Improving the quality specification of multi-layered applications

Marcelo Dias de Amorim* — Otto Carlos Muniz Bandeira Duarte**

* *Laboratoire de Recherche en Informatique (LRI)*
Université Paris XI
Bâtiment 490 – 91405 – Orsay Cedex – France
amorim@lri.fr

** *Grupo de Teleinformática e Automação – COPPE*
Universidade Federal do Rio de Janeiro
P.O. Box 68504 – 21945-970 – Rio de Janeiro – RJ – Brasil
otto@gta.ufrj.br

ABSTRACT. *Using a single parameter for the specification of the level of quality of service in multicast multi-layered applications is not adequate because of the different application requirements and the heterogeneous nature of the receivers. In this paper, we propose a three-metric approach that takes into account the type of the applications, which improves the global quality of the multicast sessions from both network and user point of views. We define a simple function that takes into account the density of satisfied users, the amount of allocated bandwidth in the multicast tree, and the degradation at the receivers. We analyze the proposed multi-criteria algorithm in an environment with other competing flows and show the improvements of the global quality of the applications.*

RÉSUMÉ. *L'utilisation d'un seul paramètre lors de la spécification du niveau de qualité de service des applications multicast en multi-couches n'est pas idéale à cause des différents types d'application et de la nature hétérogène des récepteurs. Dans cet article, nous proposons une approche avec trois métriques qui prend en compte le type d'application, ce qui optimise la qualité globale de la session multicast des points de vue de l'utilisateur et aussi du réseau. Nous définissons une fonction simple qui considère la densité d'utilisateurs satisfaits, la quantité de bande-passante allouée et la dégradation dans les récepteurs. Nous analysons le schéma proposé dans un scénario avec d'autres flux et montrons les améliorations obtenues sur la qualité globale de la communication.*

KEYWORDS: *Multicast, multi-layered applications, fairness, global quality.*

MOTS-CLÉS : *Multicast, application en multi-couches, justice, qualité globale.*

1. Introduction

The specification of the level of quality of service in multicast multi-layered applications by a single parameter is not adequate because of the heterogeneous nature of the receivers. The modern approach for providing quality of service to multimedia applications consists of classifying the applications into different pools of service. The specification of the level of quality of service in multicast multi-layered applications by a single parameter is then inappropriate because applications may have different requirements in terms of the quality of service and also because receivers are heterogeneous in their receiving capacities. Current techniques are being replaced by proposals that take into account different natures of the communication, as for example service differentiation [BLA 98, ZIV 99] and QoS routing [WAN 96, COS 00]. In this latter, to find the most suitable route, the routing algorithm considers multiple metrics such as the available bandwidth, the shortest path, and/or the minimum delay.

We argue that a multi-criteria approach should also be used in the case of multicast distribution of multi-layered applications [AMO 00]. A number of works have already dealt with QoS issues in such architectures, but they were interested in solving other related problems [SAP 00, VIC 00, LI 99, SAR 00, RUB 99, SIL 00, MCC 96, LI 98, VIC 98, AMO 99].

The receiver-driven layered multicast (RLM) [MCC 96] is a rate-adaptive protocol for the distribution of layered video using multiple IP-multicast groups where the receivers subscribe to the number of layers they want. Nevertheless, they are limited to the layers the source decides to transmit. The Layered Video Multicast with Retransmission (LVRM) [LI 98] has been proposed for the distribution of layered video over the Internet. LVRM deploys an error recovery scheme using retransmissions to adapt to network congestion and to heterogeneity using a hierarchical rate control mechanism. In [VIC 00, VIC 98], Vickers et al. propose a rate-based control mechanism in which a number of feedback control packets containing the current network congestion state are exchanged between the source and the receivers. Based on the content of these packets, the source estimates the number of video layers to transmit and their respective rates. The single-loop packet-merging algorithm proposed in [AMO 99] is an algorithm to improve the fairness in multicast distribution of multi-layered video. The control information scheme is similar to the approach proposed in [VIC 00, VIC 98]. However, intermediate nodes perform packet merging in only one loop and use the concept of virtual layers. With virtual layers, the nodes consider that the source is able to generate more than the actual number of layers.

