

Bandwidth Allocation Policies for Unicast and Multicast Flows

A. Legout, J. Nonnenmacher and E. W. Biersack

Institut EURECOM

B.P. 193, 06904 Sophia Antipolis, FRANCE

{legout,nonnen,erbi}@eurecom.fr

Abstract—Using multicast delivery to multiple receivers reduces the aggregate bandwidth required from the network compared to using unicast delivery to each receiver.

To encourage the use of multicast delivery, a higher amount of bandwidth should be allocated to a multicast flow as compared to a unicast flow that share the same bottleneck, but without starving the unicast flow. We investigate three bandwidth allocation policies for multicast flows and evaluate their impact on the bandwidth received by the individual receivers.

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 maximizing the receiver satisfaction and keeping fairness high.¹

Keywords—Unicast, Multicast, Bandwidth Allocation, Quality of Service

I INTRODUCTION

There is an increasing number of applications such as software distribution, audio/video conferences, and audio/video broadcasts where data sent by the source is destined to multiple receivers. During the last decade, multicast routing and multicast delivery have evolved from being a pure research topic [1] to being experimentally deployed in the MBONE [2] to being supported by major router manufacturers. 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 links that lead 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 link the packet traverses from the sender to the R receivers for a particular transmission (unicast or multicast). For shortest path routing between source and receivers for unicast and multicast, the multicast gain for the model of a full o -ary multicast tree is:

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

¹ Copyright 1999 IEEE. Published in the Proceedings of INFOCOM'99, 21st - 25th March 1999, New York, USA. Personal use of this material is permitted. However, permission to reprint/republish this material for advertising or promotional purposes or for creating new collective works for resale or redistribution to servers or other works, must be obtained from the IEEE. Contact: Manager, Copyrights and Permissions / IEEE Service Center / 445 Hoes Lane / P.O. Box 1331 / Piscataway, NJ 08855-1331, USA. Telephone: + Intl. 908-562-3966.

Even for random networks and multicast trees different than the idealized full o -ary tree, the multicast gain is largely determined by the logarithm of the number of receivers [3].

Despite the widespread deployment of multicast capable networks, a multicast service is rarely provided and network providers keep the multicast delivery option in their routers turned off. Several reasons contribute to the unavailability of multicast. A major reason is the lack of congestion control, and the fear that multicast traffic grabs the available network bandwidth and leaves only little bandwidth to unicast traffic.

Unicast is a one-to-one communication, multicast is a one-to-many communication, and broadcast is a one-to-all communication. Therefore, unicast and broadcast can be treated as special cases of multicast, with one receiver, or all receivers, respectively. A valid bandwidth allocation policy employed for multicast should therefore also work in the extreme cases of unicast traffic or broadcast traffic.

We want to give an incentive to use multicast by rewarding the multicast gain in the network to the receivers at the edge of the network; at the same time we want to treat unicast traffic fair relative to multicast traffic.

We investigate bandwidth allocation policies that allocate the bandwidth locally at each single link between unicast and multicast traffic and evaluate globally the bandwidth perceived by the receivers.

For different bandwidth allocation policies, we examine the case where a unicast network (like the Internet) is augmented with a multicast delivery service and we evaluate the receiver satisfaction and the fairness among receivers.

The rest of the paper is organized as follows. In Section II we present the three bandwidth allocation strategies and introduce the model and the assumptions for their comparison. In Section III we analytically study the strategies for a simple network topology. In Section IV we show the effect of different bandwidth allocation policies on random network topologies. In Section V we discuss the practical issues of our strategies, and Section VI concludes the paper.

II MODEL

We examine a very basic question: How to allocate the bandwidth of a link between unicast and multicast traffic? To eliminate all side effects and interferences we limit ourselves to static scenarios. We assume a given number of unicast sources, a given number of multicast sources, different numbers of receivers per multicast source, and a given bandwidth C for each network link to be allocated among the source-destination(s) pairs.

