Simulation and analysis of packet loss in video transfers using User Datagram Protocol

Velibor Markovski, Fei Xue, and Ljiljana Trajković*
School of Engineering Science, Simon Fraser University, Vancouver, British Columbia, Canada E-mail: {vmarkovs, fxue, ljilja}@cs.sfu.ca.

Abstract. Understanding packet loss patterns in Internet Protocol (IP) networks is important for achieving the desired quality of service in multimedia transfers. In this paper, we study the loss patterns in video transfers using User Datagram Protocol (UDP) in a congested packet network. We use trace-driven *ns-2* simulations to collect packet loss traces in networks, and we apply wavelet analysis to investigate the behavior of packet loss on various time-scales. We show that time-scales are essential for understanding loss behavior and that packet loss exhibits long-range dependence over the coarser time-scales.

Keywords: Packet loss, user datagram protocol, video transfer, long-range dependence, wavelets, time-scales.

1. Introduction

High quality multimedia applications in IP (Internet Protocol) networks have stringent requirements in terms of packet loss, delay, and delay jitter. In IP networks, loss characteristics impact the quality of the established video and/or voice connections, and affect the throughput of bulk data transfers. These loss characteristics cannot be captured by specifying only the loss probability and require a detailed study of loss patterns [8].

In our studies, we use *ns-2* simulator [1] and genuine traffic traces to explore the loss behavior [18, 25] of video connections over User Datagram Protocol (UDP). We use wavelet analysis [9] to study the packet loss behavior on various time-scales. Wavelet analysis proved to be a useful tool for analyzing the scale-dependent properties of packet loss via the coefficients of wavelet decomposition.

In Section 2, we first describe the simulation scenario, the topology, and the choice of the traffic sources. Section 3 presents the quantitative simulation results of packet loss. We use Gilbert model to characterize the loss episodes of the collected loss data. In Section 4, we perform the wavelet analysis of packet loss and discuss our findings. We conclude with Section 5.

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

2. Simulation scenarios

We consider a network consisting of n sources that generate video traffic. By selecting the number of video sources n in the input scripts for the *ns-2* simulator, we generate the network topology automatically. In our simulation scenarios, the link speed between the sources and the router is 10 Mbps, and the link speed between the router and the sink is 44.736 Mbps. The UDP protocol is used to transfer the packets. A typical packet size in UDP transfers is 200 Bytes. The router uses a FIFO buffer with a DropTail queue management policy. The maximum buffer size B varies from B = 128 to B = 1,024 packets (in steps of 128 packets), which corresponds to maximum queueing delays of 4.58 msec to 36.62 msec, respectively. These values for delays are reasonable for video communications where applications can sustain latencies between 5 msec and 300 msec. Delays below 10 msec are within the range of latencies specified for high-end interactive video [12].

A simple network scenario with an arbitrary number of sources n generating traffic and feeding a common router buffer R connected to a traffic sink D is shown in Fig. 1. We consider various scenarios for delivering a video trace to the traffic destination D using UDP transfers. In the topology shown in Fig. 1, we experimented with large values of traffic sources in order to cover utilization levels from 50% to 90%.

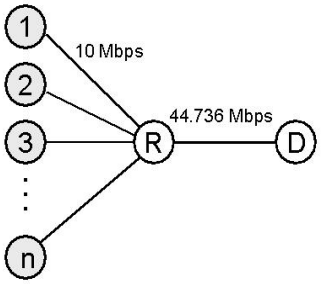

Figure 1. Simple network topology used in simulations of packet loss.

The simulations are trace-driven with genuine, MPEG-1 encoded *Star Wars* [13] and *Talk show* [19] video traces. The *Star Wars* trace is publicly available in ASCII and in *ns-2* format [2]. The *Star Wars* trace consists of more than 170,000 video frames, coded with 24 frames per second, and is approximately two hours long. The traffic pattern is shown in Fig. 2. The *ns-2* traffic trace file was produced from the original *Star Wars* trace by dividing each frame into packets. These packets are transmitted evenly spaced over approximately 20 msec. The source then remains idle for the remainder of the 1/24 second frame time [2]. The *Talk show* trace [3] is a frame size trace extracted from MPEG-1 sequences encoded with 25 frames per second. The trace consists of 40,000 frames, which is equivalent to approximately half an

hour long trace. In order to avoid synchronization of the traffic generated by the sources, each source begins sending packets at a random time instance within the traffic trace. If the end of the trace is reached before the end of the simulation, the source continues sending the packets from the beginning of the traffic trace. (I order to avoid packets resulting from the beginning scenes of the *Star Wars* movie, our simulations actually started several minutes after the beginning of the original traffic trace.)

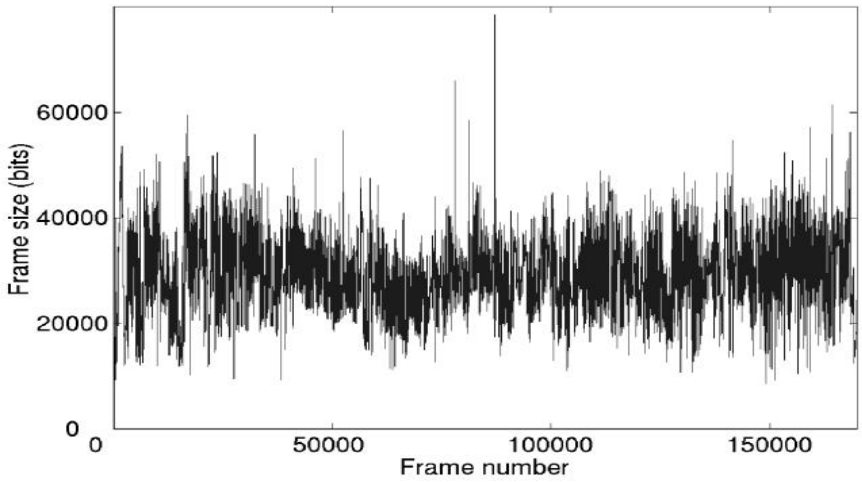

Figure 2. Traffic pattern of video trace (one video frame = 1/24 sec). Frame sizes (in bits) are shown for more than 170,000 video frames.

Table I. Encoder and trace parameters summary.

Encoder parameters	*Star Wars*	*Talk show*
Encoder input (pel)	480x504	384x288
Resolution (bits/pel)	8	8
Pattern	IBBPBBPBBPBB	IBBPBBPBBPBB
Group of Pictures (GOP) size	12	12
Frame rate (frames/sec)	24	25
Number of video frames	174,136	40,000
Peak rate (Mbps)	4.446	2.669
Mean rate (kbps)	374.4	363.4
Peak/mean rate ratio	12.23	7.13
Coefficient of variation	1.165	1.136

The encoder and trace parameters summary is given in Table I. The traffic pattern shown in Fig. 2 and data in Table I indicate that the traffic is bursty.

This is typical for a trace obtained by MPEG-1 video encoding, which produces variable bit-rates depending on the scene activity and complexity in the original uncompressed movie. In the case of the *Star Wars* trace, the peak bit-rate (calculated over one frame interval) is more than 12 times larger than the average bit-rate (calculated over the two hour trace duration).

3. Packet loss and loss episodes

Using the simulation scenario described in Section 2, we conducted a series of experiments in which we varied the number of sources n and the maximum buffer size B. We used *ns-2* simulator to record lost (dropped) packets and the corresponding times when loss occurred.

We begin loss analysis by determining the overall packet loss probability during a video transfer period in the simple network shown in Fig. 1. Our first simulation study dealt with the loss rates at the router buffer for the simulations with UDP traffic. The packet loss rates at the router buffer for the simulations with the *Star Wars* and *Talk show* traces are shown in Table II. The buffer size is varied from 25.6 kbytes to 204.8 kbytes, corresponding to maximum queueing delays of 4.58 msec to 36.6 msec, respectively. The number of sources is $n = 100$ and the packet size is 200 bytes. The loss rate is higher for the simulations with the *Star Wars* trace, due to the bursty nature of the traffic trace, as indicated in Table I). For both traces, the increase of the buffer size results in a reduction of the loss rate. However, a larger buffer implies an increase of the maximum queueing delay. The loss rates for the simulations with the two video traces are shown in Fig. 3.

