ANALYSIS OF LOSS EPISODES FOR VIDEO TRANSFERS OVER UDP

Velibor Markovski and Ljiljana Trajković*
School of Engineering Science
Simon Fraser University
Vancouver, British Columbia, Canada
e-mail: {vmarkovs, ljilja}@cs.sfu.ca
http://www.ensc.sfu.ca/research/cnl

Keywords: Communications, Statistical analysis, Discrete simulation, Network protocols, Long-range dependence.

Abstract

Understanding loss patterns in a network is important for achieving the desired quality of service during transfers of multimedia traffic. In this paper, we study the time dependence of loss patterns for video transfers over UDP in a congested packet network. Of particular interest are consecutive packet losses because they contribute to the prolonged loss episodes. We use simulation tools and scenarios to show that the length of these loss episodes increases with the increase of the average utilization levels, and that lengthy loss episodes contribute significantly to the overall loss patterns. Under fixed average utilization, during the times of higher congestion, we observe longer loss episodes and shorter loss episode distances.

1 Introduction

High quality real-time multimedia applications in Internet Protocol (IP) and Asynchronous Transfer Mode (ATM) networks have stringent requirements in terms of packet delay, delay jitter, and packet loss. Certain applications either fail to function or have poor performance if the delay or the delay jitter is larger than a certain threshold. Loss performance is also one of the critical issues for delivering Quality of Service (QoS) in multimedia networks. Throughout this text, we use the term loss performance or loss characteristics to refer to the loss behavior of a video connection. While loss characteristics sometimes refer only to loss probabilities, they may also refer to the loss patterns for particular or multiplexed connections. In ATM networks, for example, the narrower term, loss probability, is used as a QoS metric for connection admission control calculations [7, 14]. Even though loss probability does not describe the overall loss behavior for the admitted connections, it is an important parameter. The loss probability threshold, set by a specified QoS parameter, should not be surpassed by the loss probabilities of any of the admitted connections. In both ATM and IP networks, loss characteristics impact the quality of the established video and/or voice connections, or the throughput for bulk data transfers. Loss characteristics for these interactive applications depend not only on the loss probability, but also on the exhibited loss patterns or loss distributions [2, 3].

We use *ns* simulator [11] to explore the loss behavior of video connections over User Datagram Protocol (UDP). We consider only buffer overflows as a main reason for the loss. We disregard other possible reasons for loss, such as route changes, link failures, or wireless links. Because buffer overflows are the main loss generators in wireline networks, our analysis may provide insights into the overall loss behavior of the network.

In Section 2 we describe the simulation scenarios, the topology, and the choice of the traffic sources. In Section 3 we define loss episodes and loss episode distances. Section 4 presents and discusses the quantitative simulation results. We conclude with Section 5.

*This research was supported by the NSERC Grant 216844-99 and the BC Advanced Systems Institute Fellowship.

2 Simulation scenario

We consider a network topology consisting of n sources generating video traffic and feeding a common router buffer that is connected to a traffic sink. The topology is specified using a Tcl script, the basic input required by *ns* simulator [11]. By selecting the number of video sources n in the script, we generate automatically the network topology. The speeds of the links that connect the sources to the router are chosen to be larger than the peak bit-rate of the employed traffic sources. In this manner, the traffic of the video sources arrives unaltered to the router buffer. In our simulation scenarios, the link speed between the sources and the router is 10 Mbit/s, and the link speed between the router and the sink is 44.736 Mbit/s. The protocol used to transfer the packets is UDP, with packet sizes of 200 bytes. The router uses first-in-first-out (FIFO) buffer. The maximum buffer size B varies. We present results for buffer sizes ranging from $B = 25$ to $B = 200$ kB (in steps of 25 kB), which correspond to maximum queueing delay of 4.58 to 36.62 msec, respectively. The chosen values are reasonable for video communications. The values below 10 msec are within the range of latencies specified for high-end interactive video [5].

