth, etc.). However, we believe the property of interdependency of the parameters *perLAN*, *perMAN*, and *perWAN* to hold in all the cases.

In conclusion, to reap the full benefit of the *LogRD* policy, a coordinated deployment is necessary. However, as the lack of links using the *LogRD* allocation does not lead to any performance degradation for the network, an incremental deployment is possible.

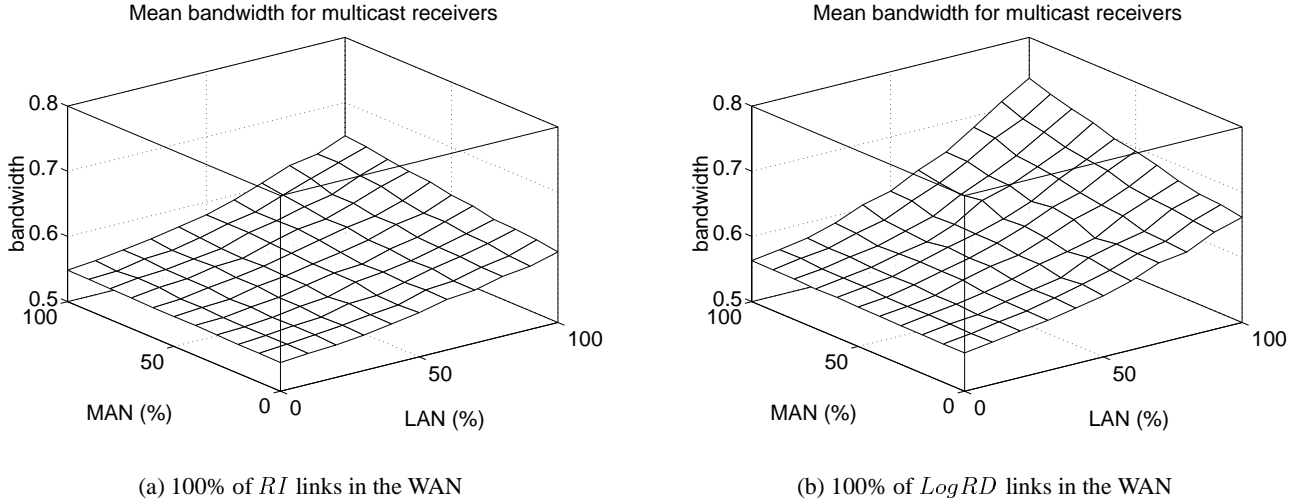


Figure 17: Influence on the mean bandwidth (Mbit/s) for the multicast receivers for an hierarchical incremental deployment of the *LogRD* policy, $k = 2000$, $M = 20$, $m = 50$.

6 Conclusion

If one wants to introduce multicast in the Internet, one should give an incentive to use it. We propose a simple mechanism that takes into account the number of receivers downstream. Our proposal does not starve unicast flows and greatly increases multicast receiver satisfaction.

We defined three different bandwidth allocation strategies as well as criteria to compare these strategies. We compared the three strategies analytically and through simulations. Analytically, we studied two simple topologies: a star, and a chain. We showed that the *LogRD* policy leads to the best tradeoff between receiver satisfaction and fairness. The striking similarities in the results for the analytical study and the simulations confirm that we had chosen valid models.

To simulate real networks, we defined a large topology consisting of WANs, MANs, and LANs. In a first round of experiments, we determined the right number of unicast receivers. We studied the introduction of multicast in a unicast environment with three different bandwidth allocation policies. The aim was to understand the impact of multicast in the real Internet. We showed that allocating link bandwidth dependent on the flows' number of downstream receivers results in a higher receiver satisfaction. The *LogRD* policy provides the best tradeoff between the receiver satisfaction and the fairness among receivers. Indeed, the *LogRD* policy always leads to higher receiver satisfaction than the *RI* policy for roughly the same fairness, whereas the *LinRD* policy leads to higher receiver satisfaction than the *LogRD* policy, however, at the expense of unacceptable decrease in fairness.

Our contribution in this paper is the definition and evaluation of a new bandwidth allocation policy called *LogRD* that gives a real incentive to use multicast. Also, the *logRD* policy gives a relevant answer to the open question on how to treat a multicast flow compared to a unicast flow sharing the same bottleneck. To the best of our knowledge, we are the first that take into account the number of multicast receivers to reward multicast flows. Moreover, we show that the deployment of the *LogRD* policy is feasible when deployed per ISP at the same time as the ISP upgrades its

network to be multicast capable.

