INTERNATIONAL COMPUTER SCIENCE INSTITUTE

1947 Center St. • Suite 600 • Berkeley, California 94704-1198 • (510) 643-9153 • FAX (510) 643-7684

Characterization of Video Traffic

Rahul Garg
rahul@cs.berkeley.edu
The Tenet Group
Computer Science Division
Department of Electrical Engineering and Computer Science
University of California
Berkeley, CA 94720, USA

TR-95-007

Abstract

ATM networks will carry a wide variety of data over the same packet switching network. A majority of this traffic is expected to be real-time video generated by video on demand, video conferencing systems, etc. We study the characteristics of video data compressed using standard coding algorithms namely, JPEG, MPEG and also popular ones such as the video conferencing software NV. A wide range of video sources from movies to a class lecture were analyzed. Most of the traces were longer than an hour. The bit rate of the traces has been characterized using the leaky bucket model. We also show a method of choosing appropriate leaky bucket parameters. Burstiness function is used to characterize the burstiness of the video traffic at different time scales. Our studies indicate that JPEG compressed has very little short term burstiness. MPEG and NV traffic shows high burstiness over small time scales. JPEG and MPEG video exhibit burstiness over long time scales, whereas NV shows no burstiness over long time scales. It is found that for constant quality JPEG and MPEG compressed video, the leaky bucket parameters depend upon the contents of the video. For JPEG video, the service rate is mainly determined by the peak rate and for MPEG the service rate is given by the peak rate of the smoothed MPEG stream. However, the traffic generated by NV can be characterized independent of the actual video. The target sending rate of the software NV determines the service rate for its traffic.

Keywords: ATM, Burstiness, Burstiness Function, Characterization, JPEG, Leaky Bucket, MPEG, Networks, NV, Packet Video, Traffic Characterization.

1 Introduction

Future networks will carry a wide variety of data over the same packet switching network. A majority of this data is expected to be video data generated by multimedia applications. The real-time traffic would need some kind of quality of service (QOS) guarantees from the network. Call admission control enables the network to provide the QOS guarantees. Various types of scheduling disciplines, QOS parameters and associated admission control methods have been proposed in [6], [21], [13], [12], [5] and [3].

Most of these schemes need a traffic descriptor (like peak bit rate, mean bit rate etc.) which characterizes the traffic. These descriptors are used in the connection setup phase to determine whether the network has enough resources to satisfy the QOS requirements of the traffic. Therefore characterization of video traffic using appropriate descriptors is very important.

It is well known that constant quality video has variable bit rate and is bursty. The burstiness of the traffic is related to the statistical multiplexing gain. The presence of larger variety of compression algorithm make the task difficult. Past studies [11] [17] have characterized one particular coding algorithm. Also it is not very clear how can the characterization help in network management functions like admission control and policing.

In this paper we characterize the video traffic using a widely known traffic descriptor - Leaky Bucket. Bandwidth and buffer assignment for leaky bucket model has been studied in great detail. Therefore leaky bucket is an appropriate descriptor to characterize the traffic. We characterize burstiness of the traffic at various time scale using the burstiness function as defined by Cruz [4].

The traces used for this paper are from a variety of sequences and encoding algorithms. The sequences were choosen to represent a range of sequences present in real life (Broadcast TV, Class Lecture, full length movie). We studied standard compression algorithm like JPEG, MPEG and traces of popular software NV.

JPEG is a compression standard which was originally designed by the Joint Photographic Experts Group for coding and storing still photographic images [20] [18]. Since then it has found its application in many other fields, including digital video. To encode digital video each frame is compressed independently as a still picture using the JPEG compression algorithm. The video consists of a series of independently compressed JPEG pictures which makes the algorithm an intra-frame compression algorithm.

MPEG is a compression standard designed for storage and transmission of video [10] [1]. It is an inter-frame coding algorithm which exploits spatial and temporal redundancy of the video to achieve the compression. It also uses a DCT based coding as in JPEG algorithm. The frames are classified into 3 types - I-frames, P-frames and B-frames. I-frames are coded as still picture and are independent of other frames. This makes I-frames largest in size. P-frames only carry the changes between the last I-frame or P-frame and the current frame. B frames which are smallest in size carry changes in the current frame with respect to last or future I-frame or P-frame. A typical order of frames is I-B-B-P-B-B-P-B-B-I-B-B-...

NV is a popular software used over the Internet to hold video conferences at low bit rates [8]. The NV software tries to limit the bit rate while sending video. This bound or target sending rate can be selected by the user and is generally set to 128 K-bits per second for sending data over the Internet. To achieve this rate control the NV software measures the actual rate at which it is sending the data. When it realizes that it has sent too much data then it becomes inactive for about 1-2 seconds to reduce its mean sending rate. The coding algorithm used by NV is a varient of Harr's transform which uses conditional replenishment for individual blocks.

It has been found that JPEG compressed video has very little short term burstiness. MPEG and NV traffic shows high burstiness over small time scales. JPEG and MPEG video exhibit burstiness over long time scales, whereas NV shows no burstiness over long time scales.

It is found that for constant quality JPEG and MPEG compressed video, the leaky bucket parameters depend upon the contents of the video. For JPEG video, the service rate is mainly determined by the peak rate and for MPEG the service rate is given by the peak rate of the smoothed MPEG stream. However, the traffic generated by NV can be characterized independent of the actual video. The target sending rate of the software NV determines the service rate for its traffic.

The rest of the paper is organized as follows. Section 2 discusses the leaky bucket model and burstiness function which is used in characterizing the traffic. Section 3 discusses the source of traces and the contents of the actual video. Sections 4, 5 and 6 discuss the characteristics of traffic generated by JPEG, MPEG and NV coding algorithms. Discussion of the results is presented in section 7. Conclusion and future work is given is section 8.

2 Traffic Descriptors

The leaky bucket mechanism is widely used to police the traffic [19] [9] [2] [15]. It has two parameters - the bucket size σ and the mean service rate ρ. A traffic stream is said to comply to a leaky bucket descriptor of parameters (σ, ρ) when the amount of data carried by the stream in any interval of length I is bounded by $\sigma + \rho I$. In other words, the descriptor bounds the bit rate such that no burst can carry more than σ bytes of data and the long term average rate is bounded by ρ. Any traffic can be made compliant to another leaky bucket descriptor with the same rate and smaller bucket size by suitably buffering and delaying some of its packets. If a traffic complies with leaky bucket descriptors with parameters (σ_1, ρ) then it can be made compliant to parameters (σ_2, ρ) $(\sigma_2 \leq \sigma_1)$ by suitable regulation. The maximum buffer space needed to do this would be $\sigma_1 - \sigma_2$ and the maximum delay incurred by any packet due to this reshaping of traffic would be $\frac{\sigma_1 - \sigma_2}{\rho}$. In the extreme case when the output of the regulator is a smooth traffic having very small bucket size, the maximum smoothing delays incurred by the traffic would be $\frac{\sigma_1}{\rho}$ and the buffer space required to do this smoothing would be σ_1.

Cruz in [4] has presented a generalized model to characterize the burstiness of the traffic by bounding its worst case behavior. The burstiness function of a traffic stream is defined as :

$$b(I) = \text{Maximum amount of data sent in any interval of length } I.$$

This burstiness function is the most general bounding descriptor of the traffic. From the burstiness function one can construct any other bounding traffic descriptor.

For this paper, we use two descriptors to characterize the traffic - sigma-rho (σ, ρ) curves and the burstiness curves. Sigma-rho curve for a stream plots the minimum bucket size (σ)

TRACE	Approx Duration	Frame rate (FPS)	Content
LECTURE	56 min.	30	Lecture of a graduate course
NEWS	122 min.	30	News program
MUSIC	122 min	30	A music program
SPORT	122 min	30	Basketball game
STWAR	114 min.	25	Entertainment movie

Table 1: Traces of JPEG compressed video

against the service rate (ρ) such that the traffic complies to the leaky bucket descriptor with parameters (σ, ρ). Thus, each point on and above the graph represents a possible characterization of the traffic using the leaky bucket model. Burstiness curves plot $\frac{b(I)}{I}$ vs. I. $b(I)$ is the maximum amount of data sent in any interval of length I. Thus, $\frac{b(I)}{I}$ is the average bit rate in the interval in which maximum amount of data is sent, which is same as the worst case average rate of the traffic taken over an interval of length I. Worst case average rate for small intervals is close to the peak rate. As the rate increases the worst case average rate tends to the eventual average rate.