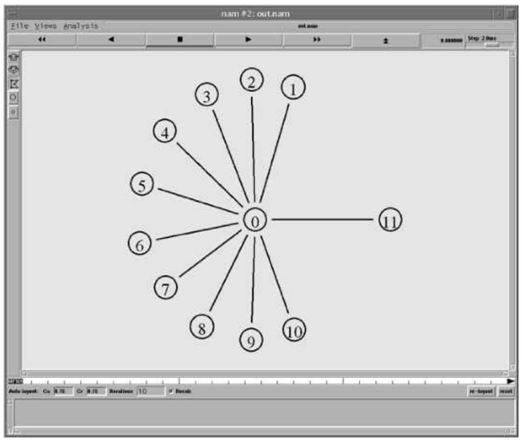

Figure 1: Screen dump from network animator *nam* for topology with $n = 10$ sources. The topology is generated automatically. Node 0 is the router, nodes 1 to 10 correspond to the traffic sources, and node 11 is the traffic sink. The sources generate UDP traffic according to a supplied trace.

Figure 1 is a screen dump from network animator *nam* [10], a companion software tool to *ns* [11]. *Nam* is used for visualization of the topology layout, packet level animation, and data inspection. In this example, a network with $n = 10$ sources, one router, and a sink is shown. The router is numbered 0, the sources are numbered from 1 to 10, and the sink is numbered 11. The number of sources in the topology shown in Figure 1 is for illustration purposes only. In our simulations, we experiment with larger values of n. While using the fixed values for the link speeds, we vary n in order to cover average utilization levels of the output link ranging from 50 % to over 90 %.

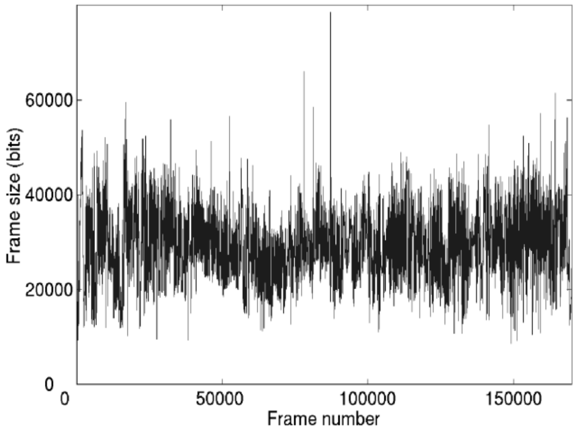

Figure 2: Traffic pattern of video trace (time scale = 1/24 sec). Frame sizes (in bits) are shown for more than 170,000 video frames. This corresponds to approximately two hours of an MPEG-1 video trace.

We use genuine video traces as a traffic source [6, 12] and present the results for sources generating MPEG-1 encoded *Star Wars* [6] and *Talk show* [12] video traces. The traffic pattern for the *Star Wars* video is shown in Figure 2. In order to avoid synchronization of the generated traffic by the sources, each source starts at random point within the traffic trace. If the end of the trace is reached before the end of the simulation, the source continues from the beginning of the trace. Detailed information about the characteristics of the *Star Wars* and *Talk show* traces is given in [6, 12].

3 Packet loss and loss episodes

With the simulation scenario described in Section 2, we conduct a series of experiments by varying the number of sources n and the maximum buffer size B. All the lost (dropped) packets were recorded as the output of the *ns* simulations. We generate a set of (time, loss) pairs, where time represents the time of the packet generation, and loss has a binary value of 1 or 0 depending whether the packet was lost or not [1]. Besides the packet generation time, we also record the instances when loss occurs at the common buffer.

Loss distance is defined as "difference in sequence numbers of two successively lost packets that may or may not be separated by successfully received packets" [8]. Loss episode begins with a lost packet if the previous packet was successfully received. For example, if packets with sequence numbers 1, 4, and 6 are successfully received and packets 2, 3, and 5 are lost, then the first loss episode begins with packet 2 and ends with packet 3, while the second loss episode begins and ends with packet 5. The length of the first loss episode is two packets, and the length of the second loss episode is one packet. The loss distance between the two loss episodes is two packets. Let $lpkt_i(t_1, t_2)$ be the number of lost packets belonging to loss episodes with length i, and $Lpkt(t_1, t_2)$ be the total number of lost packets within the time interval t_1 and t_2. Let $lep_i(t_1, t_2)$ be the number of loss episodes with length i packets and let $Lep(t_1, t_2)$ be the total number of loss episodes within (t_1, t_2). Let L be the length of the loss episode with the maximum length. In our experiments, we determine the maximum length of the loss episodes as well as the distribution of the loss episodes with lengths ranging from 1 to L. The following holds:

$$\sum_{i=1}^{L} lep_i(t_1, t_2) = Lep(t_1, t_2)$$
$$\sum_{i=1}^{L} lpkt_i(t_1, t_2) = Lpkt(t_1, t_2). \qquad (1)$$

4 Simulation results

4.1 Textured dot strip plots of loss patterns

In order to examine the patterns of $lep_i(t_1, t_2)$ and $lpkt_i(t_1, t_2)$, we first examine the overall loss pattern of a simulation run with the *Star Wars* trace, 80 sources, and buffer size of 100 kB. According to the output link speed (44.736 Mbit/s), this case corresponds to the maximum buffer delay of 18.31 msec.

