Multi-criteria arguments for improving the fairness of layered multicast applications*

Marcelo Dias de Amorim[1,2], Otto Carlos M. B. Duarte[2], and Guy Pujolle[1]

[1] PRiSM Laboratory, University of Versailles
45, Avenue des Etats-Unis – 78035 – Versailles – France
{amorim,gp}@prism.uvsq.fr
[2] GTA/COPPE/EE – Federal University of Rio de Janeiro
P.O. Box 68504 – 21945-970 – Rio de Janeiro – RJ – Brazil
otto@gta.ufrj.br

Abstract. In this paper, we address the problem of choosing the most adequate rates for the different layers in multicast distribution of multi-layered applications. We argue that depending on the type of the application that is being carried out, we face a number of problems to decide what are the most adequate bounds on the parameters to achieve the best global quality. We propose then multi-decision criteria that improve the global aggregated quality of the communication in the multicast session. We consider the quality of the application at the receiver, the bandwidth utilization, and the satisfaction index to measure the combined user/network quality of the communication.

1 Introduction

For the transport of multimedia applications through multicast sessions with heterogeneous receivers, the use of multi-layered coding has shown to be a worthwhile approach. In such a technique, the source is coded in a base layer and in one or more enhancement layers and the quality of the presentation depends on the number of layers that are played. Multi-layered coding has been applied to improve the fairness of a wide variety of multicast applications, mainly for video distribution [1–10]. This approach is interesting in multicast communications because destinations with different receiving capacities (available bandwidth, local resources, etc) can choose the layers they are able to receive.

Multi-layered applications over multicast networks require that new mechanisms are embedded in the communication protocol. One of the most difficult problems is to choose the metrics that should be considered by the algorithm that computes the rates of the video layers in order to maximize the global fairness of the application. This is a strategic task because of the subjective decision of which metric should be prioritized. For instance, in [5, 8], video rates are calculated based only on the degree of video degradation at the receivers, i.e.,

* This work has been supported by CAPES, COFECUB, CNPq, UFRJ, CNRS, and UVSQ.

the difference between the required and the received video rates. In this paper, we propose a set of other potential metrics and argue that one of them should be prioritized according to the type of the application. In our scenario, the application runs over a rate-based adaptive multicast session, where the receivers periodically send control packets to the source containing information about the video rate they want (or are able) to receive [5]. Based on these packets, the source computes the most appropriate rates to optimize the quality of the communication. However, it must be established what means a good global quality. We propose a first approach to this problem by considering the bandwidth utilization, the satisfaction index, and the degree of degradation at the receivers as the metrics to be used by the algorithm to compute the layers. We show that all metrics cannot be optimized at the same time and give some examples showing the dependency between the type of the application and the metric to be prioritized.

The following of this paper is organized in four sections. In Section 2, we present some multi-layered multicast schemes proposed in the literature. In Section 3, we discuss the limitations of using a single metric to compute the rates of the layers. The use of multiple metrics and how these metrics react to fluctuating network congestion are proposed in Section 4. Finally, Section 5 concludes this paper.

2 Multi-layered multicast schemes

Different strategies have been used to provide quality of service to applications in networks that offer no guarantees such as the Internet. One approach is to apply resource reservation, in which resources are set aside for a particular flow [11]. A more recent technique is to introduce adaptability in the applications [5–7, 12–15]. In such a case, applications are kept informed about available network resources and dynamically adapt their transmission rates to the current state of the network. This is particularly interesting in heterogeneous multicast communication in which data must be delivered to a number of destinations with different receiving capacities.

Most adaptive multimedia schemes have used multi-layered coding as the major supporting mechanism to improve the fairness of the applications. The receiver-driven layered multicast (RLM) [7] is a rate-adaptive protocol for the distribution of layered video using multiple IP-multicast groups. The source generates a certain number of video layers and transmits each layer in a different multicast session. The receivers subscribe to the number of groups they will. The communication is dynamic in the way receivers can dynamically join or leave groups. However, they are limited to the layers the source decides to transmit. Since the receivers do not have any influence on the way the rates of each layer are calculated, these values may not be suited to yield good video quality at the destinations and reasonable network utilization.