Acknowledgment

Many thanks to Jim Kurose for sharing his insights about the notions of receiver satisfaction and fairness. We also thanks the anonymous reviewers for their helpful comments. Eurecom's research is partially supported by its industrial partners: Ascom, Cegetel, France Telecom, Hitachi, IBM France, Motorola, Swisscom, Texas Instruments, and Thomson CSF.

References

- [1] J. C. Bennett and H. Zhang, "Hierarchical Packet Fair Queueing Algorithms", *IEEE/ACM Transactions on Networking*, 5(5):675–689, October 1997.
- [2] K. Bharat-Kumar and J. Jaffe, "A new Approach to Performance-Oriented Flow Control", *IEEE Transactions on Communications*, 29(4):427–435, 1981.
- [3] K. Calvert, M. Doar, and E. W. Zegura, "Modeling Internet Topology", *IEEE Communications Magazine*, 35(6):160–163, June 1997.
- [4] Cisco, "Advanced QoS Services for the Intelligent Internet", White Paper, May 1997.
- [5] R. Comerford, "State of the Internet: Roundtable 4.0", *IEEE Spectrum*, 35(10):69–79, October 1998.
- [6] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, The MIT Press, 1990.
- [7] S. E. Deering, "Multicast routing in internetworks and extended LANs", In *Proc. ACM SIGCOMM 88*, pp. 55–64, Stanford, CA, August 1988.
- [8] C. Diot, B. N. Levine, B. Lyles, H. Kassem, and D. Balensiefen, "Deployment Issues for the IP Multicast Service and Architecture", *IEEE Network magazine special issue on Multicasting*, 14(1):78–88, January/February 2000.
- [9] M. Doar and I. Leslie, "How Bad is Naïve Multicast Routing", In *Proceedings of IEEE INFOCOM'93*, volume 1, pp. 82–89, 1993.
- [10] M. B. Doar, "A Better Model for Generating Test Networks", In *Proceedings of IEEE Global Internet*, pp. 86–93, London, UK, November 1996, IEEE.
- [11] H. Eriksson, "MBONE: The Multicast Backbone", *Communications of the ACM*, 37(8):54–60, August 1994.
- [12] A. Feldman, *Welfare economics and social choice theory*, Martinus Nijhoff Publishing, Boston, 1980.
- [13] S. Floyd, "Connections with Multiple Congested Gateways in Packet-Switched Networks Part 1:One-way Traffic", *Computer Communications Review*, 21(5):30–47, October 1991.
- [14] S. Floyd and V. Jacobson, "Link-sharing and Resource Management Models for Packet Networks", *IEEE/ACM Transactions on Networking*, 3(4):365–386, August 1995.

- [15] H. W. Holbrook and D. R. Cheriton, “IP Multicast Channels: EXPRESS Support for Large-scale Single-source Applications”, In *Proc. of ACM SIGCOMM’99*, pp. 65–78, Harvard, Massachusetts, USA, September 1999.
- [16] R. Jain, D. M. Chiu, and W. Hawe, “A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Computer Systems”, Technical report 301, DEC, Littleton, MA, September 1984.
- [17] T. Jiang, M. H. Ammar, and E. W. Zegura, “Inter-Receiver Fairness: A Novel Performance Measure for Multicast ABR Sessions”, In *Proc. of ACM Sigmetrics*, pp. 202–211, June 1998.
- [18] V. P. Kumar, T. V. Lakshman, and D. Stiliadis, “Beyond Best Effort: Router Architectures for the Differentiated Services of Tomorrow’s Internet”, *IEEE Communications Magazine*, 36(5):152–164, May 1998.
- [19] A. Legout and E. W. Biersack, “PLM: Fast Convergence for Cumulative Layered Multicast Transmission Schemes”, In *Proc. of ACM SIGMETRICS’2000*, pp. 13–22, Santa Clara, CA, USA, June 2000.
- [20] S. McCanne, V. Jacobson, and M. Vetterli, “Receiver-driven Layered Multicast”, In *SIGCOMM 96*, pp. 117–130, August 1996.
- [21] J. Nonnenmacher and E. Biersack, “Asynchronous Multicast Push: AMP”, In *Proceedings of ICC’97*, pp. 419–430, Cannes, France, November 1997.
- [22] A. K. Parekh and R. G. Gallager, “A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks”, In *Proc. IEEE INFOCOM’93*, pp. 521–530, 1993.
- [23] S. Paul, K. K. Sabnani, J. C. Lin, and S. Bhattacharyya, “Reliable Multicast Transport Protocol (RMTP)”, *IEEE Journal on Selected Areas in Communications, special issue on Network Support for Multipoint Communication*, 15(3):407 – 421, April 1997.
- [24] G. Phillips, S. Shenker, and H. Tangmunarunkit, “Scaling of Multicast Trees: Comments on the Chuang-Sirbu Scaling Law”, In *Proc. of ACM SIGCOMM’99*, pp. 41–51, Harvard, Massachusetts, USA, September 1999.
- [25] E. W. Zegura, K. Calvert, and S. Bhattacharjee, “How to Model an Internetwork”, In *Infocom ’96*, pp. 594–602, March 1996.
- [26] E. W. Zegura, K. Calvert, and M. J. Donahoo, “A Quantitative Comparison of Graph-based Models for Internet Topology”, *IEEE/ACM Transactions on Networking*, 5(6):770–783, December 1997.

