Improving the Delivery Quality of MPEG Video Streams by Using Differentiated Services

Artur Ziviani[1,2], José F. de Rezende[2], Otto Carlos M. B. Duarte[2], and Serge Fdida[1]

[1] LIP6 – Université Pierre et Marie Curie
4, place Jussieu
75252 Paris Cedex 05 - France

[2] Grupo de Teleinformática e Automação – GTA
Universidade Federal do Rio de Janeiro – COPPE/EE
P.O. Box 68504 - 21945-970 - Rio de Janeiro - RJ - Brasil

{Artur.Ziviani,Serge.Fdida}@lip6.fr, {rezende,otto}@gta.ufrj.br

Abstract

The widely adoption of the MPEG format turns it into an attractive way for the distribution of audio and video over the Internet. Nevertheless, the quality perceived by an user of a video transmission may be strongly degradated by undiscriminated discards of packets under the best-effort model. The Differentiated Services (DS) architecture proposes to overcome the limitations for Quality of Service (QoS) provisioning of today's best-effort Internet. This paper evaluates how a DS mechanism can improve the delivery quality of MPEG video streams. We use a drop differentiation mechanism to avoid the error propagation effect of a dropped packet in the hierarchical structure of MPEG systems. Results show that the drop differentiation mechanism improves the delivery quality of MPEG video streams while avoiding the waste of network resources, even under heavy load conditions.

1 Introduction

The best-effort model of the current Internet becomes inadequate to face with the very diverse requirements on network Quality of Service (QoS) of the new multimedia applications. The Internet has perhaps become a victim of its own success [1]. Many of those new applications are based in the delivery of continuous media content [2]. Thus, a key point to the success of such applications is the delivery quality of the video and audio streams. In the case of a video transmission over a packet network, like the Internet, both the encoding and the transmission processes affect the QoS [3]. Video streams are typically compressed before being transported over a network. Compression is needed to reduce the storage and bandwidth requirements of digital video. Its efficiency, however, depends on the video dynamics as well as the underlying compression technique. For a video with constant quality, the video encoder produces a sequence of variable-size compressed frames. The output of the encoder constitutes a variable bit rate (VBR) stream. The transmission of VBR-coded video streams while guaranteeing a required level of QoS is a challenging problem [4, 5].

Artur Ziviani has a scholarship from CAPES/Brazil.
This work was sponsored by FUJB, CNPq, and CAPES/COFECUB.

To address the issue of providing a QoS support for the transmission of flows from applications with different QoS requirements, different approaches are currently being developed, such as adaptive applications [6, 7], QoS-based routing [8, 9], and the Differentiated Services (DS) architecture [10]. The DS proposal provides scalable discriminated services because there is no need of either one state per flow or signaling at every hop. The treatment offered by a DS-compliant node is applied to an aggregation of flows, not to individual flows. Individual flows are classified into aggregations of related flows that will be differentially served within a DS domain. Aggregations are forwarded at each node within a DS domain following a Per-Hop Behavior (PHB). The PHB is selected by a mapping between a codepoint that identifies the aggregation and the expected behavior [11]. Currently, two PHBs are in discussion in the DS research community: the Expedited Forwarding (EF) PHB [12, 13] and the Assured Forwarding (AF) PHB [14].

Although recent papers [15, 16, 17] analyze distinct aspects of differentiated services, the transmission of video traffic over a DS enabled network remains not well explored. This is true specially for the transmission of video streams based on MPEG-4 encoding, which has been pointed out as a standard for future networked interactive video applications [18]. As related research, we may cite the following works. In [19], it is suggested the use of filters in some nodes within the network to selectively drop frames of MPEG streams. The MPEG video stream, however, must be partially decoded at each filter node. Such approach introduces delays that may turn unfeasible an interactive application and also rises scalability issues. A content-based packet video forwarding mechanism based on a relative priority score that is mapped into a proportional loss-rate differentiation mechanism is proposed in [20]. However, the score has to be calculated for each packet corresponding to each Group of Blocks (GOB) in a H.263+ video stream, which may also rise problems to interactive applications.

This paper evaluates how a DS enabled network may support the forwarding of MPEG-4 compressed video streams, providing a drop differentiation of packets to improve the delivery quality of such video streams. In this context, the approach to dealing with dynamic network conditions is located in the network, as suggested in [21]. If under congestion, the network may discard a packet of the video stream. This lack of information may propagate for a number of frames due to the hierarchical structure of MPEG systems that generates video packets of unequal importance. Our goal is to protect from undiscriminated discards the packets that carry the most relevant information, avoiding the error propagation effects on the video sequence. Marking for drop differentiation of a video packets takes place at the video source based on the type of MPEG frame from where the video packet was generated, thus imposing a low overhead on delay. By using this scheme, we aim at improving the transfer of comprehensible information given the available bandwidth in comparison with the best-effort model of the current Internet. We show through simulation that such a scheme can improve the delivery quality of MPEG-4 video streams while reducing the waste of network resources to transport data that becomes useless to the decoder.