Table II. Loss rates for the UDP simulations with *Star Wars* and *Talk show* traces.

$n = 100$ Packet size $= 200$ bytes		*Star Wars* Packets arrived: 27,451,164		*Talk show* Packets arrived: 25,866,487	
Buffer size (kbytes)	Max queueing delay (msec)	Lost	Loss rate (%)	Lost	Loss rate (%)
25.6	4.58	4,512,616	16.4	1,524,967	5.90
51.2	9.16	3,279,716	11.9	845,210	3.27
76.8	13.73	2,121,275	7.73	461,867	1.79
102.4	18.31	1,130,829	4.12	223,571	$8.64 \cdot 10^{-1}$
128	22.89	577,287	2.10	93,431	$3.61 \cdot 10^{-1}$
153.6	27.47	268,873	$9.79 \cdot 10^{-1}$	24,765	$9.57 \cdot 10^{-2}$
179.2	32.05	79,563	$2{,}90 \cdot 10^{-1}$	1,406	$5.44 \cdot 10^{-3}$
204.8	36.62	13,621	$4.96 \cdot 10^{-2}$	0	0

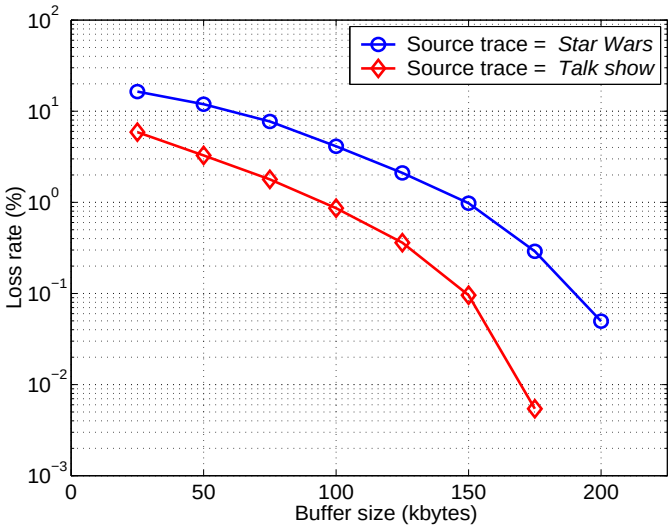

Figure 3. Loss rates at the router buffer for the UDP simulations with *Star Wars* and *Talk show* traces. The number of sources $n = 100$ and the packet size $= 200$ bytes.

The loss rates shown in Table II are calculated for the observation interval [$t_{start} = 60$ sec, $t_{end} = 1200$ sec]. Loss rates may also be calculated over a shorter interval [t_{start}, $t_{start} + T$], by sliding t_{start} to cover the observation interval, as shown in Fig. 4 (top) for $T = 5$ sec, $t_{start} = 60$ sec, and $t_{end} = 600$ sec. The time interval between 200 sec and 300 sec is enlarged and shown in Fig. 4 (bottom) for $T = 1$ sec and $T = 5$ sec. The decrease in the interval over which the loss rate is calculated shows that there exists a variability in the loss rates that cannot be captured by specifying only the long-term average loss rate. Loss rate represents an aggregate measure of the underlying loss process, and it does not provide sufficient information about the loss distribution within the observed time interval.

The loss rates shown in Fig. 4 refer to the loss at the router buffer. The loss rates for a particular source (label number 50) is shown in Fig. 5 (top). The aggregate arrival rate calculated on a time-scale of 5 sec is shown in Fig. 5 (bottom).

Our analysis of the nature of the packet loss indicated that packet loss is bursty. Therefore, we examine the loss patterns in more details, focusing on the loss episodes. A packet *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 have been successfully received, and if packets 2, 3, and 5 have been lost, then the first loss episode begins with packet number 2 and ends with packet number 3, while the second loss episode begins and ends with packet number 5. Loss episodes consist of *lost packets*. A loss episode is of length k, if it consists of k consecutively lost packets. For the above example, the length of the first loss episode is two packets, and the length of

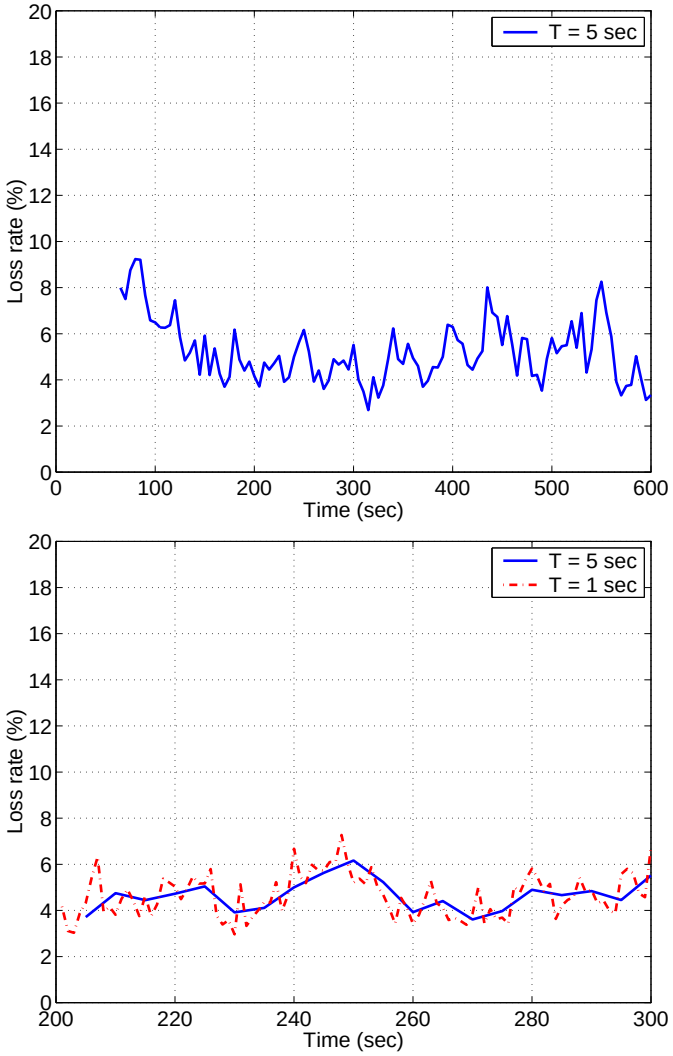

Figure 4. Loss rates at the router buffer for the UDP simulations with the *Star Wars* trace. Top: The observation interval is from 60 sec to 600 sec and the loss rate calculation interval is 5 sec. Bottom: The observation interval is from 200 sec to 300 sec and the loss rate calculation interval is 1 sec and 5 sec. In both figures, the number of sources $n = 100$, the buffer size $B = 102.4$ kbytes, and the packet size = 200 bytes.

the second loss episode is one packet. *Loss distance* is a measure of spacing between two consecutively lost packets. It is defined as "the difference in sequence numbers of two successively lost packets that may or may not be separated by successfully received packets" [15]. If these two consecutively lost packets are separated by one or more successfully received packets, then the loss distance is a measure of the spacing between two consecutive loss