Despite the effort applied to improve the quality of multicast multi-layered schemes, the usual “goodput-based” approaches studied in the literature have some limitations that arise when we consider QoS-parameters other than the global degradation at the receivers [AMO 00]. This paper proposes a single function that takes into account the density of satisfied users, the amount of allocated bandwidth in the multicast tree, and the degradation at the receivers. This function allows the user to adjust the weight of each metric in such a way that the application can be classified according to one of three "pool of services." A pool of services is characterized by its correspondent prioritized metric. A specific application that aims at saving bandwidth should give more weight to the metric that corresponds to the bandwidth allocated through the multicast tree. The same should be applied to an application that requires the optimization of the quality of the presentation at the receivers.

We have simulated our system on a network formed by one source and ten heterogeneous receivers. We have analyzed the available bandwidth and the allocated resources in the links. We have also observed the degree of allocated bandwidth through the multicast tree as an example of application of our algorithm and the results show that we can easily control the prioritized metric based on the type of the application that is being carried on. For instance, the simulation results show an improvement of about 30% on the allocated bandwidth.

This paper is organized as follows. Section 2 presents our arguments about using multiple metrics to measure the quality of the multicast session. The analysis on the three metrics used in our framework is presented in Section 3. We present in Section 4 the proposed "global quality of the multicast session," built from the independent metrics and in Section 5 the analysis of our proposal. Finally, Section 6 concludes the paper.

2. Single argument

In an conventional adaptive system, the source computes the rates of the layers based on control information sent by the receivers. One of the information is the maximum available bandwidth in the path from the source to the receivers, which corresponds to maximum required rate. Let r_i be then the rate required by receiver i, r_i^r the received rate, and N the number of receivers in the multicast session. The layers are computed in such a way to minimize the global degradation

$$\delta = \sum_{i=1}^{N} \delta_i, \qquad [1]$$

where $\delta_i = (r_i - r_i^r)$ and $r_i \geq r_i^r$.

We consider in this paper a feedback-control scheme similar to the one presented in [VIC 00, VIC 98], where control packets are transmitted from the destinations to the source to indicate the rate they want to receive. These packets are combined at intermediate nodes to avoid feedback implosion at the source. After a merging procedure, each resulting packet contains a number of entries that store a rate f_i and the number of destinations that want to receive this rate (f_i^*). The number of entries carried out by feedback-control packets corresponds to the number of layers the source can transmit. When performing a merging procedure, the node looks for the entries that must be discarded according to the merging algorithm and creates a new feedback-control packet to be sent to the source.

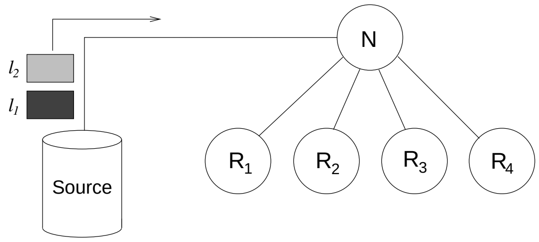

Figure 1. *Example of multicast session.*

The classical goodput approach considers the degradation at the receivers as the unique metric when computing the layers. The global degradation at the receivers is certainly one of the most important parameters for defining a fair multicast multi-layered session. Nevertheless, we argue in this paper that the global degradation should not be considered as the only one. In [AMO 00], the authors have analyzed several aspects of transmitting multi-layered applications through multicast groups based only on the global degradation at the receivers. We reproduce here a simple example with one source and four receivers $(r_1, \ldots, r_4)$ as shown in Fig. 1. In this scenario, the source transmits two layers, $L = \{l_1\ Mbps, l_2\ Mbps\}$, and that the merging algorithm only uses the global degradation at the receivers to perform the merging procedure. Consider in a first time that r_1 =1 Mbps, $r_2 = r_3$ =2 Mbps, and r_4 =3.9 Mbps. Since the rate of the base layer is limited by the slowest receiver, in this case, the global degradation at the receivers is minimized with $L = \{1, 1\}$. Therefore, the correspondent global degradation is $\delta = 1.9$ Mbps for a bandwidth utilization of 2 Mbps.