The remaining of the paper is organized as follows. Section 2 provides an overview of some characteristics of MPEG-4 video streams. In Section 3, we describe how a drop differentiation scheme based on a DS mechanism may be applied to MPEG-4 video streams. Section 4 presents the adopted evaluation metrics and our simulation results. Some final considerations are discussed in Section 5.

2 MPEG-4 Video Stream

MPEG encodes video as a sequence of frames [22]. Usually, video has a high degree of temporal redundancy, i.e., information in successive frames is highly correlated. Standard MPEG encoders generate three types of compressed frames (I, P, or B). In MPEG-4 terminology, each frame is equivalent to a Video Object Plane (VOP) [23]. An I frame is intra-coded, having no dependence on any other frames. Meanwhile, MPEG uses motion prediction and interpolation techniques to reduce the size of intermediate frames. Two types of motion prediction are used: Forward prediction, where a previous frame is used as a reference for decoding the current frame, and bidirectional prediction, where both past and future frames are used as a reference. This last technique provides a better compression. The encoding of P frames uses forward prediction and the encoding of B frames uses bidirectional prediction. As a consequence, the I frames are normally the largest in size, followed by the P frames, and finally B frames. As the frame rate transmission is constant, the output rate of a MPEG encoder composes a variable bit rate (VBR) sequence. Figure 1 depicts a sample of a frame size sequence.

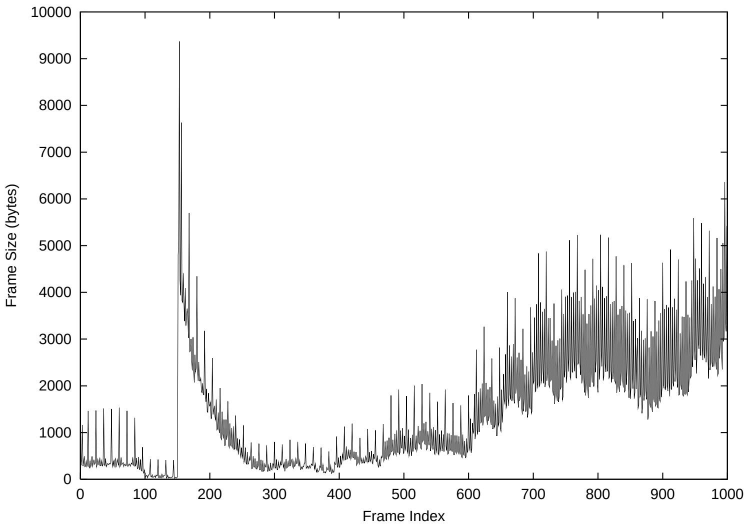

Figure 1: Sample of a frame size sequence.

The video sequence may be decomposed into smaller units which are coded together. Such units are called GOPs (Group of Pictures). Each GOP holds a set of frames or pictures that are in a continuous display order. Usually GOPs are made independently decodable units to facilitate random access. Such GOPs are called closed GOPs as they contain all relevant decoding parameters so that it can be decoded independent of other units. If a GOP needs other GOPs for decoding, it is called an open GOP.

The GOP pattern specifies the number and temporal order of P and B frames between two successive I frames. Such a GOP pattern may be characterized by two parameters: the I-to-I frame distance (N), and the I-to-P frame distance (M). In this paper, we use a MPEG-4 trace file provided by [23], where each entire scene is considered one video object. The trace is composed of open GOP units with $N = 12$ and $M = 3$. This structure and the dependency for decoding of each frame in the GOP is illustrated in Figure 2.

While each displayed frame gets the same amount of playing time and is thus equally important, there is a great difference in both size and information contents of each of the encoded frames. In order to a P or B frame be considered decodable, all the frames they depend on must also be considered decodable. Thus, an incorrect I or P frame propagates problems to all dependent frames and corrupts these as well. In the worst case, a whole GOP may be considered undecodable due to an incorrect I frame, as all other frames in the GOP depend directly

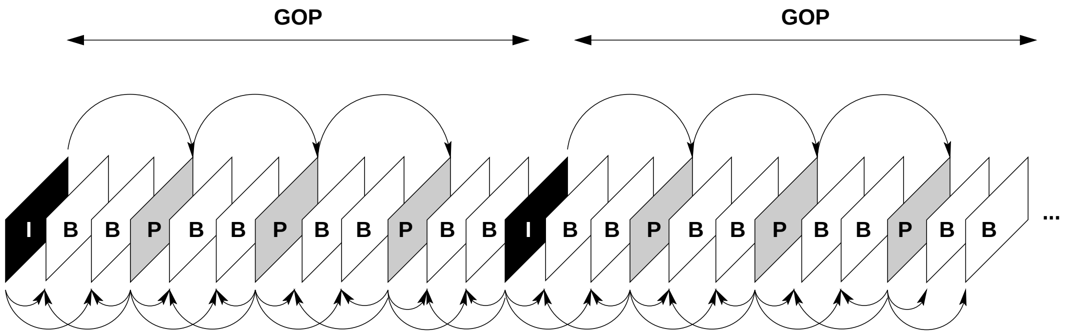

Figure 2: GOP structure of the MPEG-4 trace file.

or indirectly on the I frame. For example, the effects of an undecodable I frame of a GOP with $N = 12$ will persist for 12 frames, or almost 500 ms at a typical frame rate of 25 frames per second, which is quite noticeable to a user.