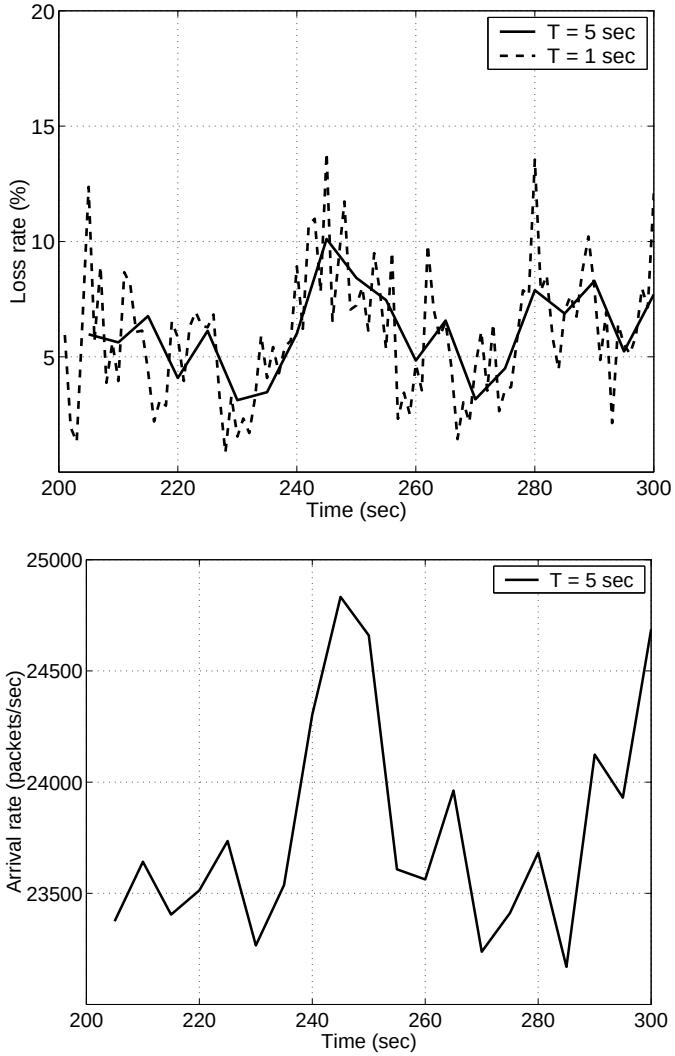

Figure 5. Loss rates of a particular source (top) and aggregate arrival rate (bottom) for the UDP simulations with the *Star Wars* trace. The observation interval is from 200 sec to 300 sec. The loss rate calculation interval is 1 sec and 5 sec and the arrival rate calculation interval is 5 sec. The number of sources $n = 100$, the buffer size $B = 102.4$ kbytes, and the packet size $= 200$ bytes.

episodes. The definition of *loss episode* and *loss distance* is illustrated in Fig. 6.

3.1. LOSS PATTERNS

To illustrate our simulation results we use textured dot strip plots [21]. The plot, depicting the loss pattern of a simulation run with 80 sources and buffer

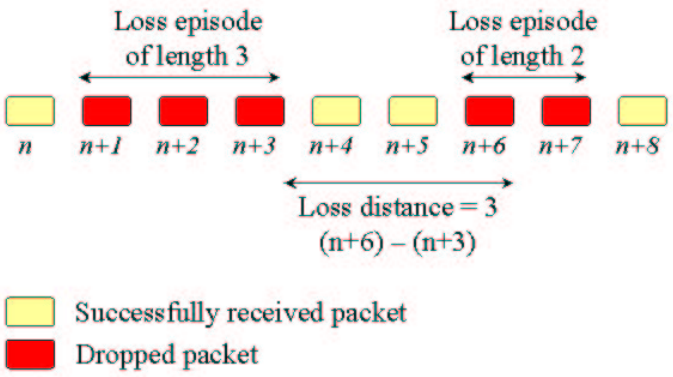

Figure 6. Definition of *loss episode* and *loss distance*.

size of 102.4 kbytes, is shown in Fig. 7. Each point represents the time instance when the loss occurred. The points are randomly placed along the vertical axis within the height of the plot for better visibility of the bands with more or less losses. We can identify periods of high congestion between 615 and 625 seconds (band A), lower congestion between 836 sec and 846 sec (band B), and long periods with no loss, as for example between 826 sec and 836 sec (band C), or between 846 sec and 854 sec (band D).

An alternative approach to observe the bands of higher or lower congestion is to count the number of lost packets within intervals (bins) covering the observation interval, as shown in Fig. 8 for fixed-size intervals of 5 sec.

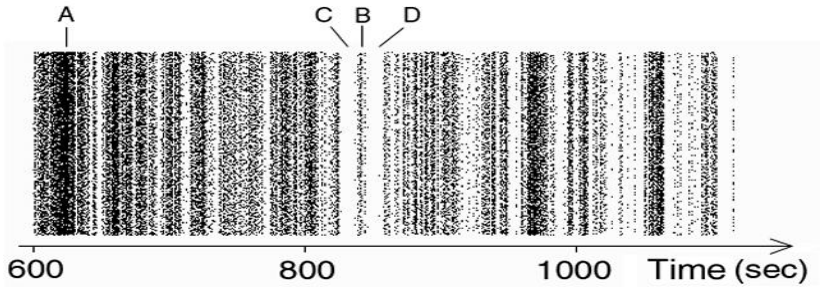

Figure 7. Textured dot strip plot of packet loss instances at the common buffer from a simulation run with n = 80 sources, buffer size B = 102.4 kbytes, and packet size = 200 bytes.

We observe two 10-second time bands A and B in Fig. 7. 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 III presents the summary for the loss episode lengths and loss episode distances for a simulation with buffer size of 102.4 kbytes (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 (in packets) of the loss episode.

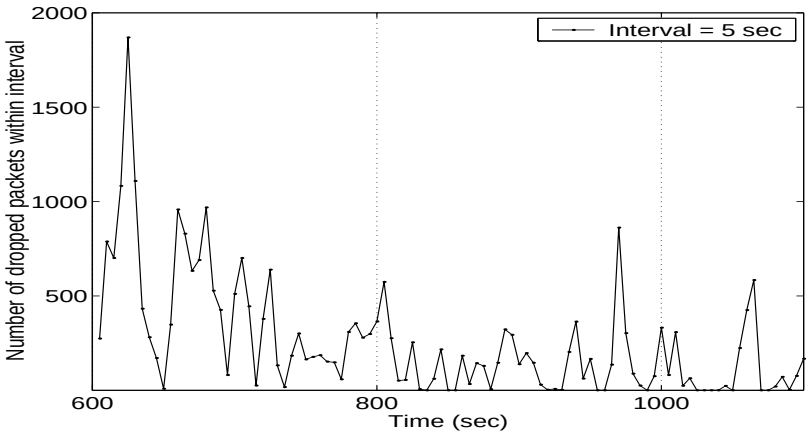

Figure 8. Aggregate number of dropped packets in intervals of duration 5 sec. Simulation run with n = 80 sources, buffer size B = 102.4 kbytes, and packet size = 200 bytes.

We can see from Table III that the mean loss episode distance is smaller for the highly congested band A. Within this band, single losses accounted for only 25.98 % 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 43.48 % of all lost packets in band B.

Table III. Contribution of lost packets (in percents) from loss bursts of length i in band A (higher congestion) and band B (lower congestion), shown in Fig. 7. The number of sources n = 80, the buffer size B = 102.4 kbytes, and packet size = 200 bytes.

i (packets)	1	2	3	4	5
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				

3.2. CONSECUTIVE PACKET LOSS

We found that consecutive packet losses represent the majority of the overall packet loss, even though loss episodes of length one represent a high percentage of all loss episodes. For example, Figs 9, 10, and 11 illustrate the contribution of lost packets from loss episodes of various lengths (left), and the contribution of loss episodes of various lengths (right).

It can be observed from Fig. 9, for the case of 80 sources and with buffer size of 25.6 kbytes (4.58 msec), that the single losses (loss episodes of length

one) represent 50 % of the total number of loss episodes. Nevertheless, they represent only 26 % of the total number of lost packets. For the case with 80 sources and buffer size 102.4 kbytes (18.31 msec) shown in Fig. 11, the single losses represent 53 % of all the loss episodes, while they contribute to only 29 % of the total number of lost packets.

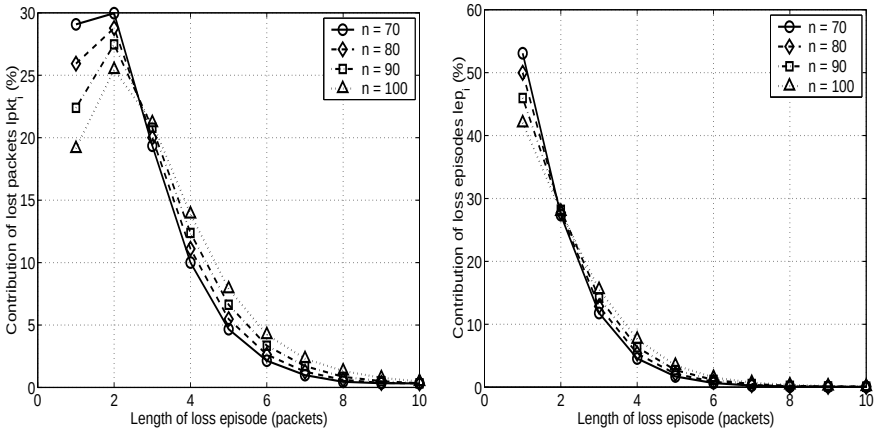

Figure 9. Contribution of lost packets (left) and loss episodes of various lengths (right) to the overall number of loss episodes. *Star Wars* trace. The number of sources = n, the buffer size $B = 25.6$ kbytes, and the packet size $= 200$ bytes.