Our 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 per flow, $R(S_i, l)$, reached via an outgoing link l . iii) A constant traffic of every flow. iv) No arriving, nor departing flows. v) Each node is making the bandwidth allocation independently, a particular receiver sees the bandwidth that is the minimum bandwidth of all the bandwidth allocation on the links from the source to this receiver. vi) The sources have the capability to send through different bottlenecks via a cumulative layered transmission [4]. For receivers of the same multicast the (bottleneck) bandwidth seen by different receivers may be different.

II-A Bandwidth Allocation Strategies

We present three bandwidth allocation policies. Important to us is 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 and on the other hand we want to assure a fairness among different receivers.

We assume a network of nodes connected via links. In the beginning we assume every network link l has a capacity of a 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 POL is $B_{POL}(S_i, l)$ the shared bandwidth of link l for the receivers of S_i downstream of l .

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

- **Receiver Independent (RI):** Bandwidth is allocated in equal shares among each flow 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.

- **Linear Receiver Dependent (LinRD):** The share of bandwidth of link l allocated to a particular stream depends linearly on the number of receivers 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

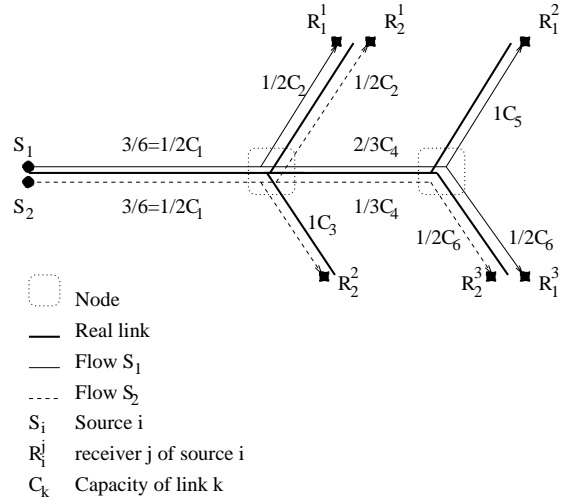


Fig. 1: Bandwidth allocation for linear receiver-dependent policy.

separate unicast flow². We allocate a share, for a multicast flow, corresponding 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 depends logarithmically on the number of receivers 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 :

Our three strategies are representatives of *classes* of strategies. We do not claim that the strategies we pick are the best representatives of its class. It is not the purpose of this paper to find the best representative of a class, we only want to study the trends between the classes.

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).

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 $B_{LinRD}(S_1, 1) = B_{LinRD}(S_2, 1) = \frac{3}{3+3} C_1 = 0.5 C_1$. For link 4, we have $R(S_1, 4) = 2$ and $R(S_2, 4) = 1$. Therefore we get $B_{LinRD}(S_1, 4) = 2/3 C_4$ and $B_{LinRD}(S_2, 4) = 1/3 C_4$. 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 $\min(1/2 C_1, 2/3 C_4, 1/2 C_6)$.

²We assume shortest path routing in the case of unicast and multicast.

II-B Measures, and Comparison criteria of 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 through receiver utility. Unfortunately, utility is a theoretical notion that does not allow to compare the utility of two different receivers and give an absolute (i.e. for all receivers) scale of utility [5]. We measure **receiver satisfaction** as the bandwidth an average receiver sees³. 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: $B_r = \min_{i=1, \dots, L} \{B_{POL}(S, l_i)\}$.

With the total number of receivers R of all sources we define the **mean bandwidth**:

$$\bar{B} = \frac{1}{R} \sum_{r=1}^R B_r \quad (1)$$

In [6] a global measure for the throughput delivered via the whole network is defined as the sum of the mean throughput over all the flows. In the global throughput measure, it is possible to weight multicast flows with a factor R^y , where R is the number of receivers and $0 < y < 1$. To the best of the authors knowledge, the approach in [6] is the only one taking into account the number of receivers of a multicast flow. While the approach in [6] takes into account the number of receivers to measure the global network throughput our approach is different in three aspects: First, we take the number of receivers into account for the *allocation* of the bandwidth on links. Second, we measure receiver satisfaction with respect to all receivers, not just the ones of a single group. Last, we use a policy (*LogRD*) that weights multicast flows in the allocation with the logarithm of the number of receivers.