The hierarchical structure of MPEG encoding with a possible error propagation through the frames imposes a great difficulty on sending MPEG video streams over lossy networks. Small packet loss rates may translate into much higher frame error rates. For example, a 3% packet loss percentage could translate into a 30% frame error [24]. This situation may seriously degradate the user perceived quality at the video reproduction. Moreover, network resources may be wasted with the transmission of information that becomes useless to the receiver. Some of the received data may become useless to the decoder as insufficient frame data is available for decoding due to losses in the network or some frame in which the current frame depends on to be considered decodable is considered undecodable.

3 Differentiated Services Scheme for MPEG-4 Video Streams

In order to improve the transfer of comprehensible information, we associate different levels of drop precedence to packets that carry information of different frames. Under this scheme, increasing drop precedences are associated with packets from I, P, or B frames, respectively. That is, packets transporting fragments of a B frame are more likely to be discarded in a congested router than packets from P frames. Meanwhile, P frame packets have precedence in discard when compared to I frame packets.

This differentiation among packets from different frames may be achieved by mapping those packets to different drop precedences in an AF class [14]. This requires an active queue management algorithm. A multiple level RED [25, 26] mechanism provides the necessary drop discrimination. On our case, we use a three level RED mechanism, that we call REDN3. Packets that carry data from each frame type are associated with a different level in the REDN3 mechanism. The three levels are labeled as red, yellow, and green in order from the highest to the lowest drop precedence. Thus, a packet from an I frame is marked as green, a P frame packet is marked as yellow, and finally a B frame packet is marked as red. The parameter set of REDN3 includes the minimum threshold, the maximum threshold, and the maximum drop probability (min, max, $Pmax$) for each level (red, yellow, and green). By using different parameter settings is possible to provide different packet drop rates to each level as shown in Figure 3. The green curve provides a lower drop probability than the yellow curve, that in turn provides a lower drop probability than the red curve.

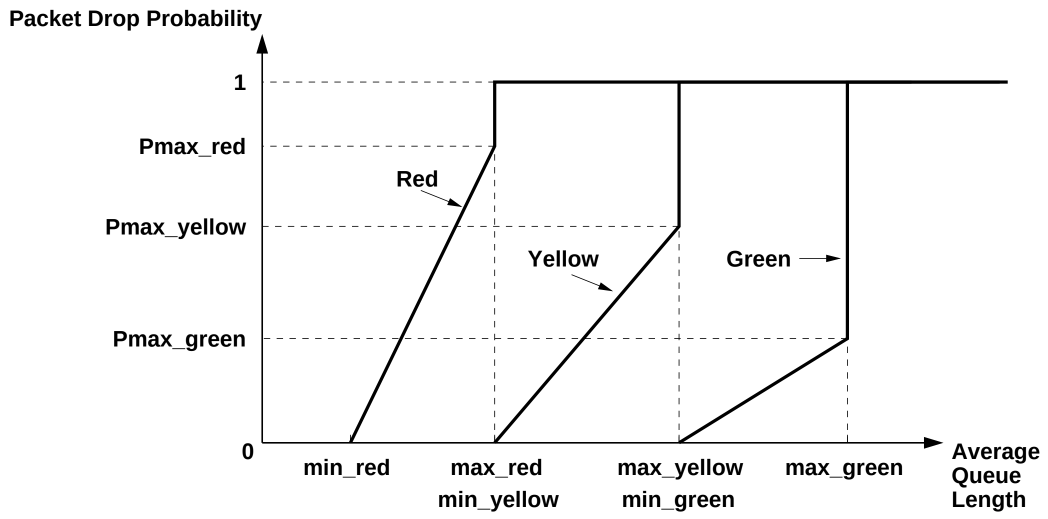

Figure 3: Parameters of a RED mechanism with 3 levels (REDN3).

4 Simulation Results

We implemented our simulation model in a modified ns-2 simulator [27]. The performance evaluation is carried out through simulations at the network scenario depicted in Figure 4. A video source generates packets based on a public available MPEG-4 frame trace file [28]. The frames are encoded in open GOP units with $N = 12$ and $M = 3$. The details of the frame trace file generation are described in [23]. We used the Starwars-IV movie sequence, encoded with high quality and MPEG-4 encoding, which has a mean bit rate of 0.28 Mbps and peak rate of 1.9 Mbps. Each frame is fragmented into packets of 200 octets. The video packets are transmitted from node V0 to node V1. A number of TCP sources are uniformally distributed among the nodes T0, T1, and T2, and send packets to nodes D0, D1, and D2, respectively. The TCP background sources use the TCP-Reno implementation. The TCP cross traffic is generated by FTP applications, which are active during the whole simulation, with packets of 1500 octets.