We used textured dot strip plots to provide the initial insight into the loss patterns. The idea of the textured dot strip plots is to depict the densities of the observed variable. In our case, the observed variable is the time when the loss occurs. We obtained these plots by using *Xgobi*, a tool for interactive graphics and data analysis [13]. Textured dot strip plots have been used also for illustrations of traffic patterns [15].

Figure 3 is a textured dot strip plot for the lost packets in a particular simulation case. It depicts

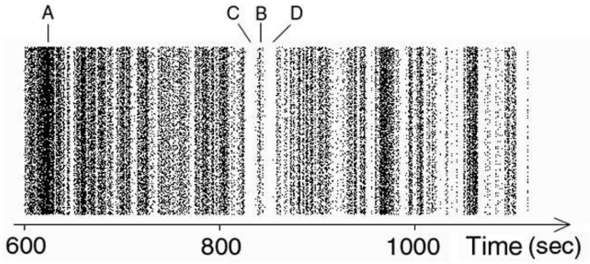

Figure 3: Textured dot strip plot of packet loss instances at the common buffer from a simulation run with $n = 80$ sources and buffer size $B = 100$ kB.

the loss instances after an initial warm-up period of 600 sec. The plot shows all the lost packets from the aggregated stream of n sources, and thus represents the losses occurring during buffer overflow. The plot provides information about the density of loss within the trace and thus reflects the periods of higher or lower congestion. For example, there is a period of high congestion between 615 and 625 sec (band A), and a period of lower congestion between 836 and 846 sec (band B). There are also long periods with no losses, as for example between 826 and 836 sec (band C), or between 846 and 854 sec (band D).

4.2 Distribution of lost packets, loss episodes, and loss episode distances

The number of lost packets from loss episodes with length i is equal to $lpkt_i(t_{start}, t_{end}) = i \cdot lep_i(t_{start}, t_{end})$. Contribution of lost packets for buffer sizes 50 and 100 kB is shown in Figure 4.

Similar pattern holds for the duration of loss episodes. Loss episodes may have duration of one or more packets. The contribution of loss episodes of length i (in packets) to the overall number of loss episodes is equal to $lep_i(t_{start}, t_{end})/Lep(t_{start}, t_{end})$, and it is shown in Figure 5. t_{start} is the start of the observation interval (in our case 600 sec) and t_{end} is the end of the observation interval (1200 sec). We can see from Figure 5 that single losses contribute significantly to the overall number of loss episodes. However, as we increase the number of traffic sources (and, accordingly, the utilization of the output link), the number of single losses decreases, while there is a significant increase in the contribution of the lengthier loss episodes.

It can be observed from Figures 4 (top) and 5 (top), for the case of 80 sources and with buffer size of 50 kB (9.16 msec), that the single losses represent 50 % of the total number of loss episodes. Neverthe-