Slope of the burstiness curve is a measure of the burstiness of the traffic. The rate of decrease of the worst case average rate with the increase in averaging interval is given by the slope of the burstiness curve. In other words, it is a measure of additional smoothing obtained when the bit rate is averaged over a longer interval. If the traffic is not bursty in short term, then averaging the bit rate over small intervals will not result in any smoothing. As a result the slope of the burstiness curves will be zero for small intervals. If the traffic is not bursty in long term then all the smoothing would have happened when the bit rate is averaged over a long enough interval. Increasing the interval size will not result into any additional smoothing and so the slope of the burstiness curve will again be zero. Thus the slope of the burstiness curves characterizes the burstiness of the traffic at all time scales.

3 Methodology of Data Collection

3.1 JPEG Compressed Video

The data contains traces of five independent video clips, 1-2 hours long as summarized in Table 1. The video clips were chosen to span a wide variety of video ranging from a lecture in class to a basketball game.

The video corresponding to the trace LECTURE was a viedotape of a graduate course lecture. Most of the time, the video shows the instructor standing next to a board. It would occationally show a students asking questions. The traces NEWS, MUSIC and SPORT correspond to real life video broadcast on popular TV channels. The programs were choosen to span a variety of programs possibly having different natures. Since these video are taken from TV channels, they are true representative a big majority of the real life video sequences. The trace STWAR is the popular trace of the movie Star Wars.

All the data except STWAR was collected in real time using DEC's JVIDEO board which was able to carry out the JPEG compression of the 232 x 320 pixel colored images

TRACE	Approx. Duration	Frame rate (FPS)	Content
LECTURE.mpg	60 min	15	Lecture of a graduate course
NEWS.mpg	58 min.	15	News program

Table 2: Traces of MPEG compressed video

at the frame rate of 30 frames per seconds (fps). The compression was carried out at the minimum possible quantization factor to get the best picture quality[1]. The quantization factor was kept constant in order to get a constant picture quality[2]. The compressed data was discarded, only the frame sizes were recorded. The process of capturing the data was timed to verify that the capture rate is actually 30 fps.

3.2 MPEG Compressed Video

We used the Berkeley software MPEG encoder to carry out the MPEG compression. The video clips were first compressed in JPEG and stored on a local disk in real-time. The disk was just able to keep up at a frame rate of 15 Fps. Then this data was compressed off-line into MPEG using a software encoder. For various reasons 1-5% of frames were dropped till the final traces were obtained. The parameters of the MPEG encoder were chosen to minimize the running time of the encoder and to get a good picture quality. The quantization factors were kept constant to get a constant picture quality[3]. Table 2 summarizes the information about the traces.

The video for the traces LECTURE.mpg and NEWS.mpg correspond to the video for the traces LECTURE and NEWS.

3.3 Traces of data sent by NV

The NV traces were collected using the *tcpdump* program. The traces are of data sent for MBONE[4] [16] video-conferences held on XUNET[5] [7]. There was an active participant on the same LAN segment on which the observations were taken. During data collection traces of all the packets on the network were captured. The fraction of packets dropped was extremely small (176 packets out of 940695 for example). Packets originating from the local participant were filtered out off-line. Since in all the traces the sender was in the same LAN segment, the traffic characteristics of the traces are expected to be close to the source traffic characteristics of the NV program. Table 3 summarizes the information about the traces.

[1]It is hard to define "constant" or "good" picture quality because it involves subjective perception of picture quality. We use quantization factor as a crude measure of quality.

[2]Default quantization matrices were used for the compression

[3]The encoded stream was a MPEG-I stream with quantization factors of 2 5 and 15 for I, P and B-frames respectively. Default quantization matrices were used.

[4]MBONE stands for Multicast backBONE which a virtual network on the internet created to experiment with multicast.

[5]XUNET stands for eXperimental University NETwork. It is a wide-area network that serves as a testbed for research on data communication techniques.

TRACE	Approx. Duration	Date & time of conference	Sender
NV0	17 min.	09/22/93, 12:05pm	law.cs.berkeley.edu
NV1	74 min.	10/20/93, 11:19am	law.cs.berkeley.edu
NV2	58 min.	10/20/93, 11:19am	law.cs.berkeley.edu

Table 3: Traces of traffic generated by NV

In the trace NV1 it was found that, after 58 minutes from the beginning, the average bit rate went up to a value significantly higher than the routine 128 Kbps. After three minutes the bit rate came back to normal. Probably the participant turned up the rate control nob of the NV software and after 3 minutes he realized his mistake and became "nice" to the network again. NV2 consists of the first 58 minutes of the NV1 trace.

The traces of NV are packet traces whereas the JPEG and MPEG traces are frame traces. Thus the NV traces are also affected by factors other than its coding algorithm. These factors include rate control built in the software, ethernet capacity etc. It is still important to consider these traces because NV is one of the very few applications which actually uses the network to send live video. Studying these traces would also help us understand the interaction of system level parameters with characterization.

4 Characteristics of JPEG

Figure 1 shows the bit rate of the trace LECTURE. The first ten minutes of the corresponding video shows clock tower of UC Berkeley. The bit rate remains fairly constant during this period. The bit rate then shoots up as the text describing the lecture is superimposed with the clock tower. The lecture then begins with the instructor standing in front of a clean board. The time at which the lecture begins can be precisely identified by a sudden drop in the corresponding bit rate. The bit rate then increases steadily as more and more text is written on the board. The bit rate drops down again when the instructor starts writing on the second board and once again the video shows the instructor next to a clean board. The lecture goes on ...

Figure 2 shows bit rate of a 130 second clip of the trace NEWS. It is interesting to note that the bit rate remains nearly the same for some time, then it suddenly changes to a new value and then it stays at its new value for some time and this behaviour is repeated. Analysis of actual video shows that in most of the cases these abrupt changes in bit rates correspond to scene changes (or cuts) in the video. For a scene showing same people in same background the bit rate remains nearly the same. As the background and orientation of people changes, the bit rate changes slowly.

From visual inspection of the bit rates it is apparent that frame sizes of consecutive frames are highly correlated. The reason for this behaviour is that the size of a compressed frame is essentially determined by the complexity (information content) of the image being compressed and consecutive images of a video sequence are very similar in visual appearance as well as in the information contents. This similarity remains for longer runs of consecutive frames. As a result the bit rate remains highly correlated over that period. In the trace LECTURE the bit rate increases as the instructor writes more and more text on the board.