Fairness

For inter-receiver fairness several measures exist, including product measure [7], and the fairness index [8], for a discussion of the different measures see [9].

In [6] inter-receiver fairness is defined 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 consider a fairness measure that takes into account all receivers of all flows.

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

$$\sigma = \sqrt{\frac{1}{R} \sum_{r=1}^R (\bar{B} - B_r)^2} \quad (2)$$

³ While there are other criteria to measure satisfaction such as delay or jitter, bandwidth is a measure of interest to the largest number of applications.

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 maximize both *receiver satisfaction* and *fairness*. 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 for the strategy p and the scenario s . 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 $\sigma_{max}(s) = \max_p \sigma(p, s)$ 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 [5] shows that *receiver satisfaction* is inconsistent with *fairness*⁴, 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 can not 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 $\frac{\sigma_A}{\sigma_B} \leq Lf$ and $\frac{\bar{B}_A}{\bar{B}_B} \geq Is$ where Lf is the maximum loss of *fairness* accepted for strategy A and Is is the minimum increase of *receiver satisfaction* for strategy A . But the choice of Lf and Is needs a fine tuning and seems to us pretty artificial.

In fact, for our study the behavior of the three strategies is so different that the evaluation of the tradeoff between *receiver satisfaction* and *fairness* does not lead to confusion.

III ANALYTICAL STUDY

We first compare the three bandwidth allocation policies from Section II for a basic network topology in order to gain some insight in their behavior. In Section IV we study the policies for random network topologies.

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 the 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.

⁴ In a mathematical economic language we can say that Pareto optimality is inconsistent with fairness criteria [5].

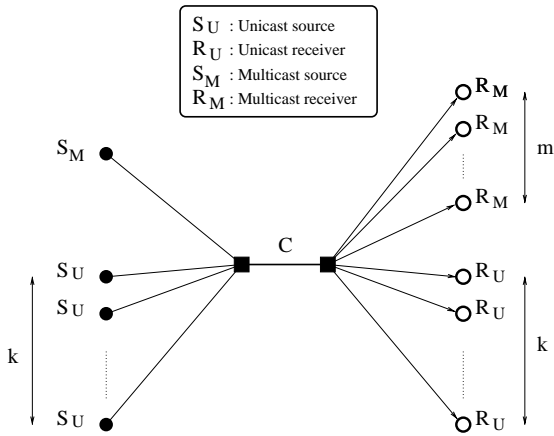


Fig. 2: One multicast flow and k unicast flows over a single link.




Fig. 3: Normalized mean bandwidth for the Star topology as a function of the size m of the multicast group; 10 unicasts.

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

$$\begin{aligned} \bar{B}_{RI} &= \frac{C}{k+1} \\ \bar{B}_{linRD} &= \frac{k+m^2}{(k+m)^2} C \\ \bar{B}_{logRD} &= \frac{k+m(1+\ln m)}{(k+m)(k+1+\ln m)} C \end{aligned}$$

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*.

We turn our attention to the case where the number of multicast receivers is increasing ($m = 1, \dots, 100$), and becomes

much higher than the number of unicasts ($k = 10$). We see in Fig. 3 that the mean bandwidth for *LinRD* and *LogRD* is increasing to multiples of the bandwidth of *RI*.

Surprisingly, we will observe nearly the same results in Section IV-C where we examine the three policies on a large random network. This indicates that the simple Star model with a single link can serve as a model for large networks.

We now briefly investigate the fairness among the receivers for the different allocation strategies and leave a more exhaustive examination to Section IV-C. With the Star model, 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 high bandwidth gains of the *LinRD* strategy result in a high unfairness for the average (unicast and multicast) receiver. For *LogRD* the repartitioning of the link bandwidth between unicast and multicast receivers is less unequal than in the case of *LinRD*, but still more pronounced than for *RI*.