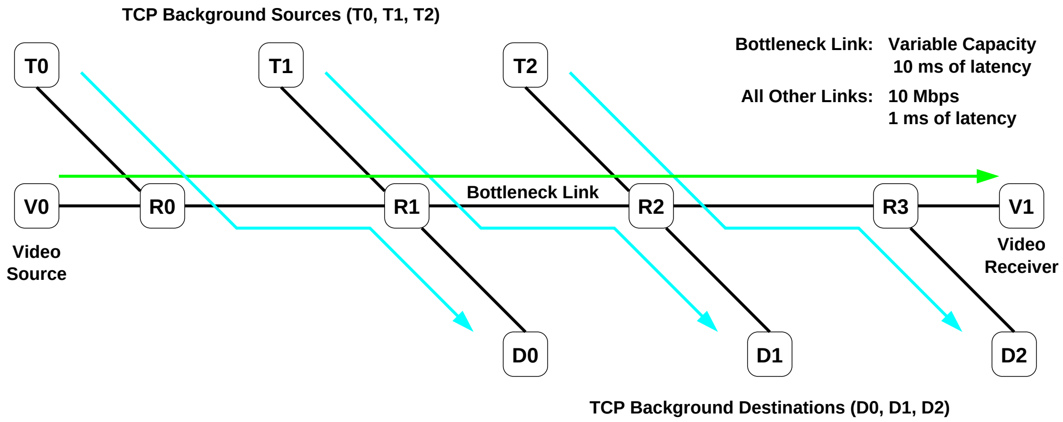

Figure 4: Network scenario.

The routers R0, R1, and R2 use either the classical Drop Tail or a RED mechanism with three levels for active queue management. The threshold parameters of the REDN3 mechanism for the three levels are configured as percentages of the total queue length ($qlen$), which is kept equivalent to 50 packets. The REDN3 parameters that we use in the simulations are ($0.2 * qlen$, $0.4 * qlen$, 0.1), ($0.4 * qlen$, $0.6 * qlen$, 0.05), and ($0.6 * qlen$, $0.8 * qlen$, 0.025) for red, yellow, and green packets, respectively. This set discriminates red packets, progressively avoiding the discard of yellow and green packets. The packets of the TCP flows are considered as red packets.

4.1 Evaluation Metrics

Our goal is to improve the quality of the video received by the user and avoid the wasting of network resources due to the delivery of information that becomes useless at the receiver. We compare the REDN3 scheme with the classical Drop Tail present in today's Internet. In order to evaluate the transmission of a MPEG-4 video stream over the network scenario, we consider the following metrics: Fraction of decodable frames and useless data received. Fraction of decodable frames reports the number of decodable frames over the total number of frames sent by the sender. Useless data received reports the fraction of the data received that is useless to the decoder. The decoder may have useless data from partial received frames or from frames that can not be decoded because they depend on other undecodable frames. We also measure the goodput of the TCP flows to see how they are affected by the Drop Tail or REDN3 queue managements in the intermediate routers. All results are presented with 99% confidence intervals.

The adopted criterion to consider a frame decodable is defined as follows. At least a fixed fraction dt (decodable threshold) of the data in each frame must be received in order to the frame be considered decodable. Further, a frame is considered decodable if and only if all the frames it depends on are also considered decodable. Therefore, when $dt = 1.0$, the decoder is completely intolerant to losses, i.e., one lost packet is enough to consider a frame undecodable. In the same manner, with $dt = 0.75$, 25% of the data from a frame may be absent due to losses in the network and the frame is still considered decodable. This criterion allows the consideration of a number of error resilience and error concealment features of MPEG-4 [18, 29]. Such features permit that a system tolerates losses in the network up a certain level and still display lossy video bitstreams with the best possible video quality.

4.2 Variation in the Load from Background Sources

Our first set of experiments evaluates the transmission of a MPEG-4 video stream through the network scenario while varying the number of TCP connections that acts as background traffic. The bottleneck link capacity is fixed at 1.024 Mbps. The TCP sources are uniformally distributed among the nodes T0, T1, and T2. Thus, one third of them compete with the video stream for bandwidth at the bottleneck link. We compare the delivery quality of the video stream when using the REDN3 mechanism and Drop Tail for different values of decodable threshold (dt).

