

Telecommunications Standardization sector
Study Group 15
Experts Group for Video Coding and Systems
in ATM and Other Network environments

Document AVC - 690
Version 1
1-4 November 1994
Singapore

Source: AT&T Bell Laboratories (Authors: Hemant Kanakia, Partho P. Mishra, Amy R. Reibman)

Title: Packet video transport in ATM networks with single-bit feedback

Purpose: Information

1. INTRODUCTION AND SUMMARY

The ATM forum recently agreed that a rate-based traffic management mechanism would be used to support Available Bit Rate (ABR) services. This contribution describes simulation results of a rate-based feedback mechanism for congestion control with real-time ABR video traffic. The results show that the rate-based mechanism is very useful for the transport of ABR video streams. We show that the visual quality of the video sequences degrades gracefully when network congestion occurs, due either to increases in the bandwidth requirements of existing video streams or because of new video streams switching on.

Feedback based congestion control mechanisms have been traditionally considered for data traffic management, where the main goal is to reduce packet losses while providing high throughput and low delay for each user. In this contribution, we examine the effectiveness of a rate-based traffic management for video traffic. The primary goal of a traffic management scheme for video traffic from the users' perspective is to maintain good perceptual quality especially during periods of network congestion. The primary goal from the perspective of a network operator is to make efficient use of network resources such as buffers and link capacities. We show that both these goals can be achieved with a rate-based traffic management scheme.

We note that for ABR services, the network will use a policer/enforcer to ensure that a source is using the feedback information as agreed. Therefore, it will be necessary for a video codec using ABR to be able to receive the feedback signals from the network and to operate its rate controller in an agreed-to manner.

Compressed video traffic has been traditionally transported as a Constant Bit Rate stream over a phone line. However, it makes better sense to generate Available Bit Rate (ABR) video streams for transport over cell switched or packet switched networks. This approach allows the end-to-end delay to be reduced, makes more efficient use of network's resources, and can potentially result in constant quality images, which are more pleasing to the human eye. However, transporting ABR video streams can result in periods of congestion during which the bandwidth requirements of active video streams exceed the available network capacity. A promising approach to congestion control for ABR video traffic is to use feedback from the network to modulate the bandwidth requirements of an ABR video stream. A detailed discussion of various possible approaches for transporting ABR video over packet switched networks is given in [1]. In this contribution, we consider feedback signals that are transmitted by the network using the EFCI bit and are reflected by the receiver back to the transmitter using resource management (RM) cells.

The performance of the single-bit feedback control mechanism is studied by using actual video traces generated by an MPEG-1 encoder to drive simulations. This allows us to study the effects of traffic overloads on the perceptual quality of video transmissions. Since widely accepted analytical models of video sources, particularly for entertainment applications, are not yet available, we chose not to use any such analytical models for traffic generation process. The main performance metrics used in this study are the signal to noise ratio (SNR) and the

perceptual quality of the video sequences.

The SNR of a single video frame is calculated by comparing the decoded frame with the original unencoded image. The average SNR is computed by averaging the per-frame SNR over the entire sequence of frames. The other performance measures used are the queue occupancy, the link utilization and the number of packets lost. The queue occupancy affects the end to end delay which is important for interactive applications while the utilization of bottleneck resources affects the efficiency with which network resources are used. The packet losses provide a crude indication of the degradation in signal quality.

The eight video sequences used in this study have a picture size of 352 by 240 pixels and a frame rate of 30 frames per second. The sequences used include bicycles, autumn-leaves, and ferris-wheel. The video traces used in these experiments were produced with $Q_p=4$. In the experiments in which the feedback control schemes were used, Q_p values were reduced to between 6 and 9 to allow a reduction in the sending rate. The data rate for the coded video sources, when they are not throttled due to network conditions, varies between 2 - 8 Mbits/sec at 30 frames/sec.

The experiments use the network configuration, shown in Figure 1. The link bandwidths are all 100 Mbits/sec except for the bottleneck link which is 50 Mbits/sec. The link propagation delays are chosen so that the round trip propagation delay is 42 ms, equivalent to the coast to coast propagation delay in the US. The receiver buffers up to 100 ms of video data in its playback buffer.

In all the simulations, multiple video sources share a common path. The multiplexing helps one study the robustness and the fairness of the (distributed) feedback control mechanism. We use bursty data sources with Poisson arrivals to provide cross-traffic on each link. The aggregate bandwidth of these sources is about 50 Mbits/sec and packets from these data sources are serviced with the same priority as those from video traffic sources. The traffic from data sources is not flow-controlled. These data sources allow us to study the effect of uncontrolled sources with stochastic packet arrivals on the service quality seen by the controlled connections; this also serves to test how well the service rate and buffer estimation techniques work.

Congestion is induced by reducing the bandwidth of the S4-H2 link from 50 Mbits/sec to 30 Mbits/sec. This abrupt reduction in bandwidth occurs 2000 ms after the start of the simulation and lasts until the end of the simulation run. The change in bandwidth is designed to stress the transient overload characteristics of the congestion control scheme and approximates a situation where there is a reduction in capacity for existing connections as a result of new video connections or an increase in the sending rate of several existing connections due to scene changes.