We can conclude that among the three strategies *LogRD* meets best the tradeoff between receivers satisfaction and fairness.

IV SIMULATION

We now study the allocation strategies on network topologies that are richer in connectivity.

The generation of realistic network topologies is subject of active research ([10, 11, 12, 13]). It is commonly agreed that hierarchical topologies better represent a real Internetwork than do flat topologies. We use *tiers* ([11]) to create hierarchical topologies consisting of three levels: WAN, MAN, and LAN that aim to model the structure of the Internet topology [11].

For details about the network generation with *tiers* and the used parameters the reader is referred to Appendix A.

IV-A 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 we add at random locations of the LAN-leaves unicast senders and unicast receivers. The number of unicast flows ranges from 50 to 4000 unicast flows. Each simulation is repeated five times and averages are taken over the five repetitions. Confidence intervals are given for 95%.

We see in Fig. 4 that the 3 allocation policies give the same allocation. Indeed there are only unicasts flows and the differences of behavior between the policies depend only on the number of receivers downstream a link for a flow. Here the number of receivers is always one.

For a small number of unicast flows we have high standard deviation (Fig. 4) 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. 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




Fig. 4: Standard deviation of all receivers for an increasing number of unicast flows, $k = [50, \dots, 4000]$

too small on a large network, results in links shared only by a small number of flows. The statistical measure becomes meaningless. When the network is lightly loaded adding one flow can heavily change the bandwidth allocated to other flows and there is high heterogeneity in the bandwidth seen by the 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. 4 we chose our unicast environment with 2000 unicast flows to obtain a low bias due to the random location of the pairs sender-receiver.

IV-B Simulation Setups

For our simulation we proceed as follows: i) 2000 unicast sources and 2000 unicast receivers are chosen at random locations among the hosts. ii) 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. iii) We use shortest path routing [14] through the network to connect the 2000 unicast source-receiver pairs and to build the source-receivers multicast tree [15]. As routing metric, the length of the link as generated by `tiers` is used. iv) For every network link, the number of flows through the 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. v) 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*. vi) In order to determine the bandwidth received by each receiver, the minimum bandwidth (see (1)) allocated on all the links of the path from source to receiver is taken as the bandwidth B_r seen by that receiver.

The result of the simulation gives \bar{B} for the three bandwidth allocation strategies. We conduct different experiments.

Single multicast group

In Section IV-C we add one multicast group to the 2000 unicast flows. The size of the multicast group increases from 1

up to 6000 receivers. 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.

Multiple multicast groups

We did two experiments with multiple multicast groups. In the first one we add to the 2000 unicast sessions multiple multicast groups of the same size $m = 100$. In a second experiment we add to the 2000 unicast sessions multiple multicast groups of the same size $m = 20$, this experiment aims to model small conferencing groups. These experiments lead to a conclusion that does not significantly differ (for the purpose of this paper) from the single multicast group experiment. Due to space limitations, we cannot present the results for these experiments.

IV-C Single Multicast Group


We add a multicast session and vary the size of this session from 1 to 6000 receivers. There are 70 hosts on each LAN, the number of potential senders and receivers is therefore 12600.

In this section 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 over-sized compared to the 2000 unicast flows, this case gives a good indication about the robustness of the policies. For ease of reading, we display the plots with a logarithmic x-axis.


Fig. 5(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 (due to the high amount of unicast receivers compared to multicast receivers) the mean bandwidth. We claim that receiver-dependent policies increase receiver satisfaction.

A more accurate analysis needs to distinguish between unicast and multicast receivers. Due to space limitations we do not give a plot for the mean bandwidth of the unicast receivers. This plot is very simple: the mean bandwidth for the *RI* and the *LogRD* policies remains constant, around $600Kb/s$, with changing m , whereas the mean bandwidth for *LinRD* policy is the same than the one for *RI* until $m = 100$ then starts decreasing toward zero (for $m = 6000$, $\bar{B} = 100Kb/s$ for *LinRD*).