As the number of TCP connections increases, the percentage of decodable frames is seriously affected when the Drop Tail is used, as shown in Figure 5(a). Even with $dt = 0.50$, where the receiver tolerates the lost of up to 50% of the information in each frame, there is a significant degradation on the quality of the delivered video. This behavior is due to undiscriminated discards of packets caused by Drop Tail. Losses in a P or I frame may turn other frames undecodable, heavily degradating the perceived quality. Meanwhile, the REDN3 mechanism performs better and provides a smoother degradation in the delivery quality when facing with heavy load conditions, specially if the system tolerates some losses ($dt = 0.75$ or $dt = 0.50$).

Figure 5(b) shows that a great amount of network resources may be wasted when using Drop Tail. As the REDN3 mechanism improves the decodable frame percentage, it avoids the waste of network resources in delivering information that will become useless at the receiver. Again, this effect is strongly emphasized if a number of losses is tolerated.

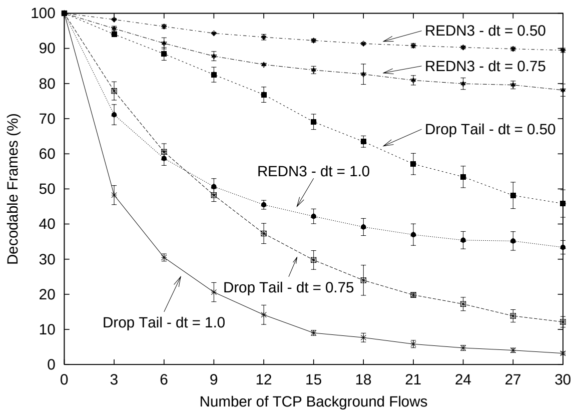

(a) Decodable Frames

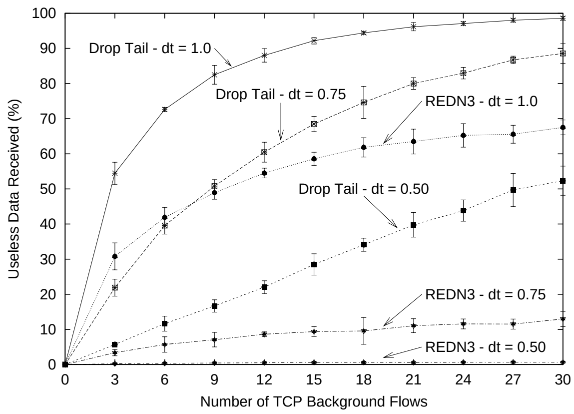

(b) Useless Data Received

Figure 5: Varying the number of TCP background flows.

Providing a smaller drop rate to packets that carry information of P and I frames avoids wasting network resources and improves the decodable frames percentage. This is achieved by penalizing the TCP connections that share links with the video stream. Such effect is shown in Figure 6, which depicts the total goodput of the TCP connections in the bottleneck link. Using REDN3, the total goodput of TCP is smaller than under Drop Tail. By imposing a small penalty on the goodput of TCP connections is possible to achieve a significant gain in the delivery of comprehensible information of the video stream.

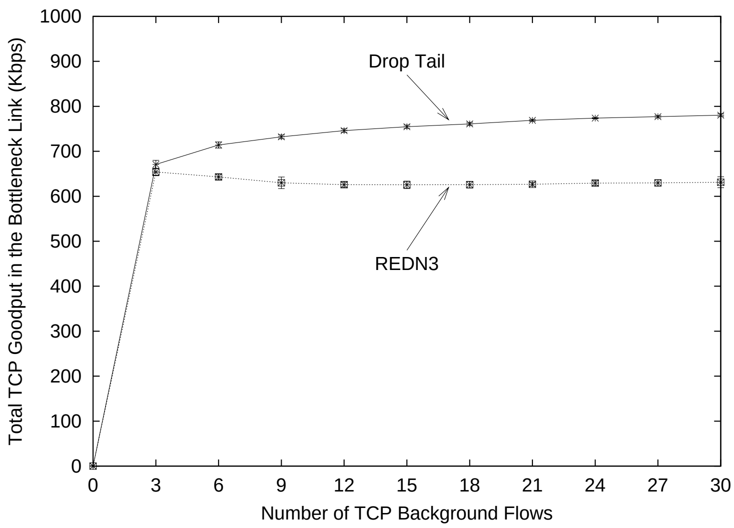

Figure 6: Total goodput of TCP background flows in the bottleneck link.

4.3 Variation in the Bottleneck Capacity

The second set of experiments evaluates the performance of REDN3 and Drop Tail with different bottleneck capacities. There are 15 concurrent TCP connections, 5 at each TCP source node. The results for the delivery of the video stream, both percentage of decodable frames and useless data received, are presented in Figure 7.