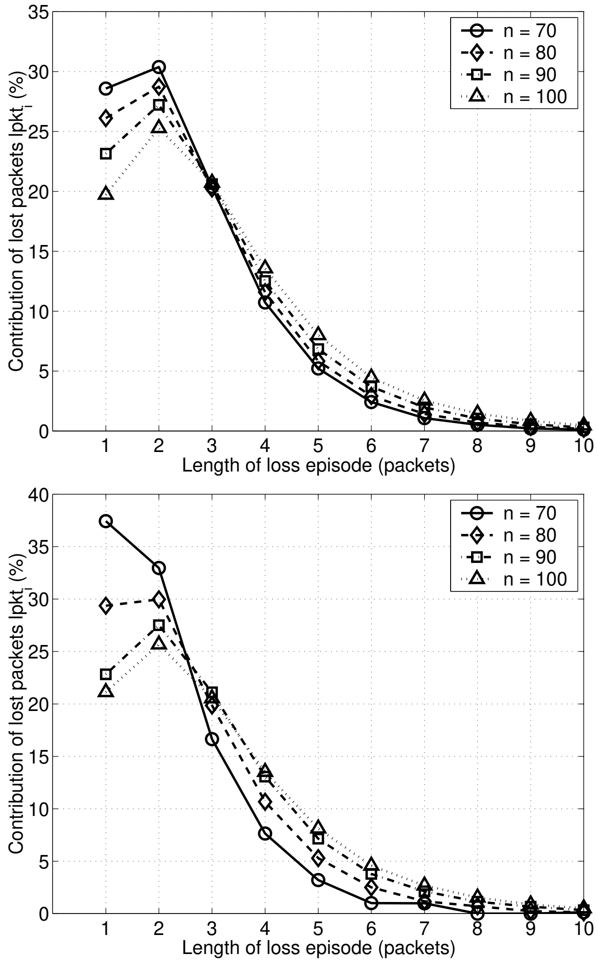

Figure 4: *Star Wars* trace, the number of sources $= n$, the buffer size $B = 50$ kB (top) and $B = 100$ kB (bottom). Contribution of lost packets from loss episodes of various lengths to the overall number of loss episodes.

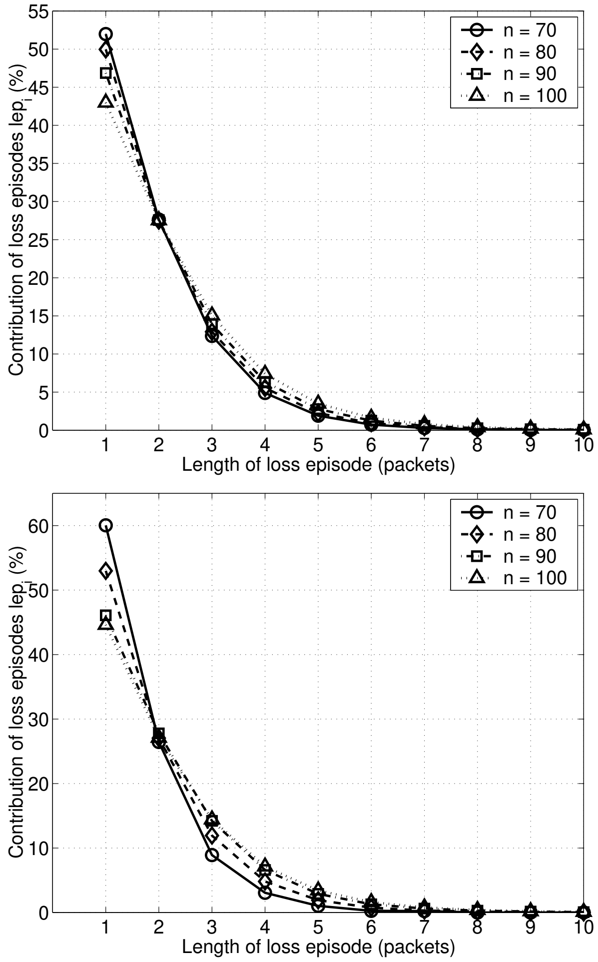

Figure 5: *Star Wars* trace, the number of sources $=$ n, the buffer size $B = 50$ kB (top) and $B = 100$ kB (bottom). Contribution of loss episodes of various lengths to the overall number of loss episodes.

less, they represent only 26 % of the total number of lost packets. For the case with 80 sources and buffer size 100 kB (18.31 msec) shown in Figures 4 (bottom) and 5 (bottom), single packet loss episodes contribute to only 29 % of the total number of lost packets, although the loss episodes of length one represent 53 % of all the loss episodes.

We also calculated the contribution of loss episodes and lost packets for the *Talk show* trace. The corresponding graphs for the case with buffer size of 100 kB are shown in Figure 6. The results are qualitatively similar to those for the *Star Wars* trace.

The number of arrived and lost packets at the router, the average loss at the router buffer, as well as the utilization of the router output link for an extended observation interval of 1140 sec ($t_{start} = 60$ sec and $t_{end} = 1200$ sec) are given in Table 1. We would like to point out that the short-range dependent (or summary) statistics does not capture the behavior of the observed loss process [9, 15].

For the cases shown in Figures 4–6, the values for the lost packets and loss episodes are obtained for the total observation interval of 600 sec. If narrower intervals were selected, one could observe the loss behavior particular to various bands of a loss trace obtained with fixed buffer size and fixed number of sources.