The Layered Video Multicast with Retransmission (LVRM) [9] has been proposed for the distribution of layered video over the Internet. LVRM deploys an error recovery scheme using retransmissions with upper bound on recovery time and adaptive playback point. LVRM also implements adaptation to network congestion and heterogeneity using a hierarchical rate control mechanism. In this mechanism, the information about network congestion is exchanged among the sender, the receivers, and some agents in the network. This scheme has the advantage of scalability since each network component stores only the information relevant to itself.

In [5], 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. In this scheme, the source periodically multicasts control packets to the receivers. As these packets traverse the network, intermediate nodes mark them with the amount of available bandwidth. Each receiver copies the packet's content into a returning feedback control packet and sends it back to the source. To avoid feedback implosion at the source, intermediate nodes implement an algorithm to merge the feedback packets returned by the receivers.

The single-loop packet-merging algorithm proposed in [8] is an algorithm to improve fairness for multicast distribution of multi-layered video. The control information scheme is similar to the approach proposed in [5]. 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. In several circumstances, this leads to an important improvement in the global video quality at the receivers since a reduced number of merging procedures are performed and consequently relevant information are not lost.

All of the above techniques have shown to improve the fairness of multicast applications. However, these algorithms consider either the user or the network when computing the video layers, and we argue that both of them should be taken into account.

3 Application-Dependent Quality

In this section, we evaluate the need for other metrics when computing the rates of the video layers other than the global video degradation (Δ), which is the aggregate of all local degradations, given by the difference between the required video rate (r_i) and the actual video rate (r_i^r) received by destination i,

$$\Delta = \sum_{i=1}^{N} \delta_i, \tag{1}$$

where $\delta_i = (r_i - r_i^r), r_i \geq r_i^r$, and N is the number of receivers in the session.

We consider in this paper a feedback control scheme similar to the one presented in [5], where feedback control packets are transmitted by the destinations to the source to indicate the video 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 feedback control packet contains a number of entries that store a video rate (f_i) and the number of destinations that want to receive this rate (f_i^*), as shown in the example of Fig. 1(a). The number of entries in the feedback control packet corresponds to the number of video 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 packet to be sent to the source. Fig. 1(b) shows an example of the merging procedure using the single-loop algorithm for two incoming feedback control packets. In this example, the algorithm has to discard two entries since the two incoming packets together have five entries and the source is supposed to transmit no more than three layers. The first entry of the incoming packet 1 is then combined to the first entry of the incoming packet 2 and entries 2 and 3 of the incoming packet 1 are merged.[1]

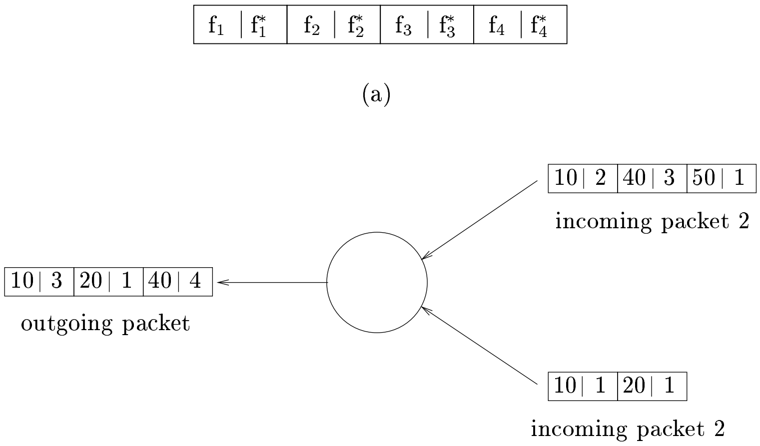

(b)

Fig. 1. (a) Feedback control packet and (b) an example of the merging procedure for a three-layer video.