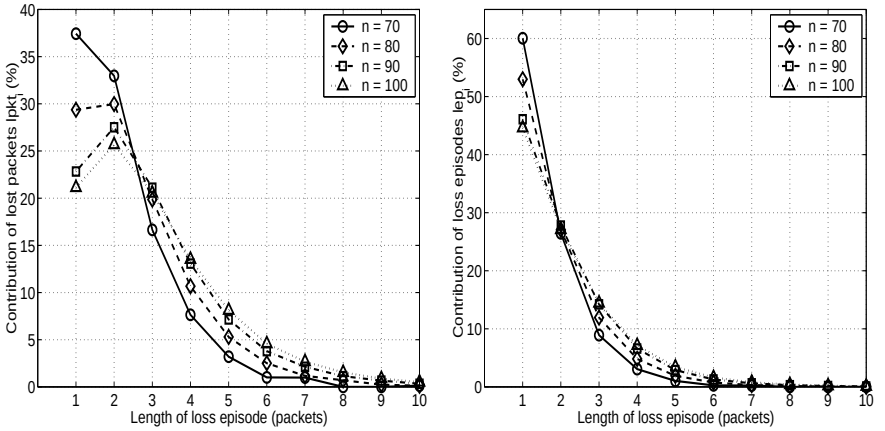

Figure 10. Contribution of lost packets (left) and loss episodes of various lengths (right) to the overall number of loss episodes. *Star Wars* trace. The number of sources = n, the buffer size $B = 102.4$ kbytes, and the packet size $= 200$ bytes.

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 102.4 kbytes are shown in Fig. 12. The results are qualitatively similar to those for the *Star Wars* trace.

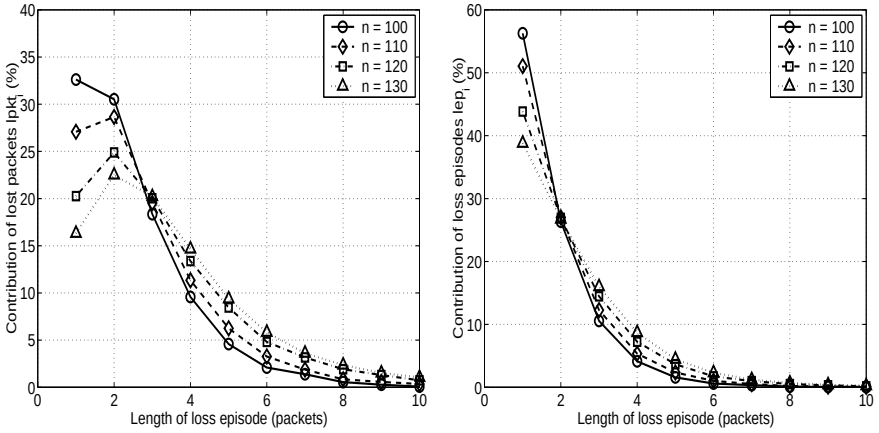

Figure 11. Contribution of lost packets (left) and loss episodes of various lengths (right) to the overall number of loss episodes. *Star Wars* trace. The number of sources = n, the buffer size $B = 204.8$ kbytes, and the packet size $= 200$ bytes.

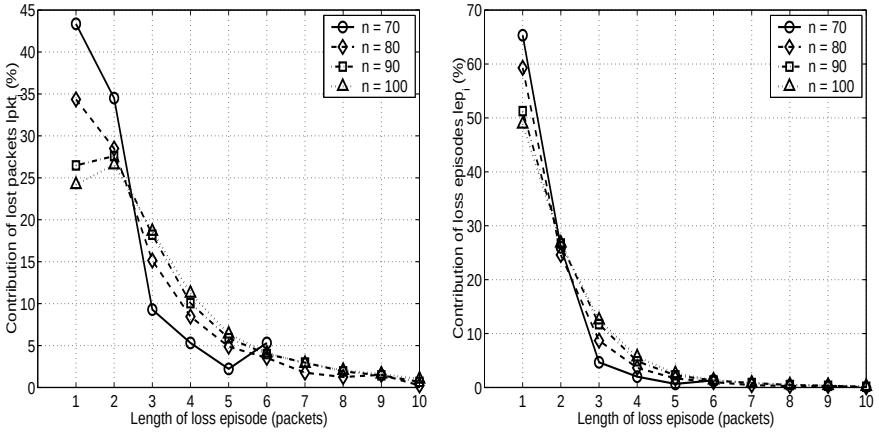

Figure 12. Contribution of lost packets (left) and loss episodes of various lengths (right) to the overall number of loss episodes. *Talk show* trace. The number of sources = n, the buffer size $B = 102.4$ kbytes, and the packet size $= 200$ bytes.

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 IV. We emphasize that the first-order (or summary) statistics does not capture the behavior of the observed loss process.

For the cases shown in Figs. 10–12, the values for the lost packets and loss episodes are obtained for the total observation interval of 600 sec. If narrower intervals are 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.

Table IV. Total number of arrived and lost packets, utilization, and average loss for n sources at the router for the *Star Wars* and *Talk show* trace. The buffer size B = 102.4 kbytes.

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 %

3.3. MODELING LOSS EPISODES

Measurements in the Internet indicated that the probability of loss episodes of length k decreases approximately geometrically with increase of k [7, 27]. This behavior can be captured by a two-state homogeneous Markov chain [14, 20]. The state diagram of the model (also known as the Gilbert model) is shown in Fig. 13 (top). State 0 represents the state of successful packet arrival, while state 1 represents the state of packet being lost. State transition probabilities are:

$$\begin{aligned} p_{01} &= P(X=1|X=0) = P(\text{packet n is lost} \mid \text{packet } n-1 \text{ is received}) \\ p_{10} &= P(X=0|X=1) = P(\text{packet n is received} \mid \text{packet } n-1 \text{ is lost}). \end{aligned} \quad (1)$$

In steady-state, transition and state probabilities can be expressed in matrix form:

$$\begin{bmatrix} 1-p_{01} & p_{10} \\ p_{01} & 1-p_{10} \end{bmatrix} \begin{bmatrix} P(X=0) \\ P(X=1) \end{bmatrix} = \begin{bmatrix} P(X=0) \\ P(X=1) \end{bmatrix}. \quad (2)$$

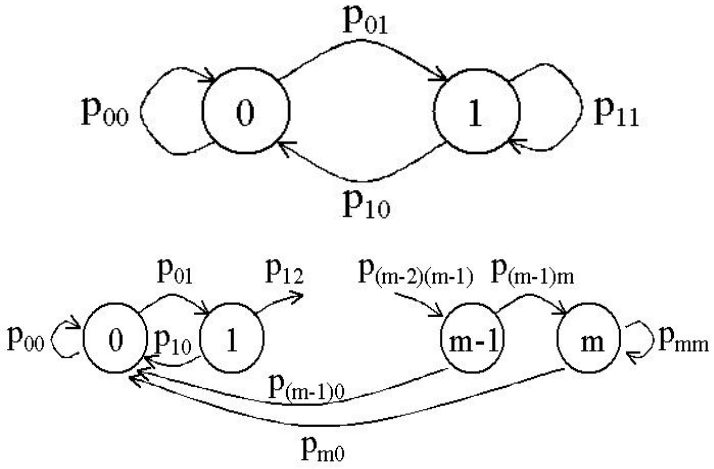

Figure 13. State diagram of the Gilbert model (top) and Extended Gilbert model (bottom).