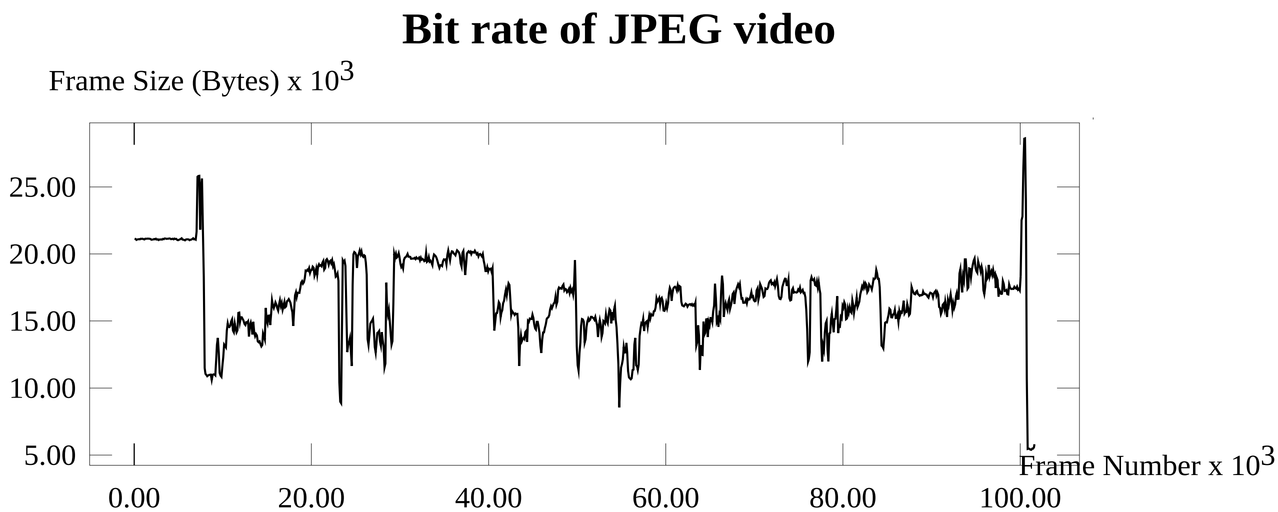

Figure 1: Bit rate of the trace LECTURE

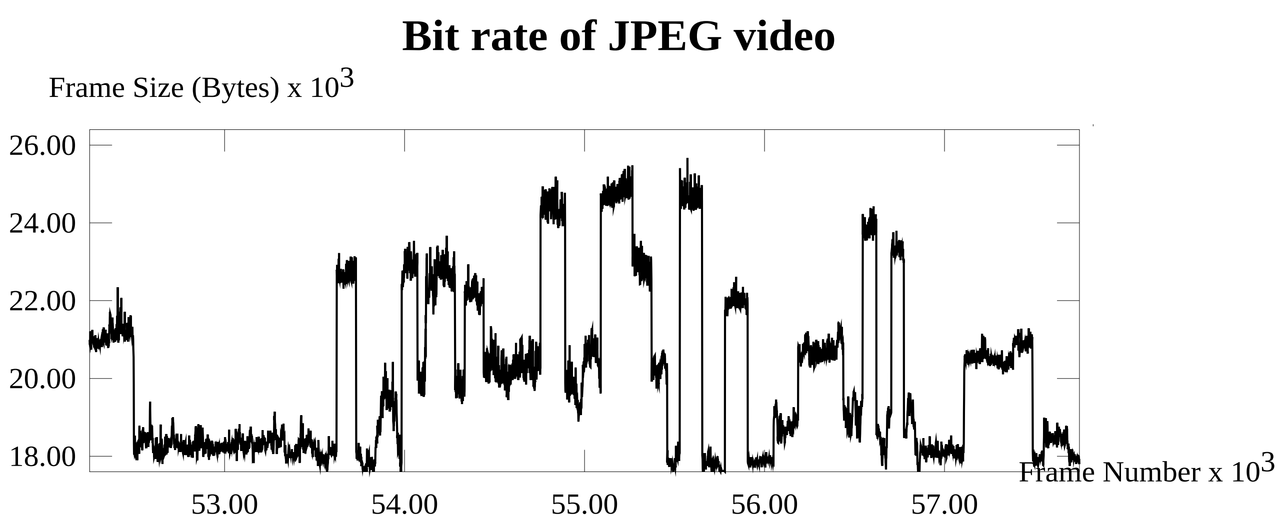

Figure 2: Bit rate of the trace NEWS

As more text is written on the board, there is more information is the corresponding image. This results into a lower compression and hence a higher bit rate. Scene changes mean an abrupt transition from one sequence of pictures to another. The two pictures at the boundary of a scene change are not similar to each other, so the sizes of the corresponding compressed frames are also uncorrelated. Therefore the bit rate changes abruptly at the boundary of scene changes in the video.

The peak bit rate of a trace corresponds to a scene of high information content in the corresponding video. The data remains close to the peak rate for the duration of the scene, which will usually be of the order of seconds. Thus the worst case average bit rate over an interval of length smaller than the peak scene length will be close to the peak rate. The burstiness curves in Figure 3 substantiate this claim. Figure 3 plots the worst case average bit rate of the traffic taken over an interval against the interval length. If the interval length is small, then the worst case average rate will be close to the peak rate. In the trace LECTURE, the worst case average rate remains close to the peak rate for interval of length up-to 10 seconds. This is the approximate duration of the scene showing clock tower along with the text describing the lecture. For other traces this interval is smaller. As the interval length increases the worst case average rate comes closer to the eventual average rate. For an interval of length equal to the length of the trace the worst case average rate is equal to the eventual average rate. It is interesting to note that the graphs in Figure 3 are very similar to each other for different video sequences. All the graphs start off with zero slope which remains close to zero for intervals of length 300 ms to 10 seconds. As the interval length increases, the magnitude of slope decreases. Towards the end, the slope converges to a small negative value.

The same information is depicted numerically in Table 4. In Table 4, *average* corresponds to the eventual average rate (in megabits per second) of the compressed video. All the traces of JPEG compressed video have average rates in the range 5-6.7 Mbps. The column labelled *burstiness* corresponds to the ratio of peak rate to the eventual average rate which is between 1.56 and 2.82. The peak bit rate varies from 19 Mbps to 7 Mbps. The following columns list the ratio of worst case average rate taken over an interval to the (eventual) average rate. As the interval size increases this ratio becomes closer to 1. It should be noted that this ratio for intervals of length 100 ms is close to the burstiness of the stream. These numbers behave in a similar way for each of the four different traces of JPEG compressed video. For example the ratios of the worst case average rate over an interval of 10 min to the average rate for different traces are between 1.00 and 1.25.

Figure 4 plots sigma-rho curves for the traces. For each trace, the service rate ρ varies from a little more than peak rate to the eventual average rate. An interesting observation is the knee points in these plots. In all of these curves, the position of the the rightmost knee corresponds to peak rate of the trace. The bucket size σ grows rapidly as soon as the service rate ρ becomes less than the peak rate. For example, for the trace LECTURE if the service rate is 80% of the peak rate then the bucket size becomes as large as 1.6 M-Bytes.

A reasonable characterization of JPEG compressed video traffic using a leaky bucket descriptor would require that the service rate be close to the peak rate. If one characterizes the video using a lower service rate the bucket size required will be very large which is not an accurate characterization of the actual traffic. A traffic characterized by a large bucket size can potentially inject a large burst of data in the network whereas the behaviour of

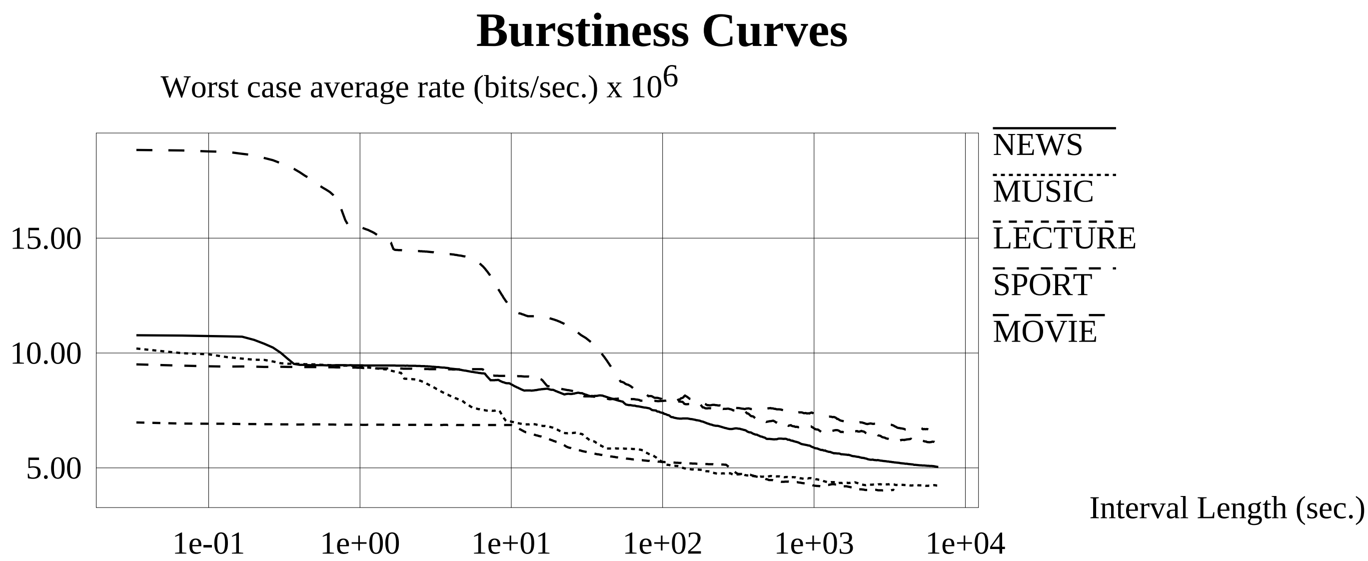

Figure 3: Burstiness curves for JPEG compressed video

Trace	Average (Mbps)	Burstiness	Worst case I = 100 ms	average rates I = 1 sec.	/ Eventual I = 100 sec.	average rate I = 10 min
NEWS	5.04	2.13	2.13	1.88	1.47	1.24
MUSIC	4.23	2.41	2.35	2.22	1.24	1.09
LECTURE	4.02	1.73	1.72	1.71	1.31	1.09
SPORT	6.09	1.56	1.54	1.54	1.30	1.14
MOVIE	6.67	2.82	2.81	2.32	1.20	1.13
LECTURE.mpg	0.61	9.88	5.53	1.95	1.24	1.15
NEWS.mpg	0.51	11.84	6.74	3.54	1.88	1.20
NV0	0.127	43.48	31.16	2.82	1.03	1.01
NV1	0.121	47.54	33.12	9.03	5.60	2.06
NV2	0.101	48.22	35.37	3.64	1.26	1.12