The algorithms proposed in [5, 8], however, are based only on the video degradation at the receivers. In this paper, we propose the use of other metrics when performing the merging procedure. Consider for example a multicast session with one sender and four receivers as shown in Fig. 2. In this scenario, suppose that the source transmits two video layers and that the merging algorithm uses only the global video degradation at the receivers to perform the merging procedure.

[1] For further details on the merging algorithms, please refer to [5, 8].

Since the rate of the base layer is limited by the slowest receiver, in our example the source transmits the two video layers at L_1=1Mbps and L_2=1Mbps. Note in Table 1 that these values lead to the lowest global video degradation.

Table 1. Video degradation for layer 2 transmitting at 1 and 2.9 Mbps.

	L_2=1Mbps			L_2=2.9Mbps		
Receiver	required	received	degradation	required	received	degradation
r_1	1	1	0	1	1	0
r_2	2	2	0	2	1	1
r_3	2	2	0	2	1	1
r_4	3.9	2	1.9	3.9	3.9	0
Global degradation			1.9 Mbps			2 Mbps

Consider now that the capacity of receiver 4 changes from 3.9 to 4.1Mbps. The new video layers that result in the lowest global degradation are L_1=1Mbps and L_2=3.1Mbps, and the correspondent global video degradation is Δ=2Mbps. The possible value for the video layers are shown in Table 2. Note that for a variation of 5% in the receiving capacity of receiver 4, the video layer 2 has increased its rate of 310%, 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 ~50%), while in the second scenario 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 1Mbps in order to have lower bandwidth utilization even if the global video degradation is greater than the optimum value.

Table 2. Video degradation for layer 2 transmitting at 1 and 3.1 Mbps.

	L_2=1Mbps			L_2=3.1Mbps		
Receiver	required	received	degradation	required	received	degradation
r_1	1	1	0	1	1	0
r_2	2	2	0	2	1	1
r_3	2	2	0	2	1	1
r_4	4.1	2	2.1	4.1	4.1	0
Global degradation			2.1 Mbps			2 Mbps

In next section we will present a qualitative analysis on the use of different metrics to improve the global quality of the application. We propose the global

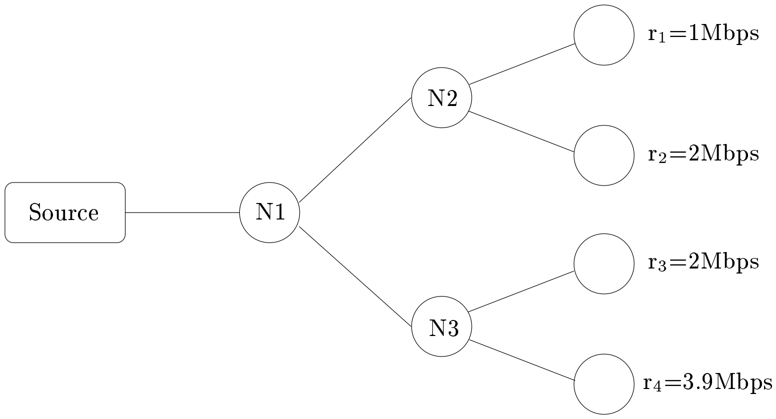

Fig. 2. Example of multicast session.

degradation, the bandwidth utilization, and the user satisfaction as the metrics to be used by the algorithm in order to improve the quality of the communication.

4 Multi-criteria multicast

It is likely that distinct applications have different requirements about the metrics to be prioritized by the algorithm that computes the video layers. 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. The use of different criteria may then be a solution to the problem of multicast delivery of multi-layered applications if it is applied a scheme of weighted parameters, in which each one of the proposed metrics is prioritized depending on the application that is being carried out.

To introduce other metrics, some extensions must be made in the feedback control mechanism. A variation of the feedback control packet must contain other information such as the video degradation and the satisfaction index at the receivers. Besides the video rate and the number of destinations that want to receive this rate (as in the original packet), each entry of the new control packet has two other fields: the video degradation and the satisfaction index.