From Eq. (2) and the equation $P(0)+P(1)=1$, the unconditional loss probability is:

$$P(X=1)=\frac{p_{01}}{p_{01}+p_{10}}. \tag{3}$$

The conditional loss probability, equal to the probability of having a loss, given the previous packet was lost, is:

$$P(X=1|X=1)=1-p_{10}. \tag{4}$$

The probability of having a loss episode with length k (k consecutively lost packets) given that we begin with the loss state $X=1$ is:

$$p_k=(1-p_{10})^{k-1}\cdot p_{10}. \tag{5}$$

Hence, the lengths of loss episodes given by the Gilbert model are geometrically distributed.

The Gilbert model memorizes only one past event. The probability that the next event will be either a successfully received or a lost packet depends only on the previous state. The values for transition probabilities p_{01} and p_{10} in the Gilbert model are based on a packet trace containing the information about the lost and successfully received packets [20]:

$$\begin{aligned} p_{01} &= P(X=1|X=0)=\sum_{k=1}^{\infty}\frac{o_k}{a} \\ 1-p_{10} &= P(X=1|X=1)=\frac{\sum_{k=1}^{\infty}(k-1)\cdot o_k}{d-1} \end{aligned} \tag{6}$$

where o_k is the number of loss episodes of length k, a is the total number of packets (successfully received and lost), and $d = \sum_{k=1}^{\infty} k \cdot o_k$ is the total number of lost packets.

The Extended Gilbert model was also proposed [14, 20] for modeling the loss behavior, It can track exactly the probability of k consecutive losses for $k \leq m-1$, where m is the order of the model. The state diagram is shown in Fig. 13 (bottom).

3.3.1. *Comparison of simulated loss data with the Gilbert model*
Gilbert model was originally derived from UDP loss traces obtained from genuine Internet measurements. We use it here to validate the loss process generated by simulations. We compare the probability of having loss episode of various lengths with the probabilities obtained from the Gilbert model. Since the model was used for loss data obtained by end-to-end measurements, which means loss affecting a single stream of data, we use loss data from a single traffic source, rather than the aggregate loss at the router buffer. For the simulation with the *Star Wars* trace, with $n = 100$ sources and buffer size $B = 51.2$ kbytes, the data needed to compute the parameters for the Gilbert model is given in Table V. Parameter a is the total number of packets sent by a particular source (source number 50 in our example), d is the total number of lost packets for that source, and o_k is the number of loss episodes of length k.

Table V. Input data for the Gilbert model. In all cases $a = 266,509$ and $d = 47,950$.

Parameter	Value							
k	1	2	3	4	5	6	7	8
o_k	15169	6042	2522	1186	608	346	180	99
k	9	10	11	12	13	14	15	16
o_k	49	32	15	7	7	4	2	2

The parameters of the Gilbert model can be obtained from Eq. (6). The transition and steady-state matrices for the Gilbert model are:

$$\begin{bmatrix} p_{00} & p_{10} \\ p_{01} & p_{11} \end{bmatrix} = \begin{bmatrix} 0.9014 & 0.5479 \\ 0.0986 & 0.4521 \end{bmatrix} \tag{7}$$

$$\begin{bmatrix} P(0) \\ P(1) \end{bmatrix} = \begin{bmatrix} 0.8475 \\ 0.1525 \end{bmatrix}. \tag{8}$$

The conditional probability of loss $p_{11} = 0.4521$ is larger than the unconditional probability $P(1) = 0.1525$, as expected [6]. This result also confirms

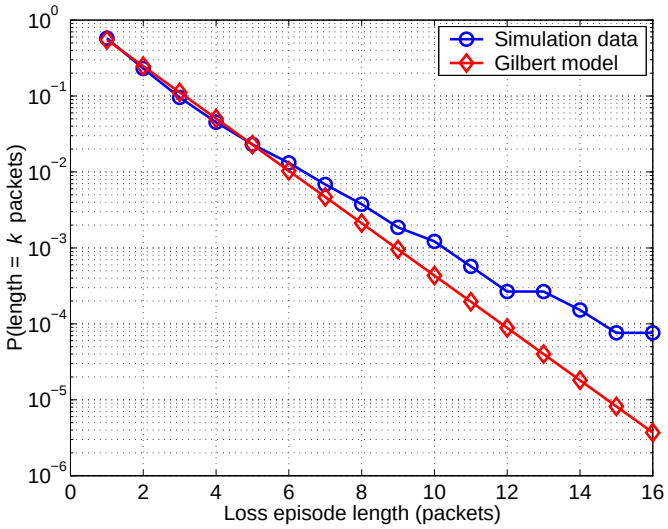

Figure 14. Comparison of the probabilities of loss-episode lengths for UDP transfers obtained from the simulation data and from the Gilbert model. The number of sources $n = 100$, the buffer size $B = 51.2$ *kbytes*, the packet size $= 200$ bytes. The loss from source number 50 is depicted.

that loss appears in bursts, i.e., if a packet is lost, there is a high probability that the next packet will be lost as well.

From Eq. (5) and Eq. (7) we can calculate the probability of having a loss episode of length k (or loss-episode length probability distribution), for $1 \leq k \leq L_{max}$, where L_{max} is the longest loss episode observed. In our example, $L_{max} = 16$.

A comparison between the distribution of loss-episode lengths derived from the Gilbert model and the distribution obtained from our simulation data is shown in Fig. 14. The Gilbert model characterizes loss-episode lengths accurately for smaller values of k, but underestimates the probabilities of having longer bursts.

4. Wavelet analysis of loss

Recent studies have shown that video traffic possesses complex correlation structure [5, 13] and exhibits long-range dependence (LRD) [10]. Wavelet analysis proved to be a useful tool to analyze the LRD properties of network traffic [4]. We investigate here the packet loss using wavelet analysis in order to understand the behavior of packet loss on various time-scales and to detect the presence of LRD in loss episodes.

4.1. LONG-RANGE DEPENDENCE

Let $X(t), t = 0, 1, ...,$ be a *covariance stationary* (sometimes called *wide-sense stationary*) stochastic process with mean m, variance σ^2, and autocorrelation function $r(k), k \geq 0$. X is said to exhibit *long-range dependence* if

$$r(k) \sim c_r k^{-(2-2H)}, \qquad k \to \infty, \tag{9}$$

where $c_r > 0$ and $0.5 < H < 1$. The key element in this description is the Hurst parameter H that measures the degree of long-range dependence. For short-range dependent processes $H = 0.5$. From Eq. (9) we see that long-range dependence is characterized by an autocorrelation function that decays hyperbolically as the lag increases. In the frequency domain, an equivalent property of an LRD process is that the spectral density $f(\lambda)$ of $X(t)$ takes the following form for a range of frequencies λ close to 0:

$$f(\lambda) \sim c_f |\lambda|^{1-2H}, \qquad \lambda \to 0, \tag{10}$$

where $c_f > 0$ and $0.5 < H < 1$.

Recently, long-range dependence has been found in network traffic measured from a wide range of high-speed networks such as Ethernet local area networks, wide area networks, and WWW [24]. It has also been found that LRD is an inherent feature of VBR video traffic, independent of scene and codec algorithms. For example, the genuine *Star Wars* video trace, which is used to drive our simulation in this paper, exhibits LRD with Hurst parameter $H \approx 0.8$ [5].

4.2. DISCRETE WAVELET TRANSFORM

The wavelet transform is the representation of a signal *X(t)* by means of its inner products with a set of basic wavelet functions $\psi_{a,t}(u) = \psi_0(\frac{u-t}{a})/\sqrt{a}$, called the wavelets, that are scaled and shifted versions of an adequately chosen mother wavelet ψ_0. *X(t)* is then transformed into a scale-time wavelet domain (*a*,*t*), $a \in R^+$ and $t \in R$. The wavelet transform presents a method for simultaneously observing a time series at a full range of different scales *a*, while retaining the time dimension of the original data. Multi-resolution analysis (MRA) theory shows that no information is lost if we sample the continuous wavelet coefficients at a sparse set of points in the scale-time plane known as the dyadic grid. This grid, defined by $(a,t) = (2^j, 2^j k),\ j,k \in N$, leads to the discrete wavelet transform. Detailed background on wavelets and MRA is given in [9].