Table 4: Effect of I on worst case average rate

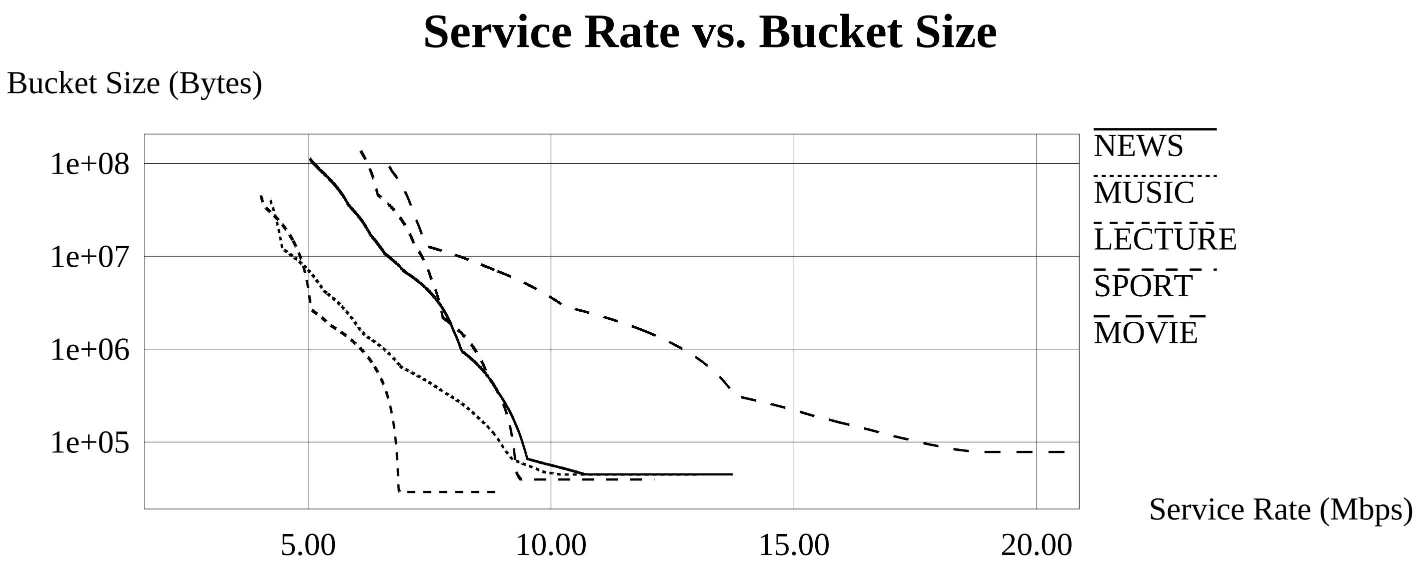

Figure 4: Sigma-Rho curves for JPEG compressed video

JPEG comressed video traffic is more regular. The large bucket in this case is not because of any bursts in the traffic, instead it is because the instantaneous arrival rate remains a little higher than the service rate for relatively long periods of time. The bursts are only due to the individual frames.

The service rate of leaky bucket for our JPEG traces was between 19 Mbps and 7 Mbps and the bucket sizes varied from 30 K-Bytes to 80 K-Bytes.

For constant quality JPEG compressed video the bit rate depends upon the complexity of the images in the video being compressed[6]. Thus, the peaks of the bit rate occur in bursts of length equal to the length of the scene having the high image complexity. This implies that, if a scene of high complexity is being transmitted over the network in real time, then the network should have enough bandwidth available over the length of the scene to support the transfer. Queueing inside the network may not be able to accommodate high bit rate bursts. Thus using a simple leaky bucket descriptor to characterize traffic would not yield any statistical multiplexing gains because it would require that the allocated rate be close to the peak bit rate of the actual video.

5 Characteristics of MPEG

Figure 5 shows the bit rate of a 20 second clip of the MPEG trace LECTURE.mpg. The bit rate of this clip is very bursty. This burstiness is due to three different types of frames generated by the MPEG compression algorithm. Note that I-frames, P-frames and B-frames can easily be recognized in the figure. I-frames are largest and B-frames are the smallest in

[6]The property discussed here is not an inherent property of the algorithms proposed by JPEG standards. It is only due to the fact that encoder used in our experiments compressed the images of individual frames at a constant picture quality.

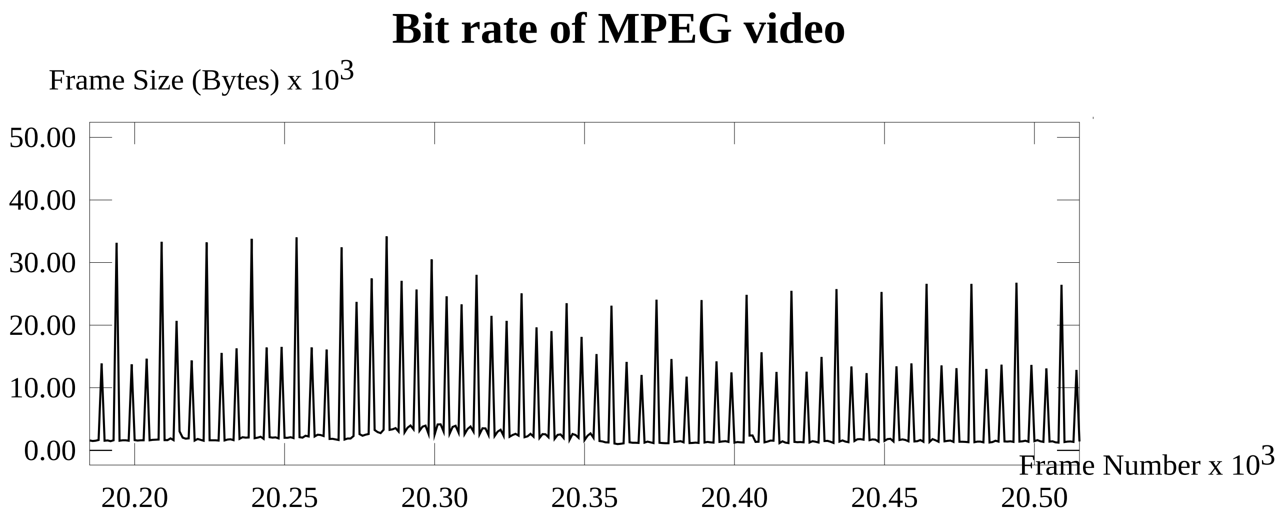

Figure 5: Bit rate of the trace LECTURE.mpg

sizes. Note that the sizes of frames of same type are highly correlated.

Figure 6 shows the average bit rate of the trace LECTURE.mpg. The average is computed by a moving average filter of length 1 second. An interesting observation is the similarity of the graphs in Figures 1 and 6. The same lecture was encoded into JPEG and MPEG separately and the variation of average bit rate of both the encodings appears to be very similar.

Figure 7 plots the burstiness curves for the MPEG traces. For small intervals the worst case average rate decreases quickly as the interval length increases. As the interval length becomes more than 1 sec., the rate of decrease of worst case average rate becomes smaller. Table 4 shows that the average rates of the MPEG video are 0.61 and 0.51 Mbps. The *burstiness* for two traces are 9.88 and 11.84. The peak rates are about 6.0 Mbps. After averaging the bit rate over an interval of 1 second the video show burstiness of 1.95 and 3.54.

Both the graphs of Figure 7 are very similar to each other. Both the graphs have large negative slope for small intervals. The magnitude of the slopes is smaller for interval lengths between 300 ms and 1 second. As the interval length increases further, the slopes converges a small negative value.