When performing the merging procedure, the node looks for the entries that have to be discarded and adds its correspondent degradation to the degradation of the entry it is going to be combined with. Consider the same topology of the example shown in Fig. 2 in which the source transmits two video layers. We denote L_i^j the value of f_i at the network element j and δ_i^k the video degradation correspondent to the rate f_i at the network element k. When generating a feedback control packet, receiver d_i sets the field $\delta_1^{d_j}$ to 0 (no degradation since no merging procedure has been performed) and sends the packet to its upstream node. The feedback control mechanism performs the following steps:

At node N_2 (no merging procedure):

$$
\begin{array}{ll}
L_1^{d_1} \rightarrow L_1^{n_2} & \delta_1^{d_1} \rightarrow \delta_1^{n_2} \\
L_1^{d_2} \rightarrow L_2^{n_2} & \delta_1^{d_2} \rightarrow \delta_2^{n_2}
\end{array}
$$

At node N_3 (no merging procedure):

$$
\begin{array}{ll}
L_1^{d_3} \rightarrow L_1^{n_3} & \delta_1^{d_3} \rightarrow \delta_1^{n_3} \\
L_1^{d_4} \rightarrow L_2^{n_3} & \delta_1^{d_4} \rightarrow \delta_2^{n_3}
\end{array}
$$

At node N_1:

$$
\begin{array}{ll}
L_1^{n_2} \rightarrow L_1^{n_1} & (\delta_1^{n_2} + \delta_2^{n_2} + (L_2^{n_2} - L_1^{n_2})) \rightarrow \delta_1^{n_1} \\
L_1^{n_3} \rightarrow L_2^{n_1} & (\delta_1^{n_3} + \delta_2^{n_3} + (L_2^{n_3} - L_1^{n_3})) \rightarrow \delta_2^{n_1}
\end{array}
$$

In this way, the information about the video degradation is carried through the multicast tree and the source can estimate the global video degradation by computing $\Delta = \sum_{i=1}^{E} \delta_i$, where E is the number of entries in the packet. Fig. 3 shows the numerical results after each one of the steps.

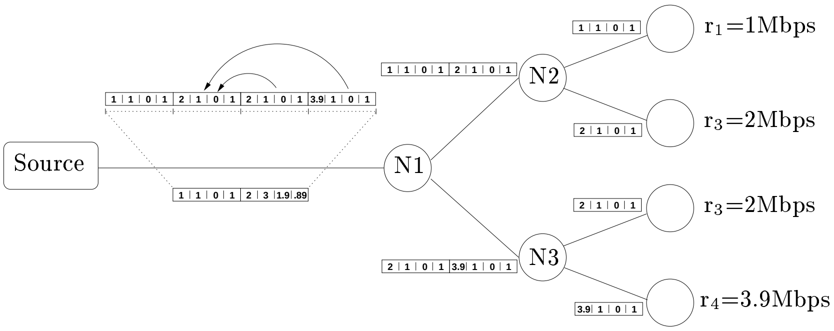

Fig. 3. The merging steps at each one of the nodes.

The degree of bandwidth utilization is computed directly at the source by combining the video rates and their respective number of requests. For a source transmitting C layers, we call weighted bandwidth allocation (Ω) the estimation of the bandwidth allocated throughout the multicast tree,

$$
\Omega = \frac{\sum_{i=1}^{C} (f_i^* \times f_i)}{\sum_{i=1}^{C} f_i^*}. \tag{2}
$$

The satisfaction index (ν) measures the average degree of satisfaction in a group of users. It corresponds to the percentage of the received video with regard to the requested video rate. Consider a scenario in which $(M-1)$ entries $\{[f_2, f_2^*, \delta_2, \nu_2] \ldots [f_M, f_M^*, \delta_M, \nu_M]\}$ are combined to entry $[f_1, f_1^*, \delta_1, \nu_1]$. The resulting entry is given by $[f_1^{'}, f_1^{*'}, \delta_1^{'}, \nu_1^{'}]$, where

$$f_1^{'} = f_1, \tag{3}$$

$$f_1^{*'} = \sum_{i=1}^{M} f_i^*, \tag{4}$$

$$\delta_1^{'} = \sum_{i=1}^{M} \delta_i + \sum_{i=2}^{M} (f_i - f_1) f_i^*, \tag{5}$$