An efficient implementation of the discrete wavelet transform is given by the filter-bank-based pyramidal algorithm. In this algorithm, the mother wavelet ψ_0 is represented in terms of the so-called scaling function ϕ_0:

$$\psi_0 = \sqrt{2} \sum_n (-1)^n h(1-n) \phi_0(2t-n). \tag{11}$$

Scaling function ϕ_0 is defined by a two-scale difference equation, known as the dilation equation:

$$\phi_0 = \sqrt{2}\sum_n h(n)\phi_0(2t-n). \tag{12}$$

The discrete-time filters $h(n)$ and $g(n) = (-1)^n h(1-n)$ have the low-pass and high-pass characteristics, respectively. Fig. 15 depicts the filter-downsample procedure of this transform. Given a time series *X(k)* of length N and the initial approximation coefficients $a_x(0,k) = X(k)$, we can compute the so-called detail coefficients $d_x(j,k)$ at scale 2^j ($1 \leq j \leq log_2 N$) by using the convolving approximation coefficients with filters $h(n)$ and $g(n)$:

$$\begin{aligned} d_x(j,k) &= g * a_x(j-1,k) \\ a_x(j,k) &= h * a_x(j-1,k). \end{aligned} \tag{13}$$

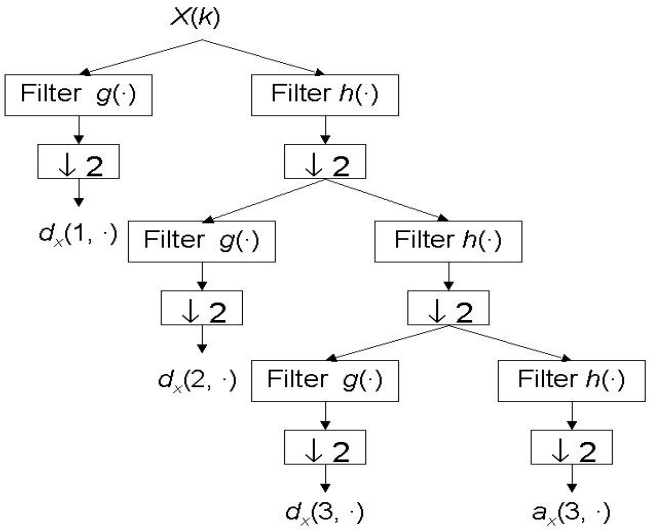

Figure 15. The filter-bank-based pyramidal algorithm. At each level j in the recursive structure, $d_x(j,k)$ and $a_x(j,x)$ are the outputs of the filters *g(n)* and *h(n)*, respectively. They occur at half the sampling rate of the input $a_x(j-1,x)$.

4.3. WAVELETS AND THE LRD DETECTION

A wavelet-based Hurst parameter estimator, suggested in [4], is based on a spectrum estimator Γ_x designed by finding a time average of $|d_x(j,k)|^2$ at a given scale level j:

$$\Gamma_x(2^{-j}\nu_0) = \frac{1}{n_j}\sum_k |d_x(j,k)|^2, \tag{14}$$

where $d_x(j,k)$ are the detail coefficients obtained from the discrete wavelet transform of time series $X(k)$ at scale 2^j, and n_j is the number of wavelet coefficients at scale level j. A linear relationship between $log_2(\Gamma)$ and j over a range $[j_1,\ j_2]$ indicates the presence of LRD behavior. We can find an estimator $\hat{H}$ for the Hurst parameter H by performing a linear regression of $log_2(\Gamma)$ on scale level j in the range $[j_1,\ j_2]$:

$$log_2(\Gamma_x(2^{-j}\nu_0)) = (2\hat{H}-1)\,j + \hat{c}. \tag{15}$$

We have evaluated the performance of this wavelet-based estimator for variety of genuine and synthetically generated traffic traces [26]. Rather than focusing on the quantitative aspects of the LRD estimation emphasized in [4, 26], we suggest here a more qualitative use of the wavelet-based analysis. We use the $log_2(\Gamma)$ vs. j plots to examine the scaling behavior of data. From the plots we observed the linear relationship between $log_2(\Gamma)$ and j over several scales that enabled us to identify the scale range over which the LRD is present. The plots also enabled us to determine the break-point scale beyond which we can accurately detect the LRD. Hence, the technique we propose may be used to detect both the presence and the location of LRD.

An example of the wavelet analysis graph for the *Star Wars* trace is shown in Fig. 16 (top). The wavelet analysis graph for the traffic trace obtained after video frames of the *Star Wars* trace were fragmented into packets is shown in Fig. 16 (bottom).

4.4. Scaling behavior of packet loss

In order to study the loss behavior on various time-scales, we consider samples of the loss processes obtained from *ns-2* simulations. Each sample represents the number of lost packets during 1 msec. Graphs shown in Fig. 17 illustrate the results of our wavelet analysis for various buffer sizes and source numbers. We observe that plots shown in Fig. 17 are similar and, hence, that the overall scaling property of packet loss is robust. Furthermore, the packet loss behavior varies over the time-scales and the loss property changes over different time-scales.

The importance of time-scales on the network traffic analysis and connection admission control has already been reported [11, 23]. For example, it has been shown in [11] that the measured Internet traffic is LRD (or asymptotically self-similar) over coarser time-scales, and that it exhibits multifractal behavior over finer time-scales. Our results presented here show that time-scales also play an important role for a comprehensive understanding of the behavior of packet loss.

In the case of UDP transfers shown in Fig. 17, the linear relationship between $log_2(\Gamma)$ and j is evident for the coarser time-scales beyond the break-point $j = 10$. Note that the same time-scale break-point ($j = 10$) was observed

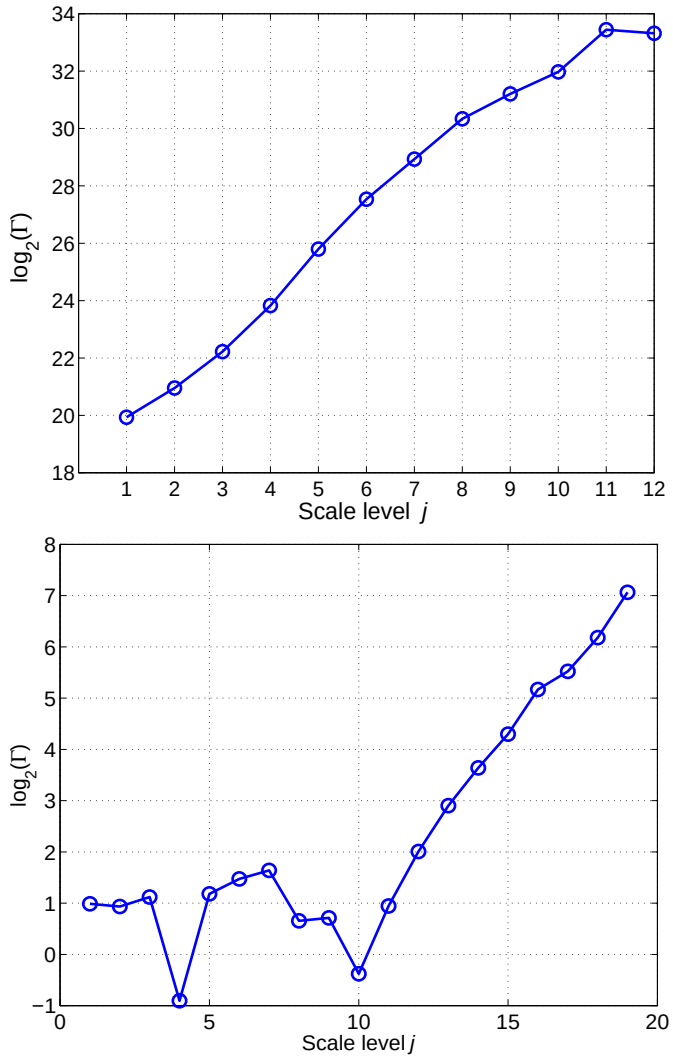

Figure 16. The $log_2(\Gamma)$ vs. j plot for *Star Wars* traffic trace (top) and for the packetized traffic trace (bottom).