Multicast receivers are rewarded with a higher bandwidth for using multicast as the comparison of the mean bandwidth for unicast receivers and the mean bandwidth for multicast receivers shows (Fig. 6). This is not surprising as our policies aim to reward using multicast. Moreover, the increase in bandwidth for multicast receivers leads to a significant decrease of bandwidth for unicast receivers for the *LinRD* policy whereas it leads to a negligible loss of bandwidth for the *LogRD* policy even in the asymptotic case. In conclusion, the *LogRD* policy is the only policy among our policies that leads to a significant increase of receiver satisfaction for the average multicast



(a) Mean bandwidth.



(b) Standard deviation.

Fig. 5: Mean bandwidth and standard deviation of all receivers for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

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. 5(b)). This unfairness is caused by the difference of the lower bandwidth received by the unicast receivers compared to the higher bandwidth of a multicast receivers (Fig. 6). The receiver-dependent curves for σ tend to flatten for large group size, since the multicast receivers determine (due to their large number) the standard deviation over all the receivers. Due to space limitations we do not give a plot for the standard deviation of the unicast receivers. The standard deviation for unicast receivers 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. 7. The sparse multicast receiver setting results in a high heterogeneity of the allocated

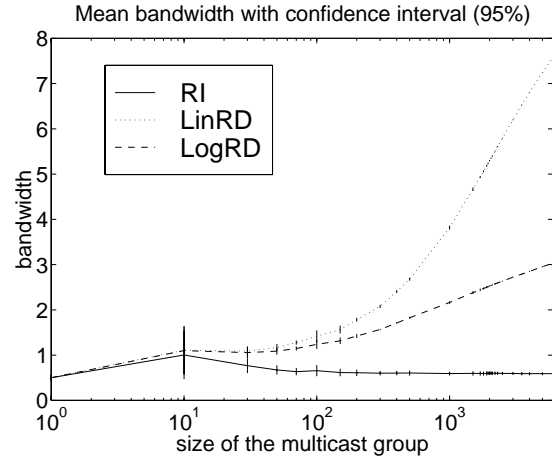


Fig. 6: Mean bandwidth of multicast receivers with confidence interval (95%) for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

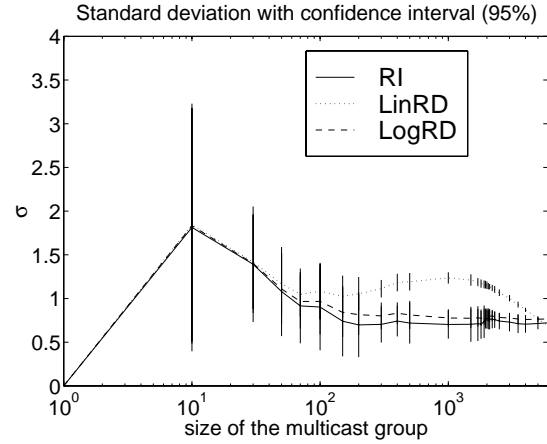


Fig. 7: Standard deviation of multicast receivers with confidence interval (95%) for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

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 receivers approaches the maximum bandwidth (the bandwidth of a LAN), see Fig. 6. 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 received bandwidth than have the unicast receivers (Fig. 7). 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 number of bottlenecks affects only a few receivers.

The standard deviation over all the receivers hides extreme




Fig. 8: Minimum bandwidth with confidence interval (95%) of the unicast receivers for an increasing multicast group size $m = [1, \dots, 6000]$, $k = 2000$, $M = 1$.