Consider now that the capacity of receiver 4 changes from 3.9 to 4.1 Mbps. The new layers that result in the lowest global degradation are l_1=1 Mbps and l_2=3.1 Mbps, and the correspondent global degradation is δ=2 Mbps. Note that for a variation of 5% in the receiving capacity of receiver 4, layer 2 increases its global rate of $\sim$300%, which results in a substantial growth of the bandwidth utilization. Moreover, in the first scenario there were three destinations receiving 100% of the required bandwidth (and only one destination receiving $\sim$50%), while in the second case only two receivers are supplied with 100% and two receivers with 50%.

It is quite intuitive that it is better to maintain l_1 and l_2 at 1 Mbps in order to have lower bandwidth utilization even if the global degradation is slightly greater than the optimum value. The above example shows that the use of a single metric may not be always adapted to every type of application. For example, in a pay-per-view video service, it is important that the subscribers receive the maximum possible video quality even if the server rejects new calls. In video surveillance, maximizing the video quality may not be the main objective, but the reduction of the allocated network resources. On another side, broadcast advertisements may require that the data are received by a number of destinations as large as possible. We show in this paper that the use of different criteria can reduce the problem of multicast delivery of multi-layered applications if a scheme of weighted parameters is applied. Each one of the proposed metrics is prioritized depending on the application that is carried out.

3. Multiple parameters

We define in this section the three parameters used in our scheme: density of satisfied users (γ), weighted bandwidth utilization (ω), and degradation at the receivers (δ).

For a receiver i requesting r_i unities of bandwidth, we say that it is partially satisfied if it receives r_i^r, with $r_i^r < r_i$. We define the degree of satisfaction for receiver i as $\gamma_i = \frac{r_i^r}{r_i}$. In a multicast session with heterogeneous receivers, we estimate the global density of satisfied users by combining γ_i of each receiver.

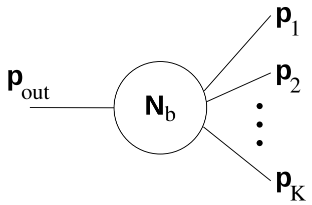

Figure 2. *Our benchmark node.*

Consider a node in the network with K incoming links and one outgoing link as shows Fig. 2. Without loss of generality, we also consider that we always have one arriving packet per incoming link. For each incoming packet p_i, $i = 1, 2, \ldots, K$, we choose a layer $l_j \in L$ such that

$$\sum_{a=1}^{j} l_a \leq f_{p_i} < \sum_{a=1}^{j+1} l_a. \qquad [2]$$

We call l_j the *adoptive layer* and do $g_{p_i} = \sum_{a=1}^{j} l_a$. The output γ_{out} is the weighted sum of the individual inputs $\gamma_1, \gamma_2, \ldots, \gamma_K$ plus the loss introduced by the merging operation $f_{p_i} \rightarrow g_{p_i}$, and is given by

$$\gamma_{out} = \frac{\sum_{i=1}^{K} m_{i,c} f_{p_i}^* \frac{g_{p_i}}{f_{p_i}} \gamma_{p_i}}{\sum_{i=1}^{K} f_{p_i}^*} \qquad [3]$$

In equation 3, $m_{i,j}$ is one element of the matrix $M_{K,L}$ respecting the following rules:

$$m_{i,j} = \begin{cases} 1 & \text{if } \sum_{a=1}^{j} l_a \leq f_{p_i} < \sum_{a=1}^{j+1} l_a, \\ 0 & \text{otherwise.} \end{cases} \quad [4]$$

Fig. 3 shows the behavior of γ in the multicast session of the example presented in Section 2, where receivers r_1 and r_4 have fix individual rates with r_1 =1 Mbps and $r_4 = 10$ Mbps. Receivers r_2 and r_3 have variable individual rate in the range $[r_1, r_4]$.