in the packetized *Star Wars* traffic trace shown in Fig. 16 (bottom). This phenomenon indicates that the time-scale granularity for possible detection of LRD in our example is close to 1,024 (2^{10}) msec. In order to gain more understanding of the break-point effect observed in packet loss processes with UDP transfers, we aggregated the loss processes over three time-scales: 10 msec, 50 msec, and 100 msec. For these three cases, the break-point scale levels are 7, 4, and 3, as shown in Fig. 18. The corresponding break-point time-scales are 1280 msec, 900 msec, and 800 msec, respectively. These results indicate

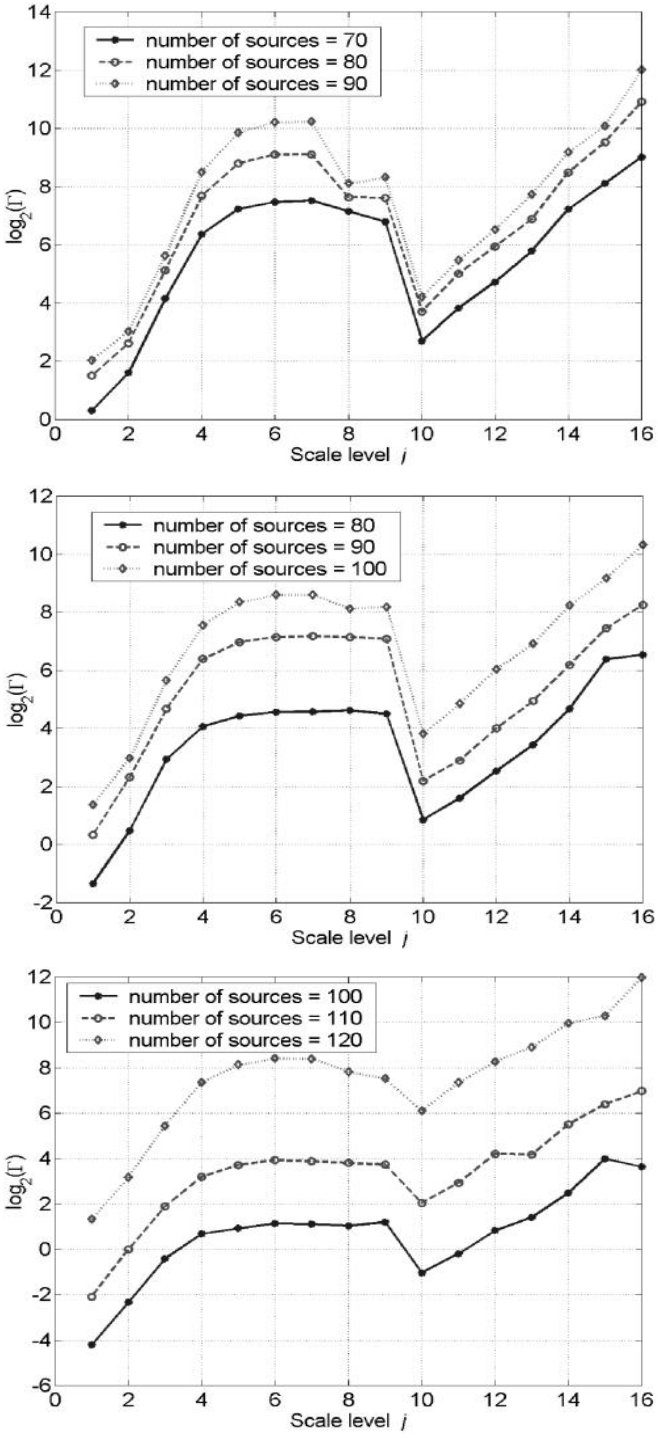

Figure 17. The $log_2(\Gamma)$ vs. j plot of packet loss process for buffer sizes: $B = 25.6$ kbytes (top), $B = 102.4$ kbytes (middle), and $B = 204.8$ kbytes (bottom). Packet size = 200 bytes.

a break-point time-scale near 1,000 msec for packet loss processes resulting from UDP transfers.

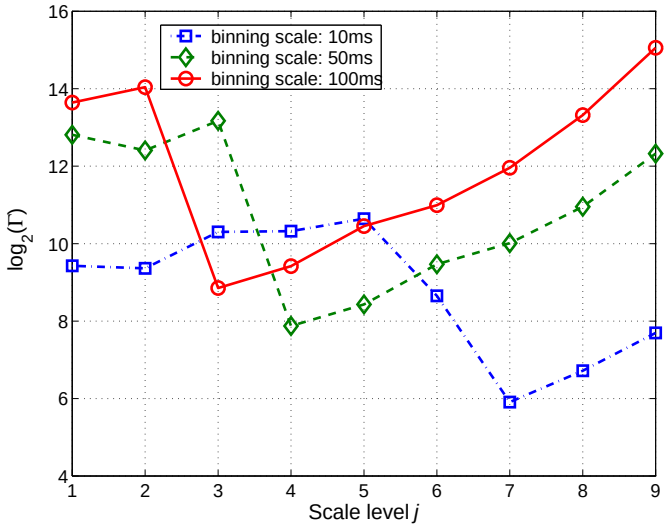

Figure 18. The $log_2(\Gamma)$ vs. j plot of packet loss process binned in various time-scales for UDP transfers: number of sources $n = 90$, buffer size $B = 102.4$ kbytes, and packet size $= 200$ bytes.

Once the time-scale break-point is observed, we can obtain new packet loss processes with a numerically more suitable time-scale. For example, we created a loss process in which each sample represents the number of loss packets over 1,000 msec. We then apply well known Hurst parameter estimators [16], such as variance-time and R/S analysis [17], to test the presence of LRD and to estimate Hurst parameters. Estimated Hurst parameters for various scenarios are given in Table VI.

We have also applied standard graphical techniques for estimating the Hurst parameter, such as the variance-time and R/S plots. The plots are obtained using the statistical package Splus tools for long-range dependence analysis [22]. From the UDP loss process obtained through simulation, a new time process was formed from the original loss trace by aggregation (binning) on a time-scale of 1024 msec. For example, for the case with $n = 100$ sources and buffer size $B = 102.4$ kbytes, the estimated value of the Hurst parameter using the variance-time plot is $\hat{H}_{variance-time} = 0.749$, while the its estimated value using the R/S plot is $\hat{H}_{r/s} = 0.789$.

Estimated Hurst parameters shown in Table VI and those obtained from the variance-time and R/S plots, are all in the range $(0.5, 1)$. This indicates packet loss process is long-range dependent over time-scales coarser than the break-point scale 10, which corresponds to time-scales of the magnitude of approximately 1,000 msec (1,024 msec) or larger. Hence, our findings

Table VI. Estimated Hurst parameters of packet loss process with 1,000 msec granularity, n sources, and buffer size B.

$B = 25.6$ kbytes			$B = 102.4$ kbytes		
n	Variance-time	R/S	n	Variance-time	R/S
70	0.8273	0.8326	80	0.7829	0.8662
80	0.6667	0.8663	90	0.6758	0.9007
90	0.6647	0.7971	100	0.7898	0.8563

$B = 204.8$ kbytes		
n	Variance-time	R/S
100	0.6716	0.9681
110	0.6927	0.8913
120	0.6030	0.8671

confirm the importance of time-scale on the quantitative estimation of packet loss.

The long-range dependence of the loss process can be related to the characteristics of the traffic feeding the router buffer. The wavelet spectrums for the loss process shown in Fig. 17 and of the aggregate input traffic shown in Fig. 16 (bottom) have similar patterns. The wavelet spectrums $\Gamma_x(2^{-j}\nu_0)$ for the loss process and the aggregate input traffic are shown in Fig. 19. Analysis of the aggregate input traffic confirms that the input traffic is also long-range dependent, with estimated Hurst parameters $\hat{H}_{variance-time} = 0.862$ and $\hat{H}_{r/s} = 0.871$.

5. Concluding remarks

In this paper, we studied packet loss in a video transfer over UDP with *ns-2* simulations driven by genuine video traffic. Our simulation results indicated that packet loss is bursty and that consecutive packet losses significantly contribute to the overall packet loss.

We performed wavelet analysis on packet loss over various time-scales. Wavelet-based analysis proved helpful not only for estimating Hurst parameter, as observed in the past, but also for finding the break-point level over which LRD detection can be performed.