We observe two 10-second time bands A and B

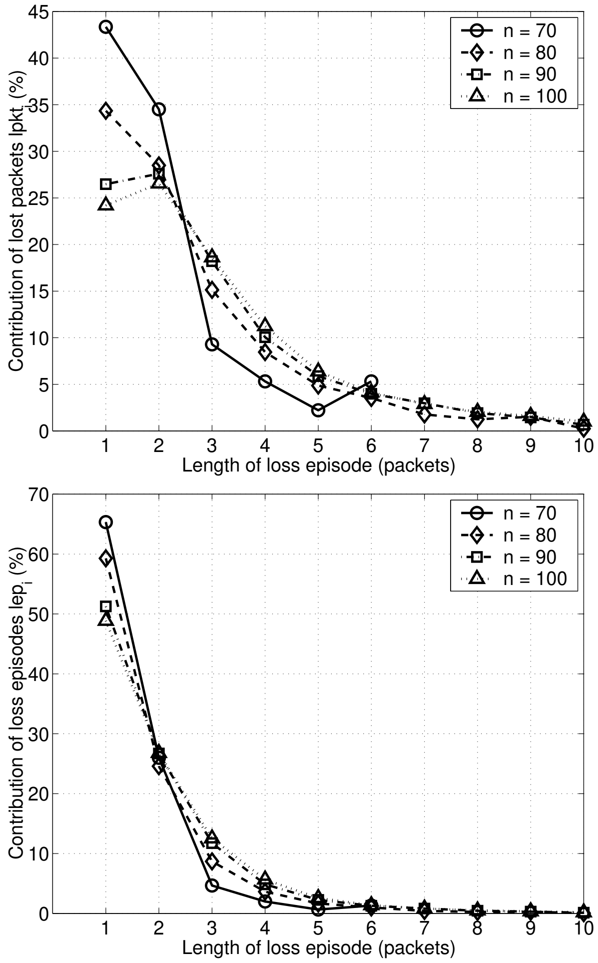

Figure 6: *Talk show* trace, the number of sources = n, the buffer size B = 100 kB. Contribution of lost packets (top) and loss episodes (bottom) of various lengths to the overall number of loss episodes.

Statistics	n	*Star Wars*	*Talk show*
arrived	70	18,286,114	19,067,183
lost		128,699	131,565
utilization		57.4 %	59.8 %
average loss		0.70 %	0.69 %
arrived	80	20,909,729	21,810,341
lost		241,639	251,057
utilization		65.6 %	68.4 %
average loss		1.16 %	1.15 %
arrived	90	23,507,614	24,517,493
lost		523,987	546,675
utilization		73.8 %	76.9 %
average loss		1.16 %	1.15 %
arrived	100	26,161,169	27,275,529
lost		886,256	924,108
utilization		82.0 %	85.6 %
average loss		3.39 %	3.39 %

Table 1: Average loss, utilization and total number of arrived and lost packets at the router for the *Star Wars* and *Talk show* trace. The buffer size B = 100 kB.

from Figure 3. Our intuitive reasoning was that during the times of high congestion, the lengths of the loss episodes were larger and the loss episode distances were shorter. Table 2 presents the summary for the loss episode lengths and loss episode distances for a simulation with buffer size of 100 kB (equivalent to the maximum queueing delay of 18.31 msec) and $n = 80$ sources (utilization = 68.1 %). The variable i in the table denotes the length of the loss episode in packets. We can see from Table 2 that the mean loss episode distance is smaller for the highly congested band A. Within this band, single losses accounted for only 26 % of the overall number of lost packets. In the episode of lower congestion, single losses have much bigger share of the overall losses and account for about 44 % of all lost packets in band B.

	Contribution of lost packets of length i to the overall number of lost packets				
Band A	25.98	28.75	19.77	12.59	6.56
Band B	43.84	28.99	16.30	7.25	3.62
Band A	Mean loss episode distance = 148				
Band B	Mean loss episode distance = 1145				

Table 2: Contribution of lost packets (in %) from loss episodes of length i, and mean loss episode distances in bands A (higher congestion) and B (lower congestion), shown in Fig. 3. Loss episode distance is given in units of packet transmission times. For our choice of the packet size and link speed, the packet transmission time is 35.77 μsec.

4.3 Effect of buffer sizes

In order to analyze the loss behavior for various maximum queueing delays, we obtained the loss episodes and loss distances for buffer sizes ranging from 25 kB (4.58 msec) to 200 kB (36.62 msec). With the increase of utilization, the frequency of single losses decreases and lengthier loss episodes become dominant.