A Discussion on Multicast Gain

To evaluate the bandwidth multicast gain, we restrict ourselves to the case of a full o -ary tree with receivers at the leaves – in this case we model a point to point network – or with broadcast LAN at the leaves. We consider one case where the unicast and the multicast cost only depends on the number of links (the unlimited bandwidth case) and one case where the unicast and the multicast cost depends on the bandwidth used (the limited bandwidth case).

Let the full o -ary tree be of height h . We assume the sender to be at the root, so there are $R = o^h$ receivers or $N = o^h$ LANs with R_N receivers on each LAN ($R = R_N \cdot N$). We define the **bandwidth cost** as the sum of all the bandwidths consumed on all the links of the tree. We define the **link cost** as the sum of all the links used on the tree,

we count n times the same link when the same data are sent n times on this link. Let C_U be the unicast bandwidth/link cost from the sender to all of the receivers and C_M the multicast bandwidth/link cost from the same sender to the same receivers.

A.1 Bandwidth-Unlimited Case

We assume that every link of the tree has unlimited bandwidth. Let C_U and C_M be the link cost for unicast and multicast, respectively.

If we consider one receiver on each leaf of the tree we have:

$$C_U = o^h + o^{h-1} \cdot o + \dots + o^1 \cdot o^{h-1} = h \cdot o^h = h \cdot R = R \cdot \log_o(R) \quad (7)$$

$$C_M = \sum_{i=1}^h o^i = \frac{o^{h+1} - o}{o - 1} = \frac{o}{o - 1}(R - 1)$$

We define the multicast gain as the ratio:

$$\frac{C_U}{C_M} = \log_o(R) \frac{R}{R - 1} \cdot \frac{o - 1}{o}$$

The multicast gain depends logarithmically on the number of receivers.

If we consider one LAN on each leaf of the tree we have:

$$C_U = h \cdot R = h \cdot N \cdot R_N = R_N \cdot N \cdot \log_o(N)$$

$$C_M = \sum_{i=1}^h o^i = \frac{o^{h+1} - o}{o - 1} = \frac{o}{o - 1}(N - 1)$$

We define the multicast gain as the ratio:

$$\frac{C_U}{C_M} = \frac{o - 1}{o} \cdot R_N \cdot \frac{1}{1 - \frac{1}{N}} \cdot \log_o(N)$$

The gain depends logarithmically on the number of LANs and linearly on the number of receivers per LAN.

A.2 Bandwidth-Limited Case

Every link of the tree has a capacity C . Let C_U and C_M be the bandwidth cost for unicast and multicast, respectively.

If we consider one receiver on each leaf of the tree we have:

$$\begin{aligned} C_U &= o \cdot C + o^2 \cdot \frac{C}{o} + o^3 \cdot \frac{C}{o^2} + \dots + o^h \cdot \frac{C}{o^{h-1}} \\ &= \sum_{i=1}^h C \cdot o = h \cdot C \cdot o = C \cdot o \cdot \log_o(R) \\ C_M &= C \sum_{i=1}^h o^i = C \cdot \frac{o^{h+1} - o}{o - 1} = C \cdot \frac{o}{o - 1}(R - 1) \end{aligned}$$

The multicast gain is:

$$\frac{C_U}{C_M} = (o - 1) \frac{\log_o(R)}{R - 1}$$

This means that there is a multicast gain smaller than 1 for large R. But, of course, in the unicast case (which is now globally less expensive), we also have much smaller receiver satisfaction due to the bandwidth-limited links close to the source. Therefore, the definition for the standard multicast gain does not make sense in the bandwidth-limited case. For the unlimited case, receivers are equally satisfied, since they receive the same bandwidth and the multicast gain *makes sense*.