We also observed that time-scales are important for estimating quantitative behavior of packet loss. For time-scales coarser than the break-point level, we applied variance-time and R/S analysis to obtain Hurst parameters of packet

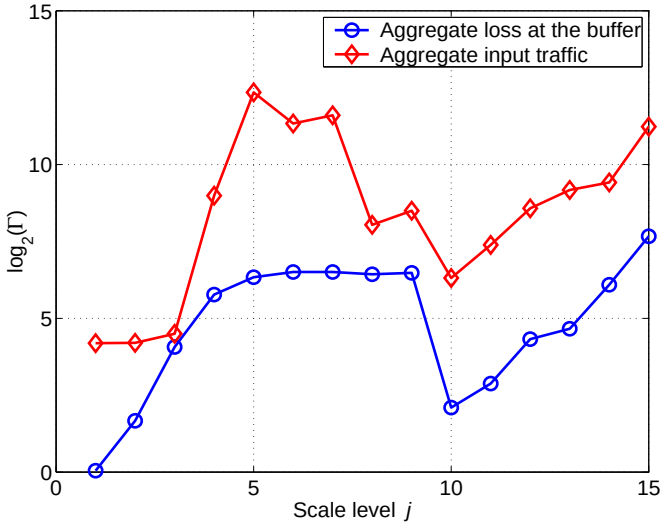

Figure 19. The $log_2(\Gamma)$ vs. j plot of the UDP packet loss process and the aggregate input traffic for buffer size $B = 102.4$ kbytes.

loss processes. Our results indicated the existence of LRD in packet loss processes on coarser time-scales. Our findings are robust and hold for various simulation scenarios. We are currently investigating the packet loss behavior on finer time-scales.

We conjecture that one possible cause for the LRD phenomena in packet loss may be that the source traffic itself (*Star Wars* trace) has been proven to be LRD. It is an open question how to establish the relationship between the LRD behavior of the source traffic and the parameters used to measure network performance (packet loss, delay, and delay jitter).

Our results imply that for our simulations with FIFO DropTail queue management policy in the router buffers, the aggregate loss process preserves the scaling properties of the aggregate input traffic process. We are currently investigating the effect of other types of queue management (e.g., Random Early Detection gateways) on the relation between the aggregate loss at the routers and the input traffic.

References

1. '*ns-2* network simulator'. <http://www.isi.edu/nsnam/ns>.
2. 'Star Wars trace'. In ASCII format <ftp://ftp.telcordia.com/pub/vbr.video.trace>; and in *ns-2* format <http://www.research.att.com/~breslau/vint/trace.html>.
3. 'University of Würzburg archive of video traces'. <http://nero.informatik.uni-wuerzburg.de/MPEG/>.

4. Abry, P. and D. Veitch: 1998, 'Wavelet analysis of long-range-dependent traffic'. *IEEE Transactions on Information Theory* **44**(1), 2–15.
5. Beran, J., R. Sherman, M. S. Taqqu, and W. Willinger: 1995, 'Long-range dependence in variable-bit-rate video traffic'. *IEEE Trans. Commun.* **43**(2/3/4), 1566–1579.
6. Bolot, J.-C.: 1993, 'End-to-end packet delay and loss behavior in the Internet'. In: *Proceedings of ACM SIGCOMM '93 Conference on Communications Architectures, Protocols and Applications*. San Francisco, CA, USA, pp. 289–298.
7. Bolot, J.-C., S. Fosse-Parisis, and D. Towsley: 1999, 'Adaptive FEC-based error control for Internet telephony'. In: *Proceedings of IEEE INFOCOM*. New York, NY, USA, pp. 1453–1460.
8. Borella, M. S. and D. Swider: 1998, 'Internet packet loss: Measurement and implications for end-to-end QoS'. In: *Proceedings of the 1998 ICPP workshops on architectural and OS support for multimedia applications/flexible communication systems/wireless networks and mobile computing*. Minneapolis, MN, USA, pp. 3–12.
9. Chui, C. K.: 1992, *An Introduction to Wavelets*. Boston, MA, USA: Academic Press.
10. Cox, D. R.: 1984, 'Long-range dependence: A review'. In: H. A. David and H. T. David (eds.): *Statistics: An appraisal*. Ames, IA, USA: The Iowa State University Press, pp. 55–74.
11. Feldmann, A., A. C. Gilbert, and W. Willinger: 1998, 'Data networks as cascades: Investigating the multifractal nature of Internet WAN traffic'. In: *Proceedings of ACM SIGCOMM '98 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication*. Vancouver, BC, CANADA, pp. 42–55.
12. Fester, M.: 1995, 'Performance issues for high-end video over ATM'. <http://www.cisco.com/warp/public/cc/sol/mkt/ent/atm/ vidat_wp.htm>.
13. Garrett, M. and W. Willinger: 1994, 'Analysis, modeling and generation of self-similar VBR video traffic'. In: *Proceedings of ACM SIGCOMM '94 Conference on Communications Architectures, Protocols and Applications*. London, UK, pp. 269–280.
14. Jiang, W. and H. Schulzrinne: 1999, 'QoS measurement of Internet real-time multimedia services'. Technical Report CUCS-015-99, Department of Computer Science, Columbia University.
15. Koodli, R. and R. Ravikanth: 2000, 'One-way loss pattern sample metrics'. Internet-draft, Internet Engineering Task Force. Expiration Date: December, 2000.
16. Leland, W. E., M. S. Taqqu, W. Willinger, and D. V. Wilson: 1994, 'On the self-similar nature of Ethernet traffic (extended version)'. *IEEE/ACM Transactions on Networking* **2**(1), 1–15.
17. Mandelbrot, B. B. and M. S. Taqqu: 1979, 'Robust R/S analysis of long run serial correlation'. In: *Proceedings of 42nd session of the International Statistical Institute*. Manila, Philippines, pp. 1–37.
18. Markovski, V. and L. Trajković: 2000, 'Analysis of loss episodes for video transfer over UDP'. In: *Proceedings of the 2000 SCS symposium on performance evaluation of computer and telecommunication Systems (SPECTS'2K)*. Vancouver, BC, Canada, pp. 278–285.
19. Rose, O.: 1995, 'Statistical properties of MPEG video traffic and their impact on traffic modeling in ATM systems'. In: *Proceedings of the 20th Annual Conference on Local Computer Networks*. Minneapolis, MN, USA, pp. 397–406.
20. Sanneck, H. and G. Carle: 2000, 'A framework model for packet loss metrics based on loss runlengths'. In: *Proceedings of the SPIE/ACM SIGMM Multimedia Computing and Networking Conference 2000 (MMCN 2000)*. San Jose, CA, USA, pp. 177–187.
21. Swayne, D. S., D. Cook, and A. Buja: 1998, 'XGobi: Interactive dynamic data visualization in the X window system'. *Journal of Computational and Graphical Statistics* **7**(1), 113–130.

22. Taqqu, M., B. Sherman, W. Willinger, and V. Teverovsky. Statistical methods for long-range dependence - Splus source code: <http://math.bu.edu/INDIVIDUAL/murad/methods/index.html>.
23. Trajković, L. and A. Neidhardt: 1999, 'Effect of traffic knowledge on the efficiency of admission-control policies'. *ACM Comput. Commun. Rev.* **29**(1), 5–34.
24. Willinger, W., M. S. Taqqu, R. Sherman, and D. V. Wilson: 1997, 'Self-similarity through high-variability: Statistical analysis of Ethernet LAN traffic at the source level'. *IEEE/ACM Transactions on Networking* **5**(1), 71–86.
25. Xue, F., V. Markovski, and L. Trajković: 2000, 'Wavelet analysis of packet loss in video transfers over UDP'. In: *Proceedings of the First International Conference on Internet Computing - IC'2000*. Las Vegas, NV, USA, pp. 427–433.
26. Xue, F. and L. Trajković: 2000, 'Performance analysis of a wavelet-based Hurst parameter estimator for self-similar network traffic'. In: *2000 SCS Symposium on Performance Evaluation of Computer and Telecommunication Systems - SPECTS'2K*. Vancouver, BC, Canada, pp. 294–298.
27. Yajnik, M., S. M. J. Kurose, and D. Towsley: 1999, 'Measurement and modelling of the temporal dependence in packet loss'. In: *IEEE INFOCOM*. New York, NY, USA, pp. 345–352.