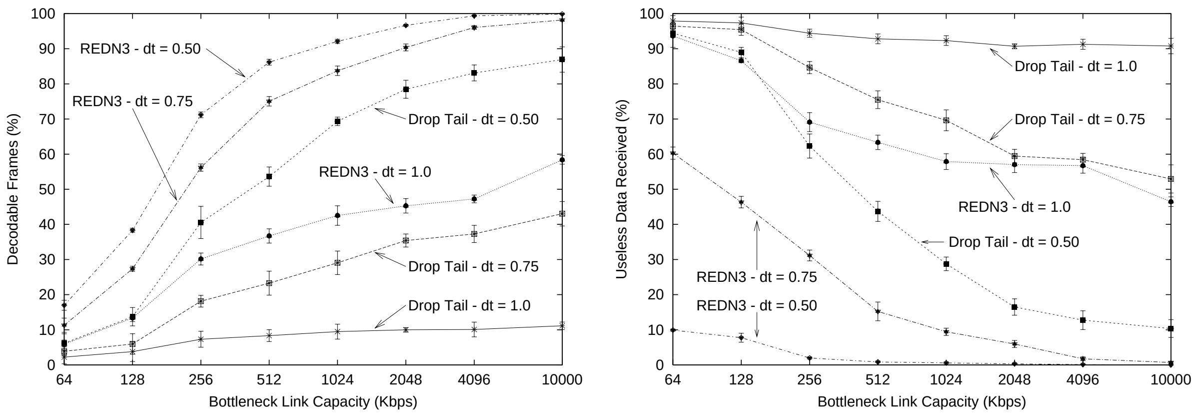

(a) Decodable Frames

(b) Useless Data Received

Figure 7: Varying the bottleneck link capacity.

In a situation of scarce bandwidth, the delivery quality of the video stream is significantly affected, even for low values of dt as shown in Figure 7(a). As more bandwidth is made available, the REDN3 mechanism profits better of the increasing amount of resources. As the available bottleneck bandwidth approaches 10 Mbps (no bottleneck at all), the REDN3 provides more than 95% decodable frames percentage for $dt = 0.75$. Note that under the same conditions, Drop Tail provides approximately just 45%. Even if a reasonable amount of bandwidth is available, the undiscriminated drops from Drop Tail impose a great degradation on the delivery quality.

At lower capacity bottleneck links, there is a great waste of network resources when using Drop Tail as the few packets that can reach the receiver are not enough to decode much frames, even with $dt = 0.50$ as depicted in Figure 7(b). As more bandwidth becomes available, the tolerance to losses improves the Drop Tail performance. Nevertheless, the REDN3 avoids the wasting of network resources by protecting the packets with most significant information. Under REDN3, packet drops occur mainly to packets from B frames, that are typically smaller. That is why, with a low bottleneck link, REDN3 presents a small decodable frames percentage, but also a small useless data received percentage for $dt = 0.50$. A great number of B frames are considered undecodable, what decreases the decodable frames percentage. However, the amount of useless data received that composes undecodable B frames is small when compared to the total data received, as B frames are normally smaller in size.

The price for the better performance of the REDN3 scheme is payed by the TCP connections that get restrict access to the available bandwidth. Figure 8 presents the total share of bandwidth obtained by the TCP connections in the bottleneck link. The protection granted to packets from I and P frames imposes more losses on the TCP connections and, as a consequence, they reduce their transmission rates. Thus, the TCP connections get a smaller share of the available bandwidth in the bottleneck link when using the REDN3 scheme.

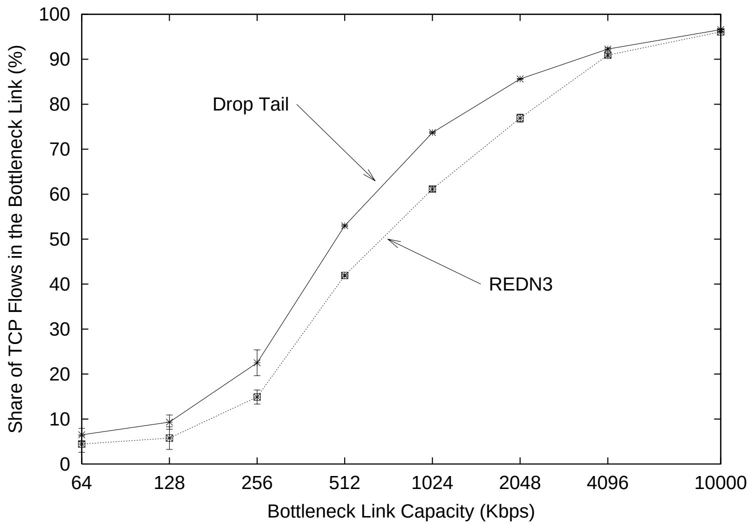

Figure 8: Share of TCP flows in the bottleneck link.