We need to define another measure that combines the satisfaction and the cost. We use cost per satisfaction. We look at the ratio of bandwidth cost per satisfaction that tells us how much bandwidth we need to invest to get a unit of satisfaction.

We now employ: $GB = \frac{\text{global cost}}{\text{global satisfaction}}$. To compute the global satisfaction, we add the satisfaction over all receivers. Let the global satisfaction be S_U for unicast and S_M for multicast.

$$\begin{aligned} S_U &= R \cdot C \cdot \frac{1}{o^{h-1}} = R \cdot C \cdot \frac{o}{o^h} = R \cdot C \cdot \frac{o}{R} = C \cdot o \\ S_M &= R \cdot C \end{aligned}$$

Then $GB = \frac{\text{global cost}}{\text{global satisfaction}}$ is:

$$GB_U = \frac{C_U}{S_U} = \frac{C \cdot o \cdot \log_o(R)}{C \cdot o} = \log_o(R)$$

$$GB_M = \frac{C_M}{S_M} = \frac{(R - 1)}{R} \cdot \frac{o}{o - 1}$$

Now the new multicast gain is:

$$\frac{GB_U}{GB_M} = \frac{o - 1}{o} \cdot \frac{R}{R - 1} \cdot \log_o(R)$$

The gain depends logarithmically on the number of receivers.

If we consider one LAN on each leaf of the multicast tree we have:

$$\begin{aligned} C_U &= o \cdot C + o^2 \cdot \frac{C}{o} + o^3 \cdot \frac{C}{o^2} + \dots + o^h \cdot \frac{C}{o^{h-1}} = C \cdot o \cdot \log_o(N) \\ C_M &= C \sum_{i=1}^h o^i = C \cdot \frac{o^{h+1} - o}{o - 1} = C \cdot \frac{o}{o - 1} (N - 1) \end{aligned}$$

The multicast gain is:

$$\frac{C_U}{C_M} = (o - 1) \frac{\log_o(N)}{N - 1}$$

Once again the multicast gain smaller than 1 for large N. The global satisfaction is:

$$S_U = R \cdot C \cdot \frac{1}{o^{h-1} R_N} = C \cdot o$$

$$S_M = R \cdot C$$

Then $GB = \frac{\text{global cost}}{\text{global satisfaction}}$ is :

$$GB_U = \frac{C_U}{S_U} = \log_o(N)$$

$$GB_M = \frac{C_M}{S_M} = \frac{N-1}{R_N \cdot N} \cdot \frac{o}{o-1}$$

Now the new multicast gain is:

$$\frac{GB_U}{GB_M} = \frac{o-1}{o} \cdot \frac{R_N \cdot N}{N-1} \cdot \log_o(N)$$

The gain depends logarithmically on the number of LANs and linearly on the number of receivers per LAN.

In conclusion, for both the unlimited and the limited bandwidth case, the multicast gain has a logarithmic trend with the number of receivers in case of point-to-point networks. For broadcast LANs at the leaves of the multicast distribution tree, the multicast gain has a logarithmic trend with the number of LANs, but a linear trend with the number of receivers per LAN. Therefore, with a small number of receivers per LAN the multicast gain is logarithmic but with a large number of receivers per LANs the multicast gain is linear.

B Global Impact of a Local Bandwidth Allocation Policy

We consider a full o -ary tree for the unlimited bandwidth case when there is a receiver per leaf. The unicast link cost is $C_U = h \cdot R$ (see Eq. 7). Now we consider the multicast link cost for the RI , the $LinRD$, and $LogRD$ policies. For instance when there are 2 receivers downstream of link l , the $LinRD$ policy allocates the equivalent of 2 units of bandwidth and the $LogRD$ policy allocates the equivalent of $1 + \ln(2)$ units of bandwidth compared to the RI policy which allocates 1 unit of bandwidth.