MPEG encoded video shows high burstiness when observed over short time scale. This is due to the fact that MPEG encoding algorithm generates different types of frames whose sizes are inherently different[7]. However, when averaged over an interval of length 1 second, the bit rate becomes smoother. This is because of the characteristics of MPEG encoding algorithm. One of the considerations while designing MPEG was the ability to have random access on the frames. In order to achieve this, intra coded frames (I-frames) are periodically present in the stream. Since I-frames are encoded independent of other frames, they provide random access points in the stream. P-frames and B-frames, which occur more frequently,

[7] It is possible to have a valid MPEG stream which consists of only I-frames or only I-frames and P-frames. Presence of all the three types of frames in the stream makes the compression ratio higher.

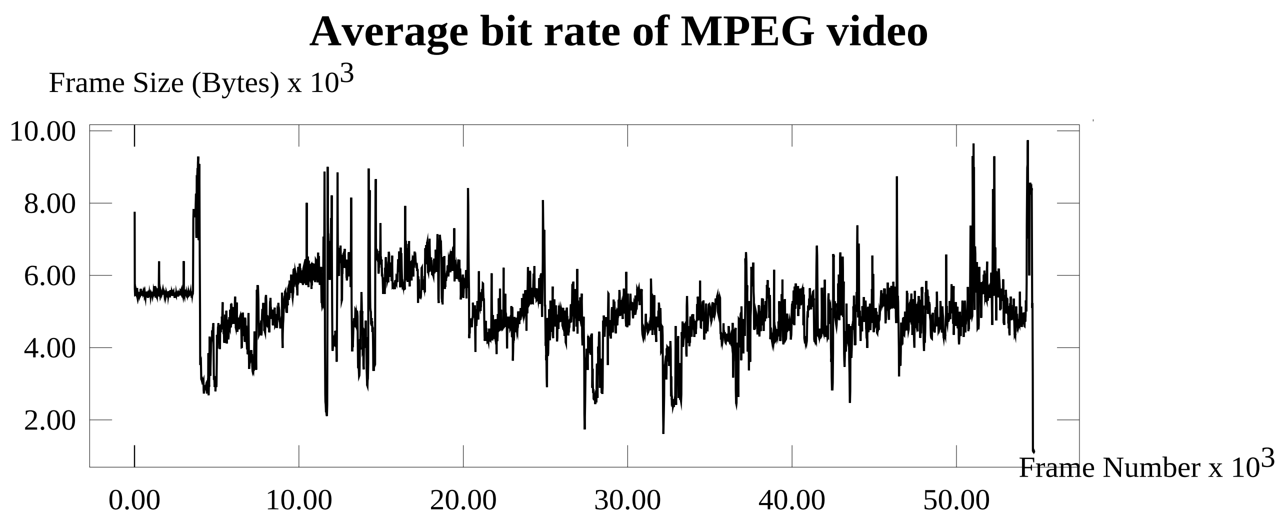

Figure 6: Average bit rate of the trace LECTURE.mpg

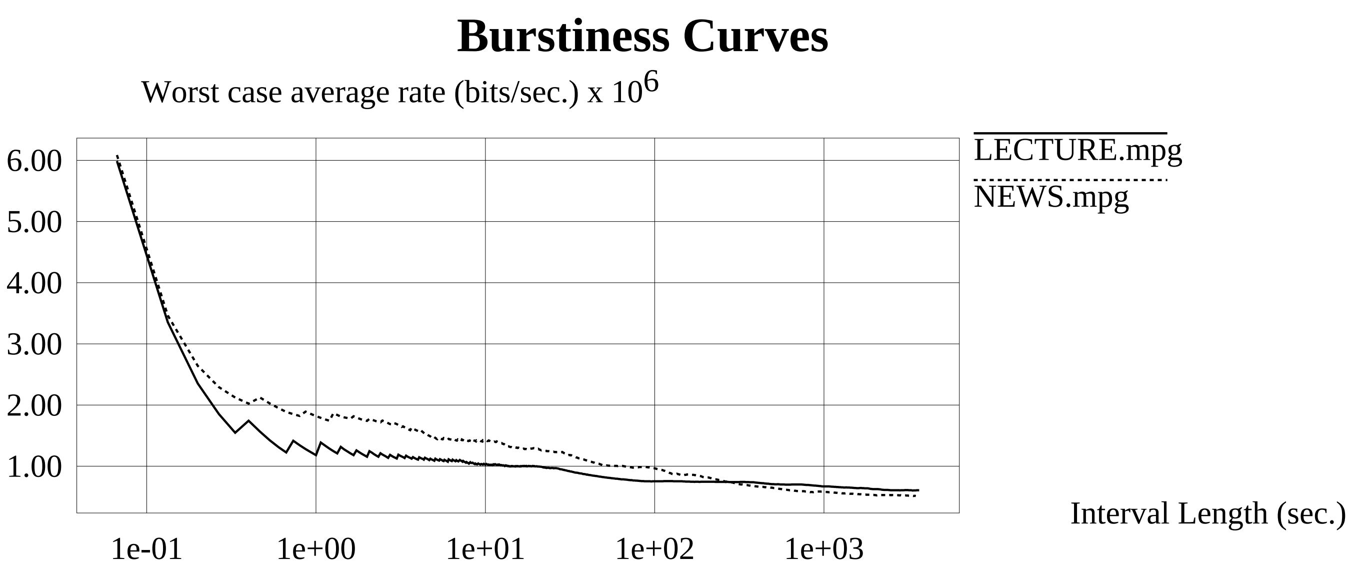

Figure 7: Burstiness curves for MPEG compressed video

constitute rest of the video. As a result MPEG video becomes a lot smoother if averaged over an interval which contains 2 or more I-frames.

After smoothing, the bit rate starts showing dependency on the contents of the video. This dependency can be clearly seen in Figure 6. Due to the periodic presence of I-frames in the stream, the bit rate of MPEG video depends upon the complexity of pictures which are sent as I-frames. Inter-frame nature of the algorithm makes the bit rate dependent on the motion in the video. In case of the video LECTURE, the motion was relatively low. As a result the bit rate averaged over periods covering multiple I-frames, showed a high correlation with the rate of corresponding JPEG trace.

Figure 8 plots sigma-rho curves for the MPEG traces. The service rate ρ for each trace varies from peak rate to the eventual average rate. It is interesting to note that the bucket size (σ) doesn't increase even if the service rate (ρ) is less than the peak bit rate. There is a knee in the curve after which the bucket size starts increasing rapidly. The rate of increase is comparable to that of Figure 4. The knee occurs at a rate which is lower than the peak rate. This is so because the largest frame is of I-type and it is always followed by a number of small B-type frames. By the time the frame next to the largest I-frame arrives, the bucket has enough space to accommodate it without necessarily serving at peak rate. The knee in the curve occurs at rate close to the worst case average rate over an interval of length one seconds.

An appropriate characterization of MPEG video using the leaky bucket model corresponds to the knee in the graphs of Figure 8. The service rate at these knees correspond to the peak rate of the MPEG video obtained after smoothing its bit rate over one second intervals. The bit rate of MPEG compressed video becomes highly correlated after smoothing and the changes in the bit rate are mostly due to the contents of the video. Therefore the knee points characterize the traffic accurately. For the trace LECTURE.mpg appropriate service rate and bucket size are 1.07 Mbps and 50 K-Bytes respectively.

MPEG compressed video exhibits burstiness on two time scales. On a short time scale the burstiness is due to the characteristics of the coding algorithm, which generates 3 types of frames of different average sizes. The short term burstiness smoothes out when the bit rate is averaged over intervals of appropriate length[8]. This length depends upon the periodicity of I-frames in the stream. After this smoothing the traffic remains bursty on a longer time scale. This burstiness for our streams were 1.95 and 3.54, which is due to the changes in motion and the complexity of scenes[9]. Since the smooth bitrate remains correlated for longer periods of time additional buffering is unable to smooth the bit rate further. The knee point in the sigma rho graphs show the same fact. The amount of additional buffering needed to be able to serve at a rate lower than the rate of knee is large. Therefore the knee points give a proper leaky bucket characterization of the video. Another implication of this is that for transmission of MPEG video over a network, the available bandwidth of the network at any instant should be more than the short term average rate of the MPEG stream. In case this bandwidth is not available, congestion would result into

[8]It must be pointed out that this interval is not equal to the delay incurred in smoothing the bit stream. The maximum delay incurred in smoothing is given by $\frac{\sigma}{\rho}$. See section 2 and 7 for details.

[9]The short term burstiness property is an inherent property of the algorithms proposed by MPEG standards. However, it also depends on the encoding process. In our case short term burstines also depends upon the ratio of the quantization factors of I, P and B-frames.