4.4 Multiple Video Sources

Until this point, we evaluated the performance of the REDN3 mechanism handling packets from one video source that competes with some TCP cross traffic. In the next set of experiments, we observe the delivery quality of one reference video stream while a number of other video streams also uses the REDN3 mechanism. In other words, there are packets marked as yellow or green sharing the same links at the network scenario that do not belong to the video stream being evaluated. The bottleneck link capacity is kept at 1.024 Mbps and there are 15 TCP background connections. The number of video sources ranges from 1 to 30. All video sources use the same Starwars-IV trace file with random starting times to avoid synchronized transmissions. Figure 9 shows the decodable frames percentage for the different frame types considering two levels of tolerance to losses, represented by $dt = 1.0$ and $dt = 0.50$.

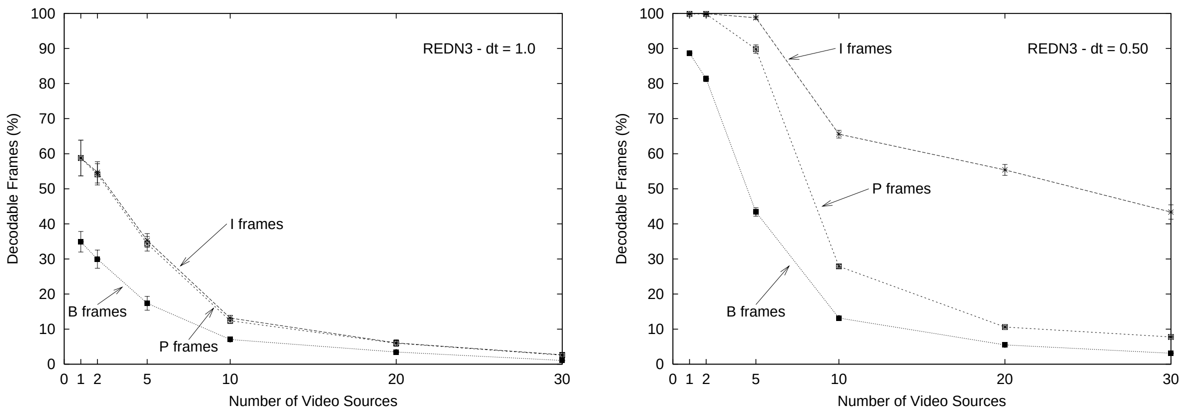

(a) $dt = 1.0$

(b) $dt = 0.50$

Figure 9: Varying the number of video sources.

As the packets from B frames compete directly against the TCP background flows and have the highest drop precedence, B frames present the worst performance, as would be expected. In Figure 9(a) is shown that for an intolerant system ($dt = 1.0$), the P and I frames show a similar performance, as the discard of just one packet can turn one entire frame undecodable. Under a system that tolerates a certain number of packet losses ($dt = 0.50$), the different frame types obtain significantly different performances. This is a result of the progressive protection from discards of packets from B, P and I frames provided by the REDN3 mechanism.

5 Conclusion

In this paper we evaluated a differentiated services scheme to improve the delivery quality of MPEG-4 video streams. The performance of a drop differentiation mechanism that avoids the propagation of errors through the hierarchical structure of MPEG encoding was evaluated in terms of the percentage of decodable frames delivered and useless data received. The improvement in the delivery quality was obtained by imposing on the TCP traffic higher drop rates and, as a consequence, a smaller goodput. Simulation results showed that the scheme under study provides a simple, but quite efficient, method to improve the delivery quality of MPEG-4 video streams while reducing the waste of network resources, even under heavy load conditions. From Figure 5(a), in a system that tolerates some losses ($dt = 0.75$) and under increasing load conditions, the studied drop differentiation mechanism provided a decodable frames percentage that levels off around 80%. The same metric under the Drop Tail mechanism continuously degradates to about 10% within the same range of simulated load conditions.

The marking of packets to different drop precedences is directly done based on the type of frame they come from, imposing a low overhead on delay. Such approach makes this scheme specially suitable for interactive applications. The study of alternatives based on the RED mechanism becomes attractive as this mechanism is being widely adopted on today's commercial routers. Thus, this paper contributes with guidelines to administrators that, in the near future, may want to configure their routers to offer a differentiated services support for MPEG-4 video streams through the RED mechanism.

References

[1] L. Mathy, C. Edwards, and D. Hutchison, "The Internet: A global telecommunications solution?", *IEEE Network*, vol. 14, pp. 46–57, July 2000.

[2] D. A. Turner and K. W. Ross, "Continuous media e-mail on the Internet: Infrastructure inadequacies and a sender-side solution", *IEEE Network*, vol. 14, pp. 30–37, July 2000.

[3] O. Verscheure, X. Garcia, G. Karlsson, and J.-P. Hubaux, "User-oriented QoS in packet video delivery", *IEEE Network*, vol. 12, pp. 12–21, Nov. 1998.

[4] G. Karlsson, "Asynchronous transfer of video", *IEEE Communications Magazine*, vol. 24, pp. 118–126, Aug. 1996.

[5] M. Krunz, "Bandwidth allocation strategies for transporting variable-bit-rate video traffic", *IEEE Communications Magazine*, vol. 37, pp. 40–46, Jan. 1999.