The multicast link cost for the RI policy is:

$$C_M^{RI} = \sum_{i=1}^h o^i = \frac{o}{o-1}(R-1)$$

The multicast link cost for the $LinRD$ policy is:

$$C_M^{LinRD} = o \cdot \frac{R}{o} + o^2 \cdot \frac{R}{o^2} + \dots + o^h \cdot \frac{R}{o^h} = h \cdot R = C_U$$

The multicast link cost for the $LogRD$ policy is:

$$C_M^{LogRD} = o \cdot (1 + \ln \frac{R}{o}) + o^2 \cdot (1 + \ln \frac{R}{o^2}) + \dots + o^h \cdot (1 + \ln \frac{R}{o^h}) = \sum_{i=1}^h o^i (1 + \ln \frac{R}{o^i})$$

We have $1 + \ln \frac{R}{o^i} \leq \frac{R}{o^i}$ and $1 + \ln \frac{R}{o^i} < \frac{R}{o^i}$ for $\frac{R}{o^i} \neq 1$. So for $h > 1$ and $o > 1$ we have $C_M^{LogRD} < C_M^{LinRD}$.

In conclusion we see that the policy that rewards multicast with its gain is the $LinRD$ policy and not the $LogRD$ policy as expected.

C Tiers Setup

We give a brief description of the topology used for all the simulations. The random topology RT is generated with `tiers v1.1` using the command line parameters `tiers 1 20 9 5 2 1 3 1 1 1 1`. A WAN consists of 5 nodes and 6 links and connects 20 MANs, each consisting of 2 nodes and 2 links. To each MAN, 9 LANs are connected. Therefore, the core topology consists of $5 + 40 + 20 \cdot 9 = 225$ nodes. The capacity of WAN links is 155Mbit/s, the capacity of MAN links is 55Mbit/s, and the capacity of LAN links is 10Mbit/s.

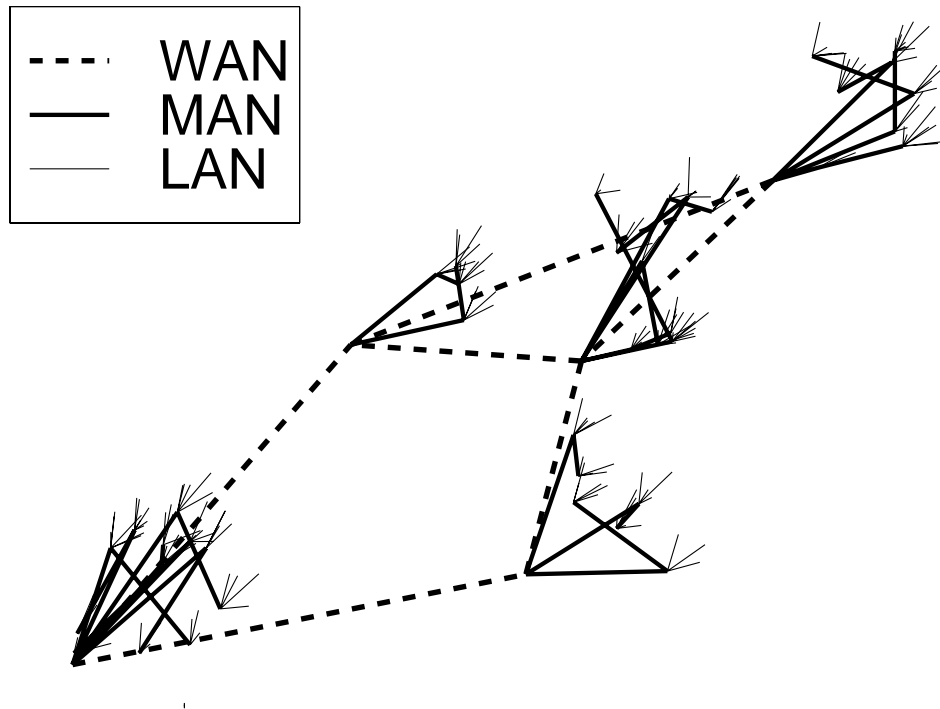


Figure 18: The random topology RT

Each LAN is represented as a single node and connects several hosts via a 10Mbit/s link. The number of hosts connected to a LAN changes from experiment to experiment to speed up simulation. However, the number of hosts is always chosen larger than the sum of the receivers and the sources all together.