behavior of isolated receivers. To complete our study, we measure the minimum bandwidth, which gives an indication about the worst case seen by any receivers. The minimum bandwidth over all the receivers is dictated by the minimum bandwidth over the *unicast* receivers (we give only one plot, Fig. 8). As the size of the multicast group increases the minimum bandwidth for the *LinRD* policy dramatically decreases, whereas the minimum bandwidth for the *LogRD* policy remains close to the minimum bandwidth 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 used by unicast if we replace the multicast flow by an equivalent number of unicast flows, heavily decreases toward zero. Therefore we note the positive impact of multicast with the bandwidth allocated, and we claim that the use of multicast *greatly improves the worst case bandwidth allocation*. The minimum bandwidth increases for *multicast* receivers with the size of the multicast group for the receiver dependent policies (we do not give the plot of the minimum bandwidth for the multicast receivers). 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 *RI* we see that the increase in the multicast group size does not influence the average user satisfaction (Fig. 5(a)), nor the fairness among different receivers (Fig. 5(b)). Also, the difference between unicast and multicast receivers is minor concerning the bandwidth both received (Fig. 6), and the unfairness (Fig. 7). The *LogRD* policy is the only policy among our policies that significantly increases receiver satisfaction (Fig. 5(a)), keeps fairness close to the one of the *RI* policy (Fig. 5(b)), and does not starve unicast flows, even in asymptotic cases (Fig. 8).

Finally, one also should note the similarity between Fig. 5(a) obtained by simulation for a large network and Fig. 3 obtained by analysis of the star topology. This suggests that the star topology is a good model to study the impact of the three dif-

ferent bandwidth allocation policies.

V PRACTICAL ASPECTS

Up to now the advantages of using the number of downstream receivers were discussed. Keeping the number of receivers in network nodes has a certain cost but has other benefits that largely outweigh this cost:

- **Establishment of a valid business model for multicast.** Multicast saves bandwidth and is currently not used by network operators. The lack of a valid charging model contributes to this [16]. By keeping the number of receivers in network nodes different charging models for multicast can be applied—also charging models that include the number of receivers.
- **Feedback implosion avoidance.** Given the number of receivers is known in the network nodes, the distributed process of feedback accumulation [17], or feedback filtering in network nodes becomes possible and has a condition to terminate upon. If a node knows the number of receivers downstream, it knows the number of feedback messages it has to collect.

Another important question is how to introduce our strategy in a real network without starving unicast flows. In section IV we show that even in asymptotic cases the *LogRD* strategy does not starve unicast flows, but we do not have a hard guarantee about the bandwidth allocated to unicast receivers. In fact we devise our strategy to be used in a hierarchical link sharing scheme (see [18], [19] for hierarchical link sharing models). The idea is to introduce our policy in the general scheduler [18] (for instance we can configure the weight of a GPS [20], [21] scheduler with our policy to achieve our goal), and to fit an administrative constraint in the link sharing scheduler (for instance we guarantee that unicast traffic receives at least 5% of the link bandwidth). Moreover in [22] the authors show that it is possible to integrate efficiently a mechanism like HWFQ ([19]) in a gigabit router, and WFQ is already available in many of the recent routers [23].

VI CONCLUSION

If we want to introduce multicast in the Internet we need to 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 a simple star topology. We showed that the *LogRD* policy always leads to the best tradeoff between receiver satisfaction and fairness. The striking similarities between the analytical study and the simulations confirm that we had chosen a good model.

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 a real Internet. We showed that: allocating link bandwidth dependent on the flows' number of receivers downstream results in a higher receiver satisfaction: the *logRD* policy provides the best trade-off 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 too, however at the expense of unacceptable decrease of fairness.

There are several open questions: Do we need to implement our mechanism in every network node, or is it possible to introduce it only in a subset of well chosen nodes? Are there better classes of policies than the *LogRD* policy? These questions will be addressed in future work.