The pattern for the relative frequency of loss episodes with length i remains similar for other values of the buffer size.

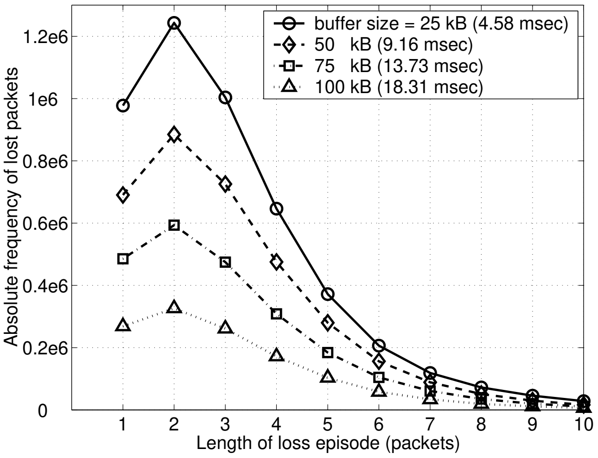

Figure 7: Total number of lost packets from loss episodes of length i packets. *Star Wars* trace, the number of sources $n = 100$. For better clarity of the figure, only the results for buffer sizes 25, 50, 75, and 100 kB are shown.

Figure 7 depicts the total number of lost packets from loss episodes of various lengths, using simulation with $n = 100$ sources. As expected, the total number of loss episodes decreases with the increase of buffer size. We can clearly see from Figure 7 the decrease of the number of losses of any length as the buffer increases. We can observe that for a fixed number of sources and variable buffer size, the largest number of lost packets emanates from loss episodes with length $i = 2$ packets. In addition, Figure 8 shows that the change of buffer size does not affect the percentage of loss episodes with length 2. We also noticed from Figure 5, that for a fixed buffer size, the percentage of loss episodes with length $i = 2$ remains approximately the same for various utilizations of the output link. In other words, the lines depicting the relative loss episode frequency are intersecting for the length of loss episode $i = 2$ packets. The intersecting point at $i = 2$ packets was also obtained for the other values of the buffer size used in our experiments. The same effect appeared in the simulation with the *Talk show* trace.

We observed two phenomena. First, for variable buffer sizes and fixed number of sources, the largest number of lost packets comes from loss episodes of length $i = 2$ packets. Second, for fixed buffer size and variable number of sources and vice versa, the percentage of loss episodes of length 2 remains the same. We are currently investigating this issue and we suspect that the explanation might be the coding scheme used for the video streams.

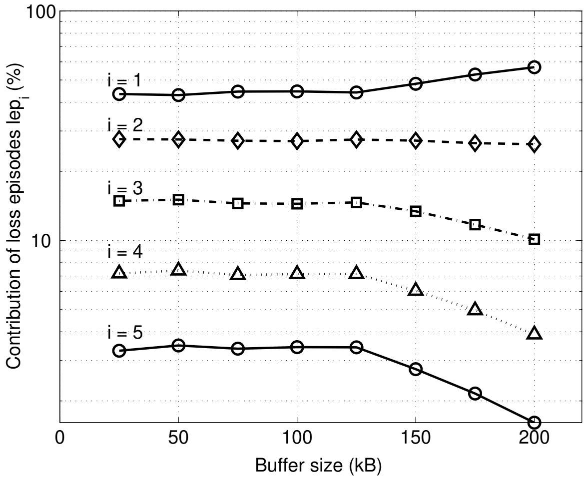

Figure 8: Percentage of loss episodes of length i ranging from 1 to 5 packets vs. the buffer size. *Star Wars* trace, the number of sources $n = 100$.

Apart from the relative and absolute frequency of loss episodes and lost packets, of particular importance is the maximum length of the loss episodes L, i.e., the maximum number of consecutively lost packets. This value depends both on the buffer size and the utilization of the output link. Smaller buffer size and higher utilization means increase in the maximum length of the loss episodes. This dependency is shown in Figure 9.

4.4 Effect of transport protocols

In addition to the loss analysis for the UDP transfers, we have also studied the loss behavior in an environment of mixed UDP and TCP sources. The graphs of loss episodes vs. length of the loss episode are qualitatively similar to the results obtained using UDP sources only. One important difference is that the maximum length of the loss episodes are always longer for the UDP transfers. This was expected, having in mind the feedback mechanism of TCP. However, when wavelet analysis is applied, a qualitative difference appears for the loss patterns of UDP and TCP transfers.