and

$$\nu_1^{'} = \frac{\sum_{i=1}^{M} (f_i^* \times \nu_i^{'})}{\sum_{i=1}^{M} f_i^*}, \tag{6}$$

where, for $j = 2 \ldots M$,

$$\nu_j^{'} = \frac{f_1}{f_j} \times \nu_j. \tag{7}$$

To illustrate that the three metrics proposed in this paper have different comportment and that the combination of them is important to improve the global quality of the communication, we show a simple example of a multicast session with one sender transmitting two video layers and three receivers as shown in Fig. 4. Receivers d_1 and d_3 have fix receiving capacities (r_1=1Mbps and r_3=4Mbps, respectively), and the available bandwidth in the path to destination d_2 varies from r_1 to r_3.

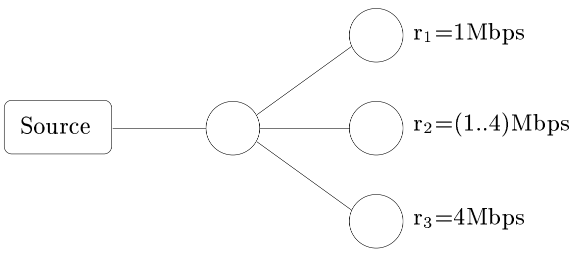

Fig. 4. Example topology.

Fig. 5 shows how the metrics respond to the different possible values of the video layers, i.e., when $\{L_1 = r_1; L_2 = r_2\}$ and $\{L_1 = r_1; L_2 = r_3\}$. The global degradation for the two possibilities are show in Fig. 5(a). Similarly, Fig. 5(b) and Fig. 5(c) depict, respectively, the weighted bandwidth utilization and the estimated satisfaction index. Note that the ideal is to have the minimum global degradation, minimum bandwidth utilization, and the maximum satisfaction index. However, the example shows that these values cannot be simultaneously obtained. In fact, the type of the application must be taken into account when the algorithm to compute the layers is being performed.

5 Conclusions

We have proposed in this paper the use of multiple metrics to compute the most adequate rates for multicast multi-layered applications. We consider the global degradation at the receivers, the bandwidth utilization, and the satisfaction index. We have also argued that one of these metrics should be prioritized to improve the quality of the communication. We showed how the metrics are calculated and how they respond to variations on the network congestion. The proposed criteria lead to a significant improvement on the quality of the application when compared with the classical approaches that consider only the degradation at the destinations.

References

1. D. Saparilla and K. W. Ross, "Optimal streaming of layered video," in *IEEE Infocom*, (Tel-Aviv, Israel), Mar. 2000.
2. T. Wong, R. Katz, and S. McCanne, "A preference clustering protocol for large-scale multicast applications," in *Networked Group Communication - First International COST264 Workshop, NGC'99*, (Pisa, Italy), Nov. 1999.
3. S. Sarkar and L. Tassiulas, "Distributed algorithms for computation of fair rates in multirate multicast trees," in *IEEE Infocom*, (Tel-Aviv, Israel), Mar. 2000.
4. D. Rubenstein, J. Kurose, and D. Towsley, "The impact of multicast layering on network fairness," in *ACM Sigcomm*, (Cambridge, Massachusets, USA), Sept. 1999.
5. B. J. Vickers, C. Albuquerque, and T. Suda, "Adaptive multicast of multi-layered video: Rate-based and credit-based approaches," in *IEEE Infocom*, (San Francisco, CA, USA), Mar. 1998.
6. B. J. Vickers, C. Albuquerque, and T. Suda, "Source adaptive multi-layered algorithms for real-time video distribution," tech. rep., Information and Computer Science – University of California at Irvine, 1999.
7. S. McCanne, V. Jacobson, and M. Vitterli, "Receiver-driven layered multicast," in *ACM SIGCOMM*, (Stanford, CA, USA), Aug. 1996.
8. M. D. de Amorim, O. C. M. B. Duarte, and G. Pujolle, "Single-loop packet merging for receiver oriented multicast multi-layered video," in *International Conference in Computer Communication*, Sept. 1999.
9. X. Li, S. Paul, and M. H. Ammar, "Layered video multicast with retransmissions (LVRM): Evaluation of hierarchical rate control," in *IEEE Infocom*, (San Francisco, CA, USA), Mar. 1998.
10. X. Li, M. H. Ammar, and S. Paul, "Video multicast over the internet," *IEEE Network Magazine*, Mar. 1999.