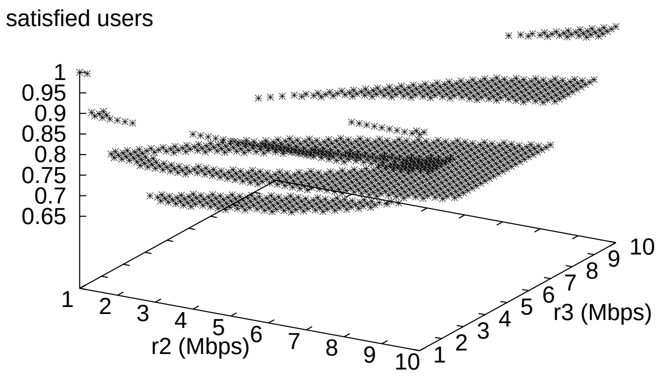

Figure 3. *Example of the behavior of γ.*

The second metric we analyze is the degradation at the receivers. As we have stated, this metric has been used as the main parameter to determine the quality of multicast multi-layered schemes. We show that the degradation at the receivers behaves quite differently from the other QoS parameters. As we have shown in Section 2, we denote the *local degradation* as the difference between the required rate and the actual rate received at the destination, formulated as $\delta_i = r_i - r_i^r$. From the source's point of view, this notion is represented by the *global degradation of the communication* defined as $\delta = \sum_{i=1}^{N} \delta_{r_i}$, where N is the number of receivers in the session. As in the analysis for the density of satisfied users, a certain number of incoming packets are merged into an output packet containing a number of entries equal to the number of layers the source is supposed to transmit. In our benchmark node of Fig. 1, each entry $p_1, \ldots, p_K$ carries also the corresponding downstream degradation, which has to be added to the degradation at the node due to the merging procedure,

$$\delta_{out} = \sum_{i=1}^{K} m_{i,c}[f^*_{p_i}(f_{p_i} - g_{p_i}) + \delta_{p_i}]. \quad [5]$$

We have observed the behavior of δ under the same circumstances used for γ in the example of Section 2. The difference is that in the case of δ we show the results for the values that minimize the sum presented in equation 5. We plot in Fig. 4 the behavior of δ in our network.

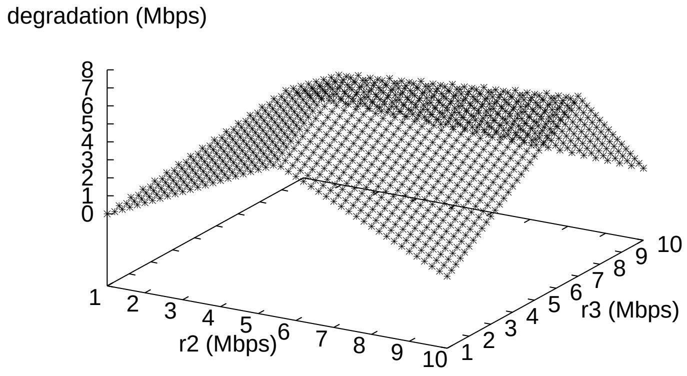

Figure 4. *Example of the behavior of* δ.

The third metric used by our extended QoS system is the *weighted allocated bandwidth*, denoted by ω. The level of bandwidth utilization in the network is computed directly at the source by combining the rates of the layers and the number of requests for each one of them. In our benchmark node, the value of ω in the output packet is the sum of the weighted allocated bandwidth in the downstream tree, given by $\omega_c^{out} = 1 + \sum_{i=1}^{K} \omega_{p_i}$. As we have done for the other two metrics, we have also evaluated the behavior of ω in our example architecture. The global allocated bandwidth is an estimation of the bandwidth allocated through the multicast tree, and is given by

$$\omega_{out} = \frac{\sum_{i=1}^{L} (f_{r_i}^{*} f_{r_i})}{\sum_{i=1}^{L} f_{r_i}^{*}}. \qquad [6]$$

The behavior of ω in our example network follows the same principle of δ in the sense the main objective is to minimize the weighted bandwidth allocation. Fig. 5 shows the obtained curve for our test.