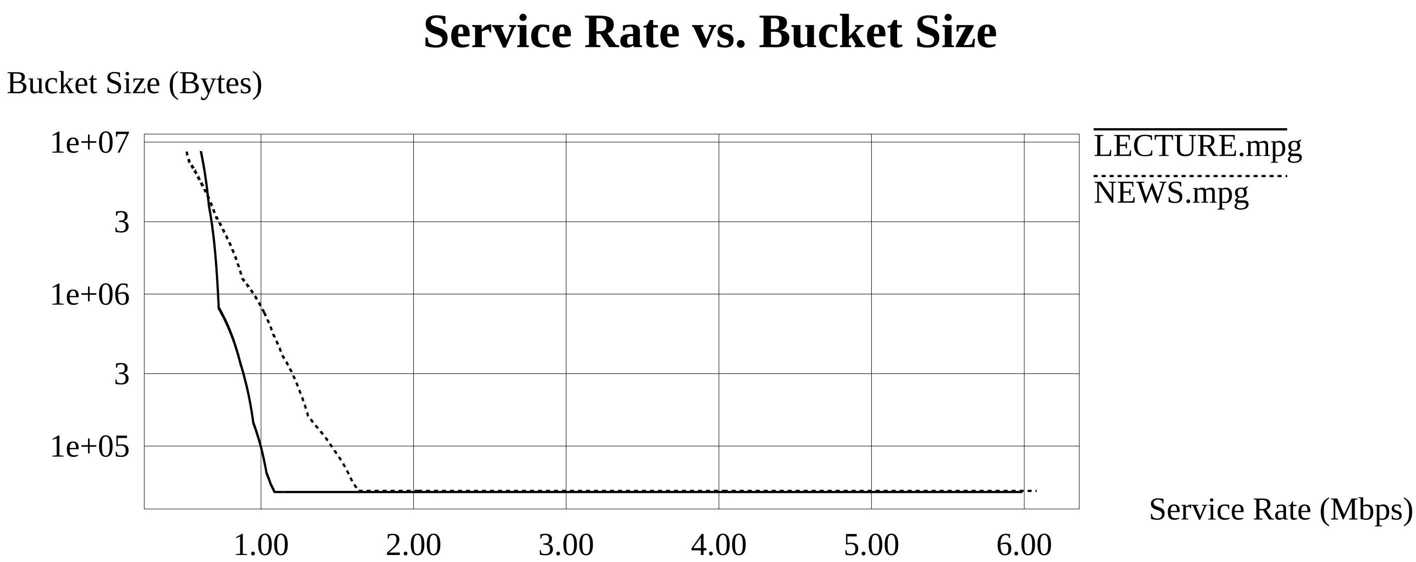

Figure 8: Sigma-Rho curves for MPEG compressed video

a series of localized packet loses. The parameter bucket size of the leaky bucket can also act as a measure of the buffering needed to smooth the short term burstiness.

6 Characteristics of NV

Characteristics of the traffic generated by NV is highly influenced by the rate control built in the software. Figure 9 shows the inter packet spacing for a series of packets generated by NV. Most of the packets are spaced at 5 ms which is probably the time it takes to send a packet. Inter-packet spacing of some packets periodically becomes as large as 0.5-2 seconds. This is due to the fact that NV stops sending data for long intervals if it realizes that it is sending the data at high rate. So there are periods of activity when NV sends the data at high rate and then there are periods on inactivity when NV "sleeps" to lower the mean bit rate. Another observation in Figure 9 is that there are intervals during which the packet spacing varies as high as 250 ms. This can be attributed to periodic waking up of other processes on the same host, or other activity in the same LAN segment. Since the NV program carries out the compression of video images in software, it shares the CPU with other processes. If other processes on the same system become active, then NV takes a longer time to produce a packet and hence the inter-packet spacing can increase. This can also result from some periodic network activity on the same LAN segment (e.g. routing updates) which would force NV to wait to send data on the Ethernet.

The packet sizes are distributed between 0 and 1 K-Bytes. And the minimum inter packet spacing is as small as 0.2 ms. So the peak rate of the traffic generated by NV is very high. However, at averaging intervals of 1 sec, the worst case average rate becomes close to the eventual average rate.

Figure 10 plots the burstiness curves for for the NV traces. The peak rate is close to

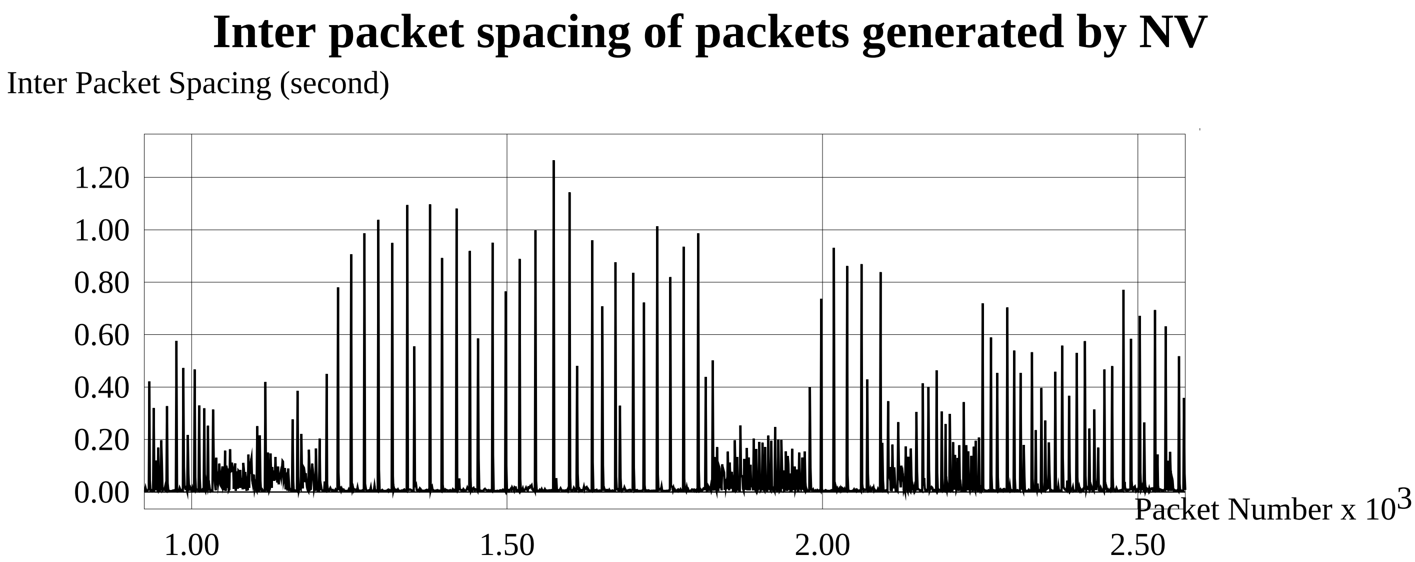

Figure 9: Inter packet spacing of packets produced by NV

6 Mbps, which is within a factor of two of of the maximum available bit rate on Ethernet. The worst case average rate is bounded by 130 Kbps for averaging intervals of 10 seconds, which is because of the rate control built in the software. For the traces NV0 and NV2, the worst case average rate decreases rapidly with interval length till the interval length is close to 3 seconds. For interval lengths between 3 and 10 seconds, the rate of change of worst case average rate decreases. At interval of length 10 seconds the worst case average rate becomes close to the eventual average rate. It is interesting to note that although NV0 and NV2 correspond to two independent traces, their burstiness curves are nearly identical. The burstiness curves have a high slope for small intervals which becomes very close to zero as the interval length increases to more than 10 seconds.

From Table 4 we can see that the ratio of worst case average rate taken over intervals of length 100 sec. to the eventual average rate for the traces NV0 and NV2 is 1.03 and 1.26 respectively. The trace NV1 is different because the target sending rate of the software was changed for 3 minutes. In Figure 10 there is a knee in the graph for NV1. This knee corresponds roughly to interval of length 3 minutes.

Figure 11 plots the sigma-rho curves for the NV traces. When the service rate is higher than the peak rate, the bucket size is close to 1 K-Bytes. As expected, the bucket size increases as the service rate decreases. Traces NV0 and NV2 have a knee in the graphs at service rate of 130 Kbps. Trace NV0 has a knee in the corresponding graph at 1 Mbps. The position of this knee gives a reasonable characterization of the traffic using the leaky bucket descriptor.