11. L. Zhang, S. Deering, D. Estrin, S. Shenker, and D. Zappala, “RSVP: A new resource reservation protocol,” *IEEE Network Magazine*, Sept. 1993.
12. B. J. Vickers, M. Lee, and T. Suda, “Feedback control mechanisms for real-time multipoint video services,” *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 3, Apr. 1997.
13. J. C. Bolot and T. Turletti, “Experience with control mechanisms for packet video in the internet,” *ACM Computer Communication Review*, vol. 28, no. 1, Jan. 1998.
14. D. G. Wadington and D. Hutchison, “A general model for QoS adaptation,” in *IEEE/IFIP International Workshop on Quality of Service (IWQoS)*, (Napa, CA, USA), May 1998.
15. C. Diot, C. Huitema, and T. Turletti, “Multimedia applications should be adaptive,” in *HPCS'95*, (Mystic, CN), Aug. 1995.

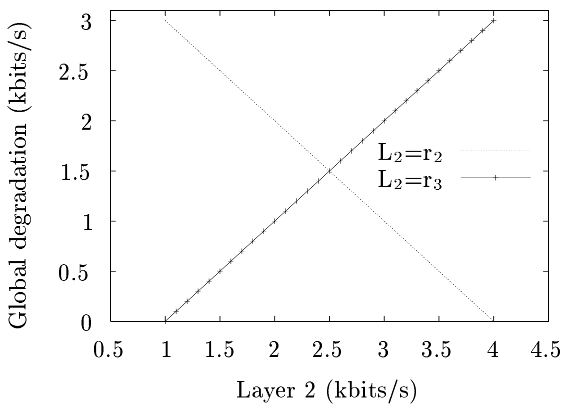

(a)

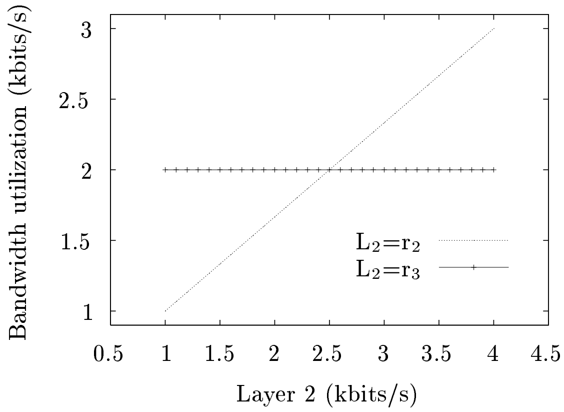

(b)

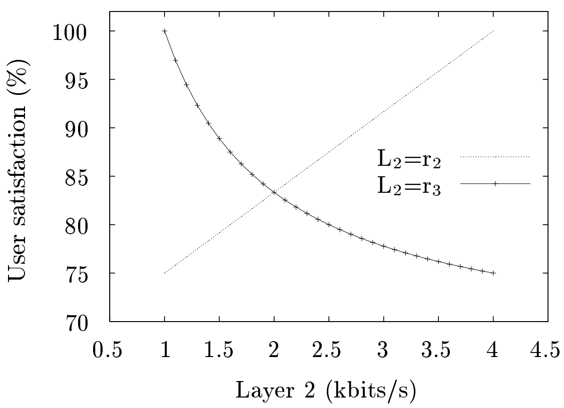

(c)

Fig. 5. (a) Global video degradation, (b) bandwidth utilization, and (c) satisfaction index.