It is easy to observe with the help of Figs. 3–5 that the optimization of the metrics cannot be obtained at the same time. For instance, δ and ω should be minimized to yield an improved global quality, but what we see is that δ typically grows when ω decreases, and vice-versa. In the next session, we will propose a mechanism to assemble the metrics in order to optimize the global quality based on the application being carried out.

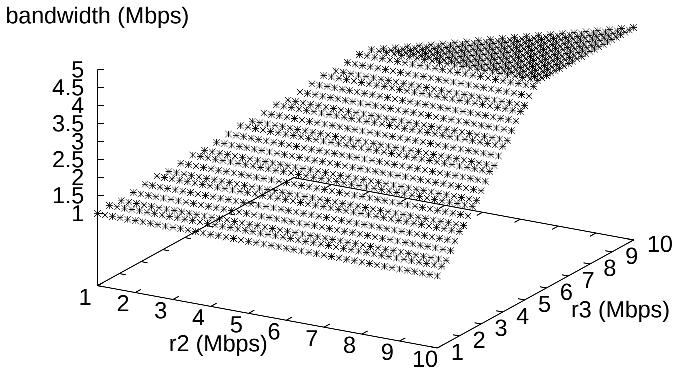

Figure 5. *Example of the behavior of* ω*.*

4. The global quality of the communication

As we have argued in section 2, applications have different requirements and the way the QoS metrics should be composed is particular in each case. We have classified the applications in three groups of interest that require the optimization of:

– The density of satisfied users.

– The allocated bandwidth through the multicast tree.

– The degradation at the receivers.

Based on the above requirements, we define as the *global quality of the multicast session* the parameter Q that is a result of mixing γ, ω, and δ. The first rule Q must follow is the direct correspondence to the metrics. As we have seen, the global quality of the multicast session must grow if:

– The density of satisfied users grows.

– The weighted bandwidth utilization decreases.

– The quality degradation decreases.

Our quality measure is the function $Q(\gamma, \omega, \delta)$. Despite some correlation between these parameters, we have considered each metric independently because of the complexity and inaccuracy that would be introduced if we had taken into account these correlations. This simplification yields a more tractable management of the metrics and does not interfere in our objective. The function Q has the following aspect

$$Q(\gamma, \omega, \delta) = f_\gamma(\gamma) + f_\omega(\omega) + f_\delta(\delta). \qquad [7]$$

We define this function to be normalized between 0 and 1. Each one of the right-sided functions of equation 7 is also normalized between 0 and 1. Since γ, ω, and δ

depend tightly on the topology of the network, we cannot define a closed function to accomplish our needs. To implement such a feature, we have to know the minimum and maximum bounds on these variables. The minimum value is 0 in all three cases, which corresponds to the cases when the source is transmitting at 0 bps (for γ and ω) and when all of the receivers are being provided with the required rate (for δ). The maximum values are:

$$\gamma_{max} = 1 \quad [8]$$

$$\omega_{max} = \lim_{\max[f^*_{p_i}] \to \sum_{i=1}^{L} f^*_{p_i}} \frac{\sum_{i=1}^{L} f^*_{p_i} f_{p_i}}{\sum_{i=1}^{L} f^*_{p_i}} = \max[f_{p_i}] \quad [9]$$

$$\delta_{max} = \sum_{i=1}^{K} (f^*_{p_i} f_{p_i} + \delta_{p_i}) \quad [10]$$

The global quality of the multicast session is:

$$Q(\gamma, \omega, \delta) = \frac{a\frac{\gamma}{\gamma_{max}} + b\left(1 - \frac{\omega}{\omega_{max}}\right) + c\left(1 - \frac{\delta}{\delta_{max}}\right)}{3(a+b+c)} \quad [11]$$

In the above equation, the variables a, b, and c represent the weight of each metric. By adjusting these values, we can place the application in the most adequate pool of service.

5. Analysis

We will analyze in this section the methodology proposed in this paper over a multicast session composed by one source and ten receivers, as shows Fig. 6. We used symmetric links and homogeneous delays and capacities in such a way to force the results to depend only on the available bandwidth. Variable resources are obtained in the links by introducing a background traffic generated by foreigner exponential sources. The available bandwidth is obtained by the difference between the link's capacity and the background traffic.