The traffic generated by the NV program is highly bursty in short term due to its on-off nature. The long term bahaviour of NV traffic is highly predictable. A plot of average bit rate taken over 10 second interval is nearly constant. Although the traces NV0 and NV2 are from two different conferences, their characteristics is nearly identical. This suggests that the characteristics of the traffic generated by NV are fairly independent of the video

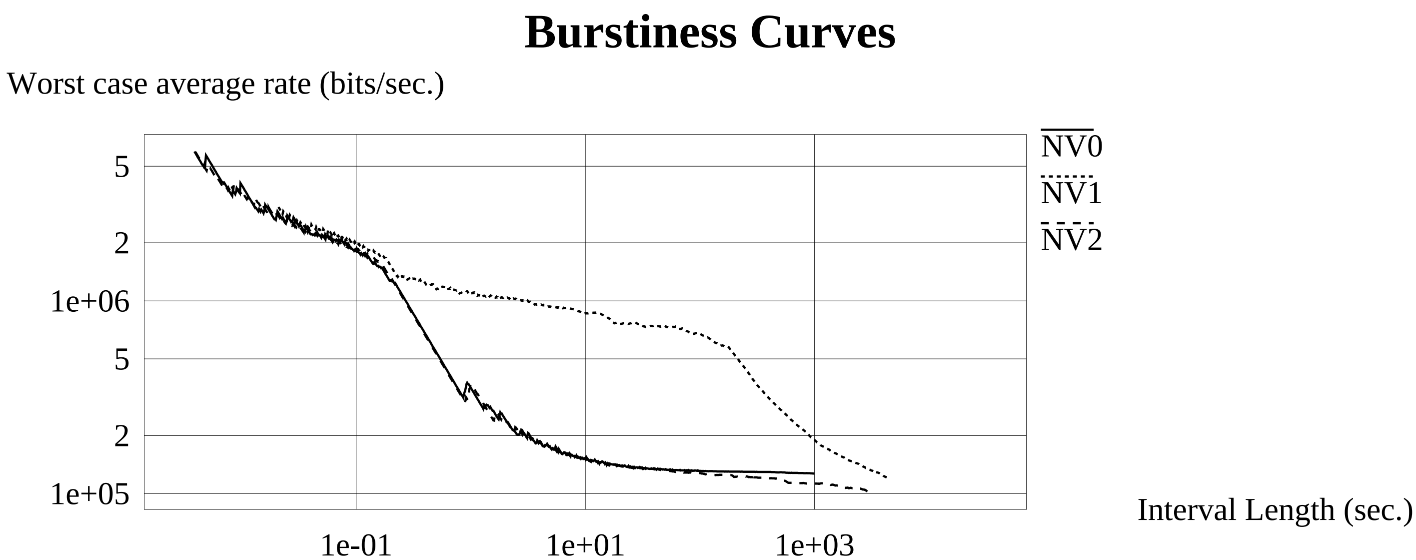

Figure 10: Burstineee curves for NV traffic

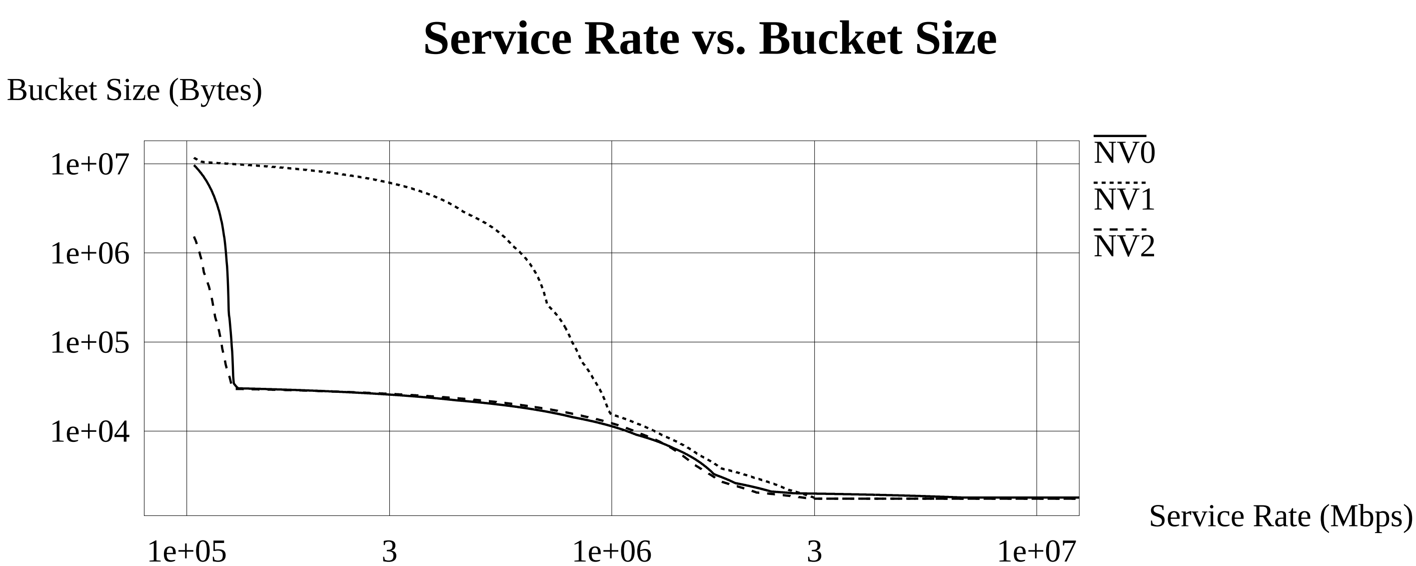

Figure 11: Sigma-Rho curves for NV traffic

which is being sent. Constant bit rate over a long time scale is achieved by compromising the quality of the video being sent. Whenever NV needs to reduce its bit rate, it reduces the frame rate of the video being sent. For a NV program sending at a target rate of 128 Kbps the appropriate service rate is 130 Kbps and bucket size is 30 K-Bytes.

7 Discussion

Sigma-Rho curves for all the three kinds of video traffic have knees. If the chosen service rate for leaky bucket is less than that of the knee point, then the bucket size becomes too large to be an accurate description of the traffic. On the other hand if the service rate is larger than that of the knee then the reduction in bucket size is small. The tradeoff of decreasing the bucket size by increasing the service rate above the service rate of the knee point may not be justified. Therefore the knee points in the sigma-rho curves give appropriate leaky bucket parameters to characterize the video traffic.

The appropriate service rate for JPEG compressed video is close to its peak rate. The service rate for MPEG compressed video is approximately equal to the peak rate of MPEG stream after smoothing. The service rate of NV traffic is slightly more than its target sending rate. For the trace NV1 the target sending rate was increased for 3 minutes. Thus the service rate for NV1 is a little more than its maximum target sending rate.

The parameter bucket size depends heavily upon the way the application delivers data to the network. In case of JPEG and MPEG compressed video, we assume that the application delivers data in form of frame at a constant frame rate. Therefore the bucket sizes obtained are close to the size of the largest frame present in the stream (30 - 80 K-Bytes). If the application decides to break each frame into packets and send the packets uniformly over one frame time then the bucket sizes will be different. The bucket size can be made arbitrarily small by suitably delaying the packets. The maximum delay incurred in this process is given by $\frac{\sigma}{\rho}$. For JPEG streams this delay found to be 33 ms, which is equal to one frame time. This means that if the burst size of JPEG video is reduced then the largest frame will be sent uniformly over one frame time. This delay for the MPEG trace LECTURE.mpg is 380 ms, which corresponds to about 6 frame times. The smoothing delays for NV is found to be 1.85 seconds, which is close to the off periods of NV.

Burstiness function characterizes the burstiness of the traffic at various time scales. JPEG video exhibits no burstiness on small time scales. This can be seen in the burstiness curves for JPEG traces by zero slopes for small intervals. However JPEG video is bursty over longer time scales. The burstiness curves for JPEG traces have a negative slope in the corresponding region.

MPEG video is bursty on short as well as long time scales. The burstiness curves for MPEG video have a high negative slope for small intervals. For larger intervals the magnitude of the slope is relatively small. This implies that MPEG video is more bursty on small time scales than on large time scales.