[6] C. Diot, C. Huitema, and T. Turletti, "Multimedia applications should be adaptive", in *HPCS'95*, Aug. 1995.

[7] M. D. de Amorim, O. C. M. B. Duarte, and G. Pujolle, "Multi-criteria arguments for improving the fairness of layered multicast applications", in *Proceedings of Broadband Communications, High Performance Networking, and Performance of Communication Networks - Networking'2000*, no. 1815 in Lecture Notes in Computer Science, pp. 1–10, May 2000.

[8] E. Crawley, R. Nair, B. Rajagopalan, and H. Sandick, "A framework for QoS-based routing in the Internet", *Internet RFC 2386*, Aug. 1998.

[9] L. H. M. K. Costa, S. Fdida, and O. C. M. B. Duarte, "Distance-vector qos-based routing with three metrics", in *Proceedings of Broadband Communications, High Performance Networking, and Performance of Communication Networks - Networking'2000*, no. 1815 in Lecture Notes in Computer Science, pp. 847–858, May 2000.

[10] S. Blake, D. L. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss, "An architecture for differentiated services", *Internet RFC 2475*, Dec. 1998.

[11] K. Nichols, S. Blake, F. Baker, and D. L. Black, "Definition of the differentiated services field (DS field) in the IPv4 and IPv6 headers", *Internet RFC 2474*, Dec. 1998.

[12] V. Jacobson, K. Nichols, and K. Poduri, "An expedited forwarding PHB", *Internet RFC 2598*, June 1999.

[13] G. Armitage, A. Casati, J. Crowcroft, J. Halpern, B. Kumar, and J. Schnizlein, "A revised expression of the expedited forwarding PHB", *Internet Draft*, Nov. 2000. `draft-ietf-diffserv-efresolve-00`.

[14] J. Heinanen, F. Baker, W. Weiss, and J. Wroclawski, "Assured forwarding PHB group", *Internet RFC 2597*, June 1999.

[15] A. Ziviani, J. F. de Rezende, and O. C. M. B. Duarte, "Towards a differentiated services support for voice traffic", in *Proceedings of the IEEE GLOBECOM'99*, vol. 1, pp. 59–63, Dec. 1999.

[16] C. Dovrolis and P. Ramanathan, "Proportional differentiated services, part II: Loss rate differentiation and packet dropping", in *Proceedings of IWQoS'2000*, June 2000.

[17] N. Semret, R. R.-F. Liao, A. T. Campbell, and A. A. Lazar, "Peering and provisioning of differentiated Internet services", in *Proceedings of IEEE INFOCOM'2000*, Mar. 2000.

[18] R. Koenen, "MPEG-4: Multimedia for our time", *IEEE Spectrum*, vol. 36, pp. 26–33, Feb. 1999.

[19] M. Hemy, U. Hengartner, P. Steenkiste, and T. Gross, "MPEG systems streams in best-effort networks", in *Proceedings of the 9th Packet Video Workshop - PV'1999*, Apr. 1999.

[20] J. Shin, J.-W. Kim, and C.-C. J. Kuo, "Content-based packet video forwarding mechanism in differentiated services networks", in *Proceedings of the 10th Packet Video Workshop - PV'2000*, May 2000.

[21] S. Bhattacharjee, K. L. Calvert, and E. W. Zegura, "Network support for multicast video distribution", Tech. Rep. GIT-CC-98/16, Georgia Institute of Technology, 1998.

[22] M. Krunz and S. K. Tripathi, "On the characterization of VBR MPEG streams", in *Proceedings of the ACM SIGMETRICS'97*, May 1997.

[23] F. H. P. Fitzek and M. Reisslein, "MPEG-4 and H.263 video traces for network performance evaluation", Tech. Rep. TKN-00-06, Technical University Berlin, Oct. 2000.

[24] J. M. Boyce and R. D. Gaglianello, "Packet loss effects on MPEG video sent over the public Internet", in *Proceedings of the ACM Multimedia 98*, Sept. 1998.

[25] S. Floyd and V. Jacobson, "Random early detection gateways for congestion avoidance", *IEEE/ACM Transactions on Networking*, vol. 1, pp. 397–413, Aug. 1993.

[26] D. D. Clark and W. Fang, "Explicit allocation of best effort packet delivery service", *IEEE/ACM Transactions on Networking*, vol. 6, pp. 362–373, Aug. 1998.

[27] K. Fall and K. Varadhan, "NS – notes and documentation", tech. rep., The VINT Project, Jan. 1999.

[28] F. H. P. Fitzek and M. Reisslein, "MPEG-4 and H.263 video traces for network performance evaluation", `http://www-tkn.ee.tu-berlin.de/research/trace/trace.html`.

[29] Y. Wang and Q.-F. Zhu, "Error control and concealment for video communication: a review", *Proceedings of the IEEE*, vol. 86, pp. 974–997, May 1998.