As we have stated throughout this paper, the system is configured to be optimized for specific types of applications. We have then many possibilities for the same multicast session. For our simulations shown below, we have chosen the weighted allocated bandwidth as the prioritized metric. We focus on the the responsiveness of the system when merging procedures in the intermediate nodes have to be performed. In our topology, since we have ten receivers connected to one node, at least one merging procedure is executed, except when all of the receivers require the same rate, which does

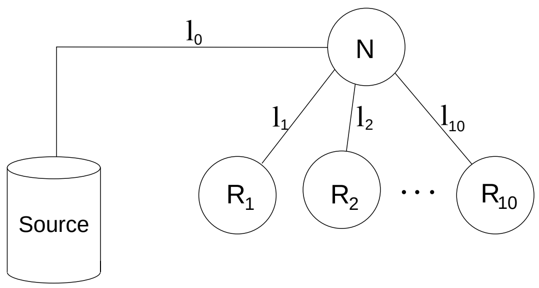

Figure 6. *The simulation topology.*

not correspond to the average case, as we will see in the obtained results. Fig. 7 shows the evolution of the quality of the session obtained after each reception of a feedback packet. Of course, this does not have any special meaning, but is a way to illustrate the evolution of the metric. What is really important is the control that we introduce in the system as we will see in the following.

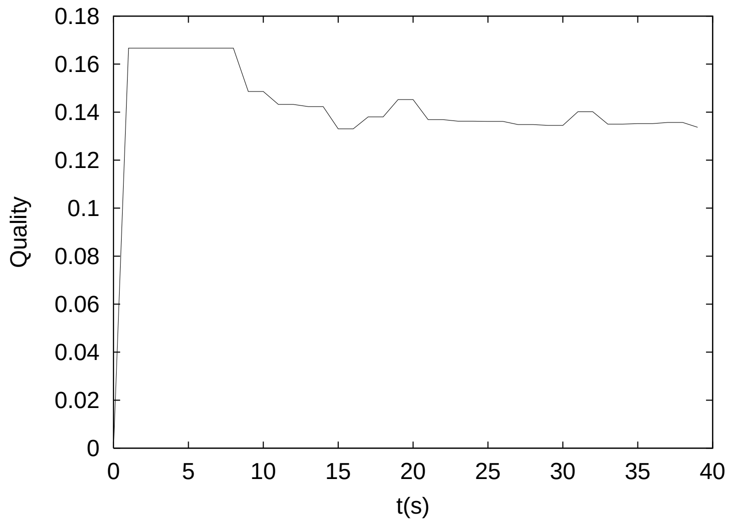

Figure 7. *The quality of the communication.*

We have selected two particular links to be represented in our graphics. They are the least and the most stringent ones, respectively links l_2 and l_7 in Fig. 6. In Fig. 8 we show the available bandwidth and the rate transmitted in link l_2. Which is interesting to observe in this curve is that the rate is in average about 50% of the available bandwidth. It means that the control packets sent by receiver 2 are being combined to the packets of a receiver requiring about a half its available bandwidth.