ACKNOWLEDGMENT

Many thanks to Jim Kurose for sharing his insights about the notions of receiver satisfaction and fairness. We also want to thank the anonymous reviewers for their 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] S. E. Deering, "Multicast routing in internetworks and extended lans," in *Proc. ACM SIGCOMM 88*, pp. 55–64. Stanford, CA, August 1988.
- [2] Hans Eriksson, "MBONE: The multicast backbone," *Communications of the ACM*, vol. 37, no. 8, pp. 54–60, Aug. 1994.
- [3] J. Nonnenmacher and E.W. Biersack, "Asynchronous multicast push: Amp," in *Proceedings of ICC97*, Cannes, France, November 1997, pp. 419–430.
- [4] S. McCanne, V. Jacobson, and M. Vetterli, "Receiver-driven layered multicast," in *SIGCOMM 96*, Aug. 1996, pp. 117–130.
- [5] Allan Feldman, *Welfare economics and social choice theory*, Martinus Nijhoff Publishing, Boston, 1980.
- [6] T. Jiang, M. H. Ammar, and E. W. Zegura, "Inter-receiver fairness: A novel performance measure for multicast abr sessions," in *Proceedings of ACM Sigmetrics*, June 1998.
- [7] K. Bharat-Kumar and J. Jeffrey, "A new approach to performance-oriented flow control," *IEEE Transactions on Communications*, vol. 29, no. 4, 1981.
- [8] R. Jain, D. M. Chiu, and W. Hawe, "A quantitative measure of fairness and discrimination for resource allocation in shared systems," Tech. Rep. TR-301, DEC, Littleton, MA.
- [9] S. Floyd, "Connections with multiple congested gateways in packet-switched networks part 1:one-way traffic," *Computer Communications Review*, vol. 21, no. 5, pp. 30–47, October 1991.
- [10] Ken Calvert, Matt Doar, and Ellen W. Zegura, "Modeling internet topology," *IEEE Communications Magazine*, vol. 35, June 1997.
- [11] Matthew B. Doar, "A better model for generating test networks," in *Proceedings of IEEE Global Internet*, London,

UK, November 1996, IEEE.

- [12] Ellen W. Zegura, Ken Calvert, and S. Bhattacharjee, "How to model an internetwork," in *Infocom '96*, March 1996.
- [13] Ellen W. Zegura, Kenneth Calvert, and M. Jeff Donahoo, "A quantitative comparison of graph-based models for internet topology," *IEEE/ACM Transactions on Networking*, vol. 5, no. 6, December 1997.
- [14] T. H. Cormen, C. E. Leiserson, and R. L. Rivest, *Introduction to Algorithms*, The MIT Press, 1990.
- [15] M. Doar and I. Leslie, "How bad is naïve multicast routing," in *Proceedings of IEEE INFOCOM'93*, 1993, vol. 1, pp. 82–89.
- [16] R. Comerford, "State of the internet: Roundtable 4.0," *IEEE Spectrum*, vol. 35, no. 10, October 1998.
- [17] 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*, vol. 15, no. 3, pp. 407 – 421, April 1997.
- [18] S. Floyd and V. Jacobson, "Link-sharing and resource management models for packet networks," *IEEE/ACM Transactions on Networking*, vol. 3, no. 4, pp. 365–386, August 1995.
- [19] Jon C.R. Bennett and H. Zhang, "Hierarchical packet fair queueing algorithms," *IEEE/ACM Transactions on Networking*, vol. 5, no. 5, pp. 675–689, October 1997.
- [20] A. K. Parekh and R. G. Gallager, "A generalized processor sharing approach to flow control in integrated services networks," in *Proc. IEEE INFOCOM'93*, 1993, pp. 521–530.
- [21] E. Hahne and R. Gallager, "Round robin scheduling for fair flow control in data communications networks," in *Proceedings of IEEE Conference on Communications*, June 1986.
- [22] Vijay P. Kumar, T. V. Lakshman, and D. Stiliadis, "Beyond best effort: Router architectures for the differentiated services of tomorrow's internet," *IEEE Communications Magazine*, vol. 36, no. 5, pp. 152–164, May 1998.
- [23] Cisco, "Advanced qos services for the intelligent internet," White Paper, May 1997.

APPENDIX

A 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 155Mb/s, the capacity of MAN links is 55Mb/s, and the capacity of LAN links is 10Mb/s.

Each LAN is represented as a single node and connects several hosts via a 10Mb/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.