Figure 10 depicts the log of the spectrum $\Gamma_x(2^{-j}\nu_0)$ at a given scale j. Γ_x can be obtained from the discrete wavelet coefficients $d_{j,k}$. Detailed information on wavelet analysis of UDP loss patterns is reported in [16]. For the UDP transfers, we can see from Figure 10 (top) that there is a breakpoint at scale $j = 10$

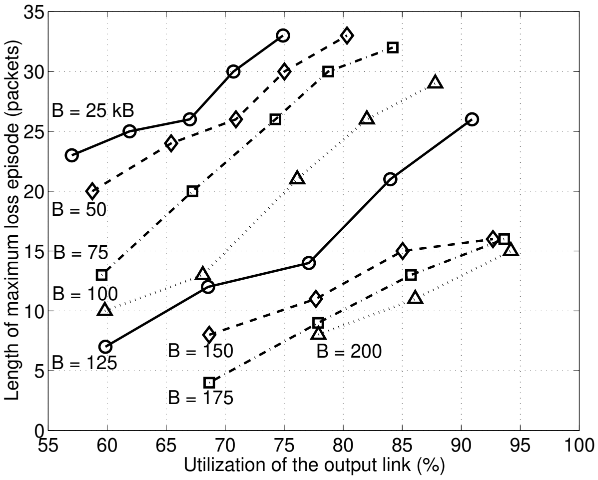

Figure 9: Length of the maximum loss episode vs. the utilization of the output link. The buffer is a parameter and ranges from 25 kB (4.58 msec) to 200 kB (36.62 msec).

(corresponding to 1,024 msec in our analysis) beyond which linear relationship between $log_2(\Gamma)$ and level j is evident. On the other hand, the linear relationship between $log_2(\Gamma)$ and j for a range of scales $[j_1, j_2]$ indicates presence of a long-range dependent behavior. In other words, for the time scales in the linear region (in our case, coarser time scales, larger than 1,024 msec), the UDP loss shows long-range dependent characteristics. For the TCP loss shown in Figure 10 (bottom), we have not observed linear segments in the region of time scales we considered in our analysis (1 msec – 65.536 sec). One possible approach to discover the reasons for the observed phenomena of the TCP and UDP loss scaling behavior is to perform local scaling analysis [4], which was successfully applied for network traffic.

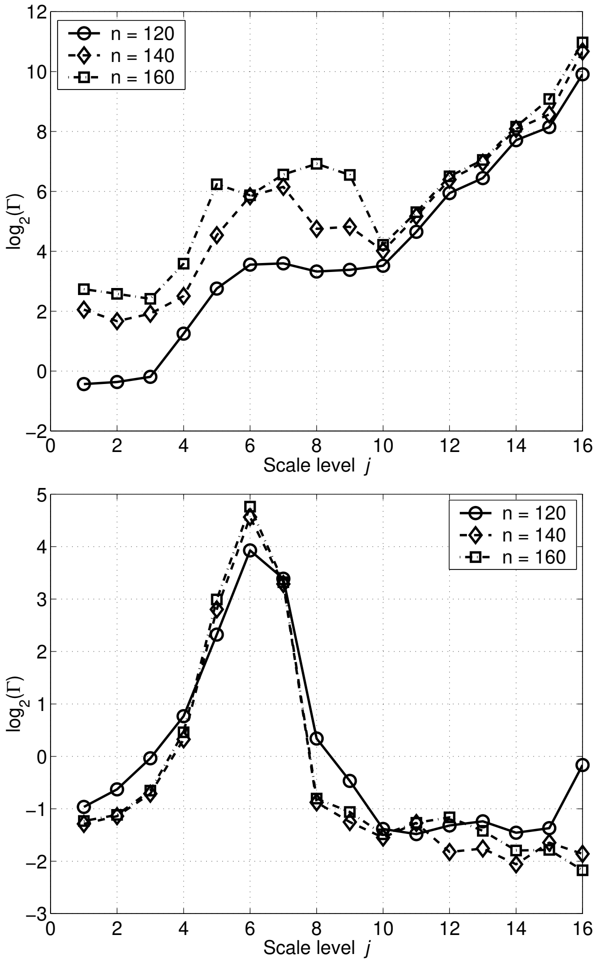

Figure 10: $log_2(\Gamma)$ vs. j plot of packet loss process for buffer size $B = 69$ kB. UDP transfers (top) and TCP transfers (bottom)

5 Concluding remarks

We used the *ns* simulator to perform quantitative analysis of the loss, loss episode lengths, and loss episode distances for a simple topology of n video sources, one router, and a sink. Our results indicate that, as the utilization of the network increases, lengthier loss episodes have more significant contribution to the overall loss, and that the contribution of loss episodes with lengths larger than one is quite high, especially for the periods of high congestion.