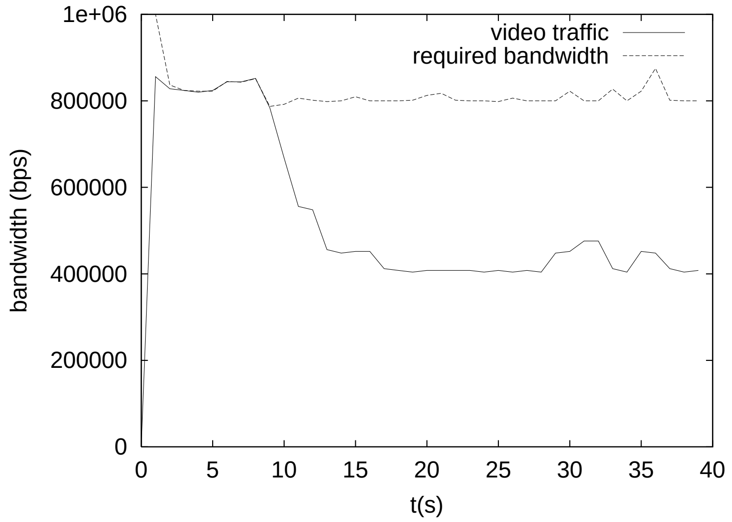

Figure 8. *Available and allocated bandwidths in link l_2.*

Another important result is the rate received by the most stringent receiver that corresponds, in our topology, to link l_7. The results are shown in Fig. 9. Observe that the required rate is almost all the time equal to the available bandwidth. The only regions where it is not true is when another receiver requires a rate under the one of receiver 7, which corresponds in the graphic to the periods when the curves do not coincide. This is expected since the merging procedure never discards the lowest rate.

We show now the bandwidth allocated through the multicast tree. As we have already stated, in our example we choose the weighted allocated bandwidth as the main metric, i.e., the algorithm prioritizes the configuration that results in the lowest allocated bandwidth. Fig. 10 shows the obtained results for both algorithms, the proposed algorithm based on three metrics and the goodput-based approach. After a temporary stabilization period, the proposed approach results in reduced allocated bandwidth. When compared with the classical approach the new approach leads to 30% of bandwidth saving.

These curves depicts the behavior of the system when a large number of merging procedures are performed in one node. This shows that the system responds well to heterogeneous receivers and also that the quality of the multicast session is quite stable during the communication.

6. Conclusion

In classical multi-layered multicast schemes, the computation of the rates of the layers are based on a single metric. In this paper, we have shown that the use of

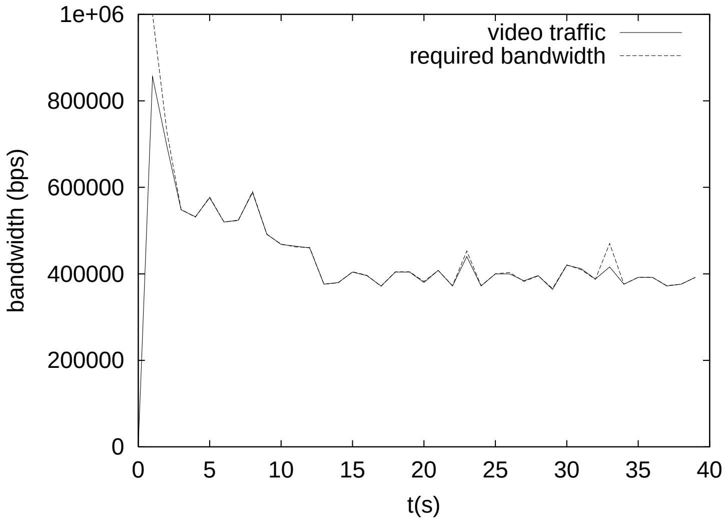

Figure 9. *Available and allocated bandwidths in link* $l_{1,7}$.

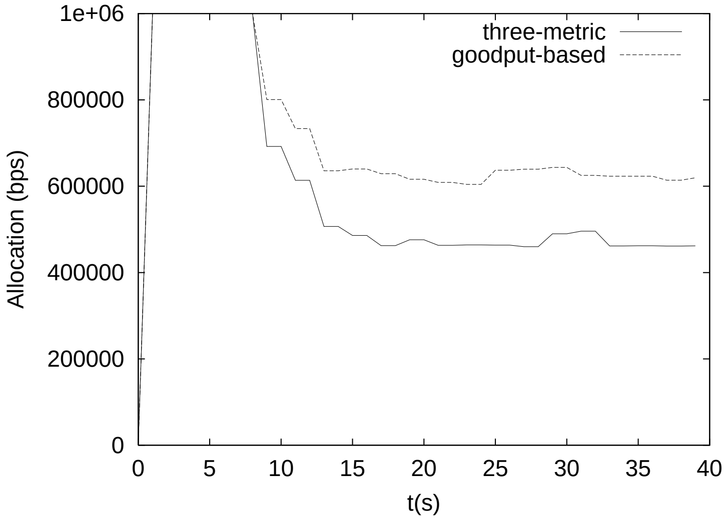

Figure 10. *Allocated bandwidth through the multicast tree.*

other metrics improves the quality of the communication when we introduce in the system the knowledge about the type of the application. We have proposed a combined metric that takes into account the density of satisfied users, the weighted bandwidth allocation, and the quality degradation at the receivers. We have shown through a number of simulations that the system responds well the a relatively large number of heterogeneous receivers connected to one merging node. We have observed that if an application requires the optimization of the allocated resources, the simulation results show that the proposed approach leads to an improvement of 30% on the allocated resources when compared with the classical approach.