Traffic generated by NV is highly bursty over short periods of time but is very smooth when averaged over an interval of 10 seconds or more. This is reflected in the burstiness curves by a large negative slope for small intervals and zero slope for intervals greater than 10 seconds. The burstiness curve for the trace NV1 is different because in case of NV1, the target sending rate was increased for 3 minutes. As a result NV1 has burstiness over a long time scale also. The burstiness curve for NV1 has a knee at interval of length 3 minutes. The slope of the burstiness curve is steeper towards the right of knee. This implies that NV1 is more bursty when observed over intervals of length a little more than 3 minutes, as compared intervals of length slightly less than 3 minutes.

8 Conclusion and Future Work

Long traces of real video compressed using three different compression algorithm namely JPEG, MPEG and NV have been studied in this paper. The paper characterizes the traces using a leaky bucket model and shows a way of choosing appropriate leaky bucket parameters. Burstiness function is used to characterize the burstiness of the traffic at various time scales. The impact of burstiness on network congestion is discussed.

It is found that the leaky bucket parameters for constant quality JPEG and MPEG compressed video depend upon the actual video streams. For JPEG streams the appropriate service rate was found close to the peak bit rate of the compressed video. As a result characterizing JPEG video using leaky bucket descriptor is not expected to give any statistical multiplexing gains. The service rate of the JPEG compressed streams under study varied from 7 to 18 Mbps. The bucket sizes were found to be between 30 and 80 K-Bytes. MPEG compressed video could be characterized at a service rate lower than its peak rate. The service rates for our traces were found to be 1.05 and 1.65 Mbps which is 2-4 times the eventual average rate. The corresponding bucket sizes were approximately equal to 50 K-Bytes. Due to the rate control built in the software, the traffic generated by NV can be characterized independent of its contents. The service rate depends upon the target sending rate set on the program. For a target rate of 128 Kbps, the required service rate was found to be 130 Kbps and the bucket size was 30 K-Bytes.

It is shown how burstiness curves can be used to characterize the burstiness of the traffic at different time scales. Using these burstiness curves, JPEG compressed video has shown to exhibit low burstiness over small time scales, whereas MPEG and NV traffic shows high burstiness. On long time scales, JPEG and MPEG video exhibit burstiness whereas NV shows no burstiness.

Short term burstiness which is a property of the coding algorithm is easier to handle. A traffic having short term burstiness can be made smoother at the cost of additional delays. Even if the traffic is not smoothed the worst case buffer requirements to serve such a traffic at a constant bit rate is small. Long term burstiness is more difficult to handle.

One way to handle long term burstiness is by not having it. Encoders are being designed which supoport constant bit rate operation. Since bit rate is a property which depends upon motion and complexity of the video being encoded, such encoders will operate at a constant bit rate by reducing the picture quality of the scenes which have high motion or information content (for eg. scenes showing a black board full of text in a class room). On the other hand the scenes which don't have much information (scenes showing a clean black board) would be compressed at extremely good picture quality. This kind of behaviour is certainly undesirable. One might want to limit the peak rate of the video to a high value by designing apropriate encoders, but in normal mode of operation one would like to have a video whose quality remains fairly constant.

Another approach for wide area networks might be to reserve the bandwidth at peak bit rate for connections carrying video. Policing in this case would involve peak rate enforcement, but finding the peak rate for interactive video in advance may be difficult. The problem can be solved if the encoder limits the peak rate of its output stream. The quality of scenes requiring higher bit rate can be compromised. The long term burstiness is found to be between 2 and 4 for our traces. By reserving the bandwidth at peak rate, the video connections can utilize the network at 25-50%. If the non real-time data traffic which can adapt to congestion in the network is able to take rest of the available bandwidth then higher network utilizations can be achieved.

One may want to reserve bandwidth somewhere between peak and average rate and "hope" that statistical multiplexing among various video streams will smooth the aggregrate bit rate. It could be hard to deal with the congestion as all the clients can potentially send their data at peak rate. Multiplexing behaviour of constant quality video needs to be studied in more detail.

Adapting the sending rate depending upon feedback from the network has also been proposed [14]. Network signals congestion to the application and based on this congestion signal the application sending video adapts to a different rate by having a different compression ratio. This scheme may work well for interactice video which encodes the video on the fly but is difficult for stored video because changing the bit rate of stored video may involve additional processing. The additional complexity of this scheme may not be justified.

9 Acknowledgements

I would like to thank Prof. Domenico Ferrari who provided me with all the facilities needed to carry out this work. Without his encouragement it would have been impossible to do this work. I would thank S. Keshav motivated me to write this paper. I would like to express my gratitude to Bruce Mah who provided me with the traces of the video conferences held on XUNET. I would also like to thank him for giving comments on the manuscript. I am also thankful to Steve McCanne who provided me the necessary code to use DEC JVIDEO board. I would thank Amit Gupta for valuable discussions regarding the work. Thanks to Puneet and Rosen who helped me in writing this paper. I would acknowledge Mark Garrett for making his Star Wars data publicly available for research purposes.

References

[1] ISO/IEC JTC 1/SC 29/WG 11. Final text for ISO/IEC 11172-2, information technology - Coding of moving pictures and associated audio for digital storage media at up to about 1,5 Mbits/s - part 2: Video, May 1993.

[2] D. Luan A. Eckberg and D. Lucantoni. Bandwidth management: a congestion control strategy for broadband networks - characterizing the throughput-burstiness filter. In *Proceedings of the ITC Specialist Seminar*, September 1989.

[3] David Clark, Scott Shenker, and Lixia Zhang. Supporting real-time applications in an integrated services packet network: Architecture and mechanism. In *Proceedings of SIGCOMM'92*, Baltimore, Maryland, August 1992.

[4] Rene L. Cruz. A calculus for network delay, part I : Network elements in isolation. *IEEE Transaction of Information Theory*, 37(1):114–131, 1991.

[5] Rene L. Cruz. A calculus for network delay, part II : Network analysis. *IEEE Transaction of Information Theory*, 37(1):131–141, 1991.

[6] Domenico Ferrari and Dinesh Verma. A scheme for real-time channel establishment in wide-area networks. *IEEE Journal on Selected Areas in Communications*, 8(3):368–379, April 1990.

[7] Alexander G. Fraser, Chuck R. Kalmanek, A.E. Kaplan, William T. Marshall, and R.C. Restrick. Xunet2: A nationwide testbed in high-speed networking. In *Proceedings of INFOCOM*, pages 582–589, Firense, Italy, May 1992.

[8] Ron Frederic. An internet posting: Some info on the nv 2.x packet format, December 1992.

[9] G. Rigolio G. Gallassi and L. Fratta. ATM: bandwidth assignment and bandwidth enforcing policies. In *Proceedings of IEEE Globecom'89*, November 1989.

[10] D. Le Gall. MPEG: A video compression standard for multimedia applications. *Communications of the ACM*, 34(4):305–313, April 1991.

[11] Mark William Garrett. *Contributions Toward Real-Time Services on Packet Switched Networks*. PhD dissertation, Columbia University, 1993.

[12] J. Hyman, A. Lazar, and G. Pacifici. Real-time scheduling with quality of service constraints. *IEEE Journal on Selected Areas of Communications*, pages 1052–1063, September 1991.

[13] Sugih Jamin, Scott Shenker, Lixia Zhang, and David Clark. An admission control algorithm for predictive real-time service. In *Proceedings of the Third International Workshop on Network and Operating System Support for Digital Audio and Video*, San Diego, CA, November 1992.

[14] H. Kanakia, P.P. Mishra, and A Reibman. An adaptive congestion control scheme for real-time packet video transport. *Computer Communication Review*, 23:20–31, October 1993.

[15] E. Cavallero M. Butto and A. Tonietti. Effectiveness of the leaky bucket policing mechanism in ATM networks. *IEEE Journal on Selected Areas of Communications*, 9(3):335–342, 1991.

[16] M.R. Macedonia and D.P. Brutzman. Mbone provides audio and video across the internet. *Computer*, 27(4), 1994.

[17] P. Pancha and M. El Zarki. A study of MPEG coding for vbr video transmission. *IEEE Communication Magazine*, 1994.

[18] William B. Pennebaker and Joan L Mitchell. *JPEG Still Image Data Compression Standard*. Van Nostrand Reinhold, 1993.

[19] Erwin P. Rathgeb. Policing of realistic vbr video traffic in a ATM network. *International Journal of Digital and Analog Communication Systems*, 6:213–226, 1993.

[20] Gregory K. Wallace. The JPEG still picture compression standard. *Communications of the ACM*, 34(4):31–44, April 1991.

[21] Hui Zhang. *Service desciplines for packet-switching integrated-services.* PhD dissertation, University of California at Berkeley, 1993.