In the first group of experiments, each of the connections starts at intervals of 33 ms (roughly one frame interval). There is 200 Kbytes of buffers shared among all the active connections at each of the output queues. The queue occupancy threshold is 25 Kbytes. Figure 2 shows the evolution of the queue occupancy and the link utilization at the bottleneck. These curves show that there is some packet loss following the reduction in the bandwidth. The figure also demonstrates the presence of the characteristic oscillations caused by delayed feedback information. Due to the large buffers the only effect of the oscillation in the queue occupancy is to make the delay more variable - this has little effect in our experiments because of the playout buffering. The average SNR values for each of the video sources is shown in Table 1. Figures 3 and 4 show the variation in the per frame SNR for the adaptive scheme for the bicycles sequence. In Figure 3, note that because of the interframe coding used to generate the sequences, the signal quality stays degraded even past the point where packets are not lost any more. This situation is corrected only when a refresh frame (I frame) is sent out. In table 1, we note that one is not able to achieve complete fairness, with the degradation in average SNR values ranging from about 1dB for the horse rider sequence to 3 dB for the ferris wheel sequence.

Video Sequence	Random	Rate Based	Perfect
	Loss	Feedback	Channel
autumn	30.15	36.61	37.30
ball of wool	28.33	37.55	39.42
bicyclists	24.86	35.14	37.91
birches	35.02	34.50	37.21
ferris wheel	25.69	36.20	39.30
horse riders	28.39	37.84	38.83

Table 1: Average SNR values of video streams with large buffers and out of phase sources

A measure of how well the rate based scheme does is provided by the comparison with the SNR values obtained (column 1) when no feedback is generated and packets are lost when the bandwidth at the bottleneck is exceeded. The perfect channel column refers to the situation where there is always adequate bandwidth available and consequently no cell loss occurs.

In the second set of experiments, only 50 Kbytes of buffers are available at a switch with the other experimental parameters being the same as in the previous experiment. The thresholds are also picked to be proportionally lower. At a speed of 50 Mbits/sec and a round-trip time of 42 ms, 50 Kbytes of buffers can store less than one fifth of the round-trip delay-bandwidth product. Due to the smaller number of buffers in this scenario, a greater number of packets are lost when the bandwidth reduces abruptly. Moreover due to the oscillations in buffer occupancy packets are again lost later. As a result, the degradation in signal quality is slightly worse than when a greater number of buffers are available. The average SNR values observed for each of the sequences is shown in Table 2.

Video Sequence	Random	Rate Based	Perfect
	Loss	Feedback	Channel
autumn	29.96	36.29	37.30
ball of wool	28.04	37.52	39.42
bicyclists	24.69	34.34	37.91
birches	24.60	32.28	37.21
ferris wheel	24.91	35.55	39.30
horse riders	28.31	38.06	38.83

Table 2: Average SNR values of video streams for video streams with small buffers and out of phase sources

In the third set of experiments, there are 400 buffers available per output queue as in the first set but the video connections start up at the same time. As a result, the I frames of all the video sources will arrive in-phase. This experiment is used to study if the synchronization among traffic sources produces a significant degradation in the performance. The phase synchronization was observed to cause some oscillations in the buffer occupancy; however, the effect was not significant enough to affect the average signal quality.

Video Sequence	Random	Rate Based	Perfect
	Loss	Feedback	Channel
autumn	29.86	36.58	37.30
ball of wool	27.42	37.83	39.42
bicyclists	24.94	35.20	37.91
birches	25.30	35.01	37.21
ferris wheel	26.41	37.25	39.30
horse riders	27.94	37.92	38.83

Table 3: Average SNR values of video streams for video streams with large buffers and in phase sources

Although we have used a relatively simple experimental configuration, the simulated network is adequate for highlighting the impact of congestion on the perceptual quality of individual video streams. Our simulations are designed to study a worst-case scenario where an abrupt drop in the capacity of the bottleneck link occurs. In real networks, where the switching on off new connections gradually reduces the capacity available for existing connections, the transients would have a less serious impact on the perceptual quality.

The results of our experiments for a number of other topologies, traffic models and various other single-bit rate based schemes are omitted here due to space constraints. Some of these results are reported in [1] and [2]. In particular, we have investigated the impact of congestion due to scene changes on the perceptual quality using very long video traces. We have also examined the impact of longer propagation delays ranging from 20 to over 200 ms round-trip time. In our experiments, we have looked at a wide range of parameter values for the rate-based proposal and found that the performance is quite robust over a wide range of parameters. All of these results support the basic assertion that the rate-based traffic management scheme is quite effective for ABR video streams.

As pointed out earlier, the rate-based proposal as described provides unfair service in that there is unequal degradation in the perceptual quality for video streams sharing the same bottleneck link. We argue that using a round-robin scheduling discipline is not expected to improve fairness due to the time-varying bandwidth requirements of ABR video sources. We have investigated modifications that improve on the fairness of the feedback scheme for real-time video traffic. The "fair" version does not require any changes to the traffic management proposal under consideration.

[1] Kanakia, Mishra and Reibman, "An Adaptive Congestion Control Scheme for Real-Time Packet Video", Proceedings of ACM Sigcomm, Sep 1993, San Francisco, California.

[2] Kanakia, Mishra and Reibman, "Packet Video transport in ATM networks with single-bit feedback", Presented at the Sixth International Workshop on Packet Video, Sep 1994, Portland, Oregon.




Figure 1: Network configuration used in the simulations




Figure 2: Buffer occupancy and link utilization at the bottleneck




Figure 3: SNR plots for bicycle sequence with a) the rate-based scheme b) with the random cell loss model.