Acknowledgment

The authors would like to thank CAPES, COFECUB, UFRJ, and CNRS for their support of this work.

7. References

[AMO 99] DE AMORIM M. D., DUARTE O. C. M. B., PUJOLLE G., "Single-Loop Packet Merging for Receiver Oriented Multicast Multi-Layered Video", *International Conference in Computer Communication*, Tokyo, Japan, 1999.

[AMO 00] DE AMORIM M. D., DUARTE O. C. M. B., PUJOLLE G., "Multi-Criteria Arguments for Improving the Fairness of Layered Multicast Applications", *Lecture Notes in Computer Science 1815*, Paris, France, 2000, p. 1–10.

[BLA 98] BLAKE S., BLACK D. L., CARLSON M., DAVIES E., WANG Z., WEISS W., "An Architecture for Differentiated Services", *Internet RFC 2475*, , 1998.

[COS 00] COSTA L. H. M. K., FDIDA S., DUARTE O. C. M. B., "Distance-vector QoS-based Routing with Three Metrics", *Lecture Notes in Computer Science 1815*, Paris, France, 2000, p. 847–858.

[LI 98] LI X., PAUL S., AMMAR M. H., "Layered Video Multicast with Retransmissions (LVRM): Evaluation of Hierarchical Rate Control", *IEEE Infocom*, San Francisco, CA, USA, 1998.

[LI 99] LI X., AMMAR M. H., PAUL S., "Video Multicast over the Internet", *IEEE Network Magazine*, , 1999.

[MCC 96] MCCANNE S., JACOBSON V., VITTERLI M., "Receiver-Driven Layered Multicast", *ACM Sigcomm*, Stanford, CA, USA, 1996.

[RUB 99] RUBENSTEIN D., KUROSE J., TOWSLEY D., "The Impact of Multicast Layering on Network Fairness", *ACM Sigcomm*, Cambridge, Massachusets, USA, 1999.

[SAP 00] SAPARILLA D., ROSS K. W., "Optimal Streaming of Layered Video", *IEEE Infocom*, Tel-Aviv, Israel, 2000.

[SAR 00] SARKAR S., TASSIULAS L., "Distributed Algorithms for Computation of Fair Rates in Multirate Multicast Trees", *IEEE Infocom*, Tel-Aviv, Israel, 2000.

[SIL 00] DA SILVA GONÇALVES P. A., DE REZENDE J. F., DUARTE O. C. M. B., "An Active Service for Multicast Video Distribution", *Journal of the Brazilian Computer Society*, vol. 7, num. 1, 2000.

[VIC 98] VICKERS B. J., ALBUQUERQUE C., SUDA T., "Adaptive Multicast of Multi-Layered Video: Rate-Based and Credit-Based Approaches", *IEEE Infocom*, San Francisco, CA, USA, 1998.

[VIC 00] VICKERS B. J., ALBUQUERQUE C., SUDA T., "Source-Adaptive Multi-Layered Multicast Algorithms for Real-Time Video Distribution", *IEEE/ACM Transactions on Networking*, vol. 8, num. 6, 2000.

[WAN 96] WANG Z., CROWCROFT J., "Quality of Service Routing for Supporting Multimedia applications", *IEEE Journal on Selected Areas in Communications*, vol. 14, num. 7, 1996.

[ZIV 99] ZIVIANI A., DE REZENDE J. F., DUARTE O. C. M. B., "Towards a Differentiated Services Support for Voice Traffic", *IEEE Globecom*, Rio de Janeiro, RJ, Brazil, 1999, p. 59–63.