For a fixed number of sources and buffer sizes, we observed that the episodes of lower congestion are characterized with more frequent single loss episodes and with larger loss episode distances. For example, in a scenario with 100 sources, buffer size of 100 kB, and 82 % utilization of the output link, loss episodes with two or more consecutively lost packets constitute almost 80 % of the overall number of lost packets.

The maximum length of the loss episodes depends on the buffer size and the number of sources feeding the buffer. For example, with buffer size of 25 kB (4.58 msec) and 74 % utilization of the output link, the maximum length of the loss episode reached 33 packets.

We are currently exploring the influence of time scales on the loss patterns in a congested network using various traffic sources. In addition, we are investigating the reason for the observed qualitatively different behavior of UDP and TCP loss patterns on various time scales.

References

[1] G. Almes, S. Kalidindi, and M. Zakeuskas, "A one-way packet loss metric for IPPM," RFC 2680, IETF, Sept. 1996.

[2] J-C. Bolot and A. Vega Garcia, "The case for FEC-based error control for packet audio in the Internet," *Proc. ACM Int. Conf. Multimedia*, Sept. 1996.

[3] M. S. Borella, D. Swider, S. Uludag, and G. B. Brewster, "Internet packet loss: measurement and implications for end-to-end QoS," *Proc. Int. Conf. Parallel Processing*, Minneapolis, MN, Aug. 1998, pp. 3–15.

[4] A. Feldmann, A. C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A study of the role of variability and the impact of control," *Proc. ACM SIGCOMM'99*, Aug. 1999, pp. 301–313.

[5] M. Fester, "Performance issues for high-end video over ATM," Aug. 1995: http://www.cisco.com /warp/public/cc/sol/mkt/ent/atm/vidat_wp.htm.

[6] M. Garrett and W. Willinger, "Analysis, modeling and generation of self-similar VBR video traffic," *Proc. ACM SIGCOMM'94*, London, U.K., Aug. 1994, pp. 269–280.

[7] E. Knightly and N. Shroff, "Admission control for statistical QoS: Theory and practice," *IEEE Network*, vol. 13, no. 2, pp. 20–29, Mar. 1999.

[8] R. Koodli and R. Ravikanth, "One-way loss pattern sample metrics," Internet draft, IETF, Oct. 1999, ftp://ftp.ietf.org/internet-drafts/draft-ietf-ippm-loss-pattern-02.txt .

[9] W. Leland, M. Taqqu, W. Willinger, and D. Wilson, "On the self-similar nature of Ethernet traffic (extended version)," *IEEE/ACM Trans. Networking*, vol. 2, pp. 1–15, 1994.

[10] Nam network animator, 1999: http://www-mash.cs.berkeley.edu/nam .

[11] ns-2 network simulator, 1999: http://www-mash.cs.berkeley.edu/ns .

[12] O. Rose, "Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems," *Proc. 20th Annual Conference on Local Computer Networks*, Minneapolis, MN, Oct. 1995, pp. 397–406.

[13] D. S. Swayne, D. Cook, and A. Buja, "XGobi: Interactive dynamic data visualization in the X window system," *J. Computational and Graphical Statistics*, vol. 7, no. 1, pp. 113–130, Mar. 1998.

[14] Lj. Trajković and A. Neidhardt, "Effect of traffic knowledge on the efficiency of admission-control policies," *ACM Comput. Commun. Rev.*, vol. 29, no. 1, pp. 5–34, Jan. 1999.

[15] W. Willinger, M. S. Taqqu, R. Sherman and D. V. Wilson, "Self-similarity through high-variability: statistical analysis of Ethernet LAN traffic at the source level," *IEEE/ACM Trans. Networking*, vol. 5, no. 1, pp. 71–86, 1997.

[16] F. Xue, V. Markovski, and Lj. Trajković, "Wavelet analysis of packet loss in video transfers over UDP," *IC'2000*, June 2000, Las Vegas, Nevada, USA.