

Multiplexing Video Traffic using Frame-Skipping Aggregation Technique

Alan Yeung

Soung C. Liew

Department of Information Engineering
The Chinese University of Hong Kong
Shatin, N.T., Hong Kong.

Department of Information Engineering
The Chinese University of Hong Kong
Shatin, N.T., Hong Kong.

Abstract

MPEG, the Moving Picture Experts Group, is a popular compression standard used in many multimedia applications. In video distribution systems, a CBR communications channel is often shared among several VBR MPEG video streams. Since there are fluctuations of data rates in video streams, traffic congestion may occur and video quality may be affected. We propose a frame-skipping aggregation technique to multiplex several MPEG streams (outside the network) to feed into a CBR channel (within the network). The salient feature of the scheme is that the transmission of B frames will be skipped (outside the network) before the onset of traffic congestion so as to minimize the image quality degradation. Our result shows that using this technique, more users can be supported compared to regular video transmission systems. This technique allows us to adopt a simple call admission strategy in which video requests are granted based on their mean rates.

1 Introduction

In video transmission, several video streams are often multiplexed in order to use the channel efficiently. Generally, it is better to multiplex video traffic outside the communication network than inside [1, 4, 6]. There are two basic ways of multiplexing video, namely, lossy and lossless. Lossy multiplexing discards data when the channel bandwidth is insufficient, while lossless multiplexing avoids data loss by delaying data delivery. Besides multiplexing, compression is also used to reduce bandwidth requirements. The Moving Picture Experts Group (MPEG) coding scheme [2, 5] is a popular video compression scheme. There are three types of encoded pictures (frames) in an MPEG stream: I (intra-coded), P (predicted), and B (bi-directionally predicted) frames. These frames form a pre-defined sequence called a group of pictures (GOP) [2, 5], which specifies the display or-

der of frames. Some possibilities are IBBPBBPBB, IBPBPBPB, etc. The GOP usually (although not necessarily) repeats itself for the whole video duration. In this paper, we propose a frame-skipping lossy multiplexing technique for video streams. When the bandwidth is not enough for the transmission of frames from every stream for some period of time, the transmission of B frames will be skipped to conserve bandwidth. We choose to skip B frames rather than I or P frames because B frames will not be referenced by other frames during decoding, and thus the image quality of other decoded frames will not be affected. This paper proposes a particular method of skipping B frames that is based on the status of the displaying process at the receivers.

2 Framework of Frame-Skipping Lossy Aggregation

2.1 Video frames delivery using round-robin scheduling

The architecture of the frame-skipping video transmission system is depicted in Fig. 1. Frames belonging to different video sources are forwarded to a central controller, the output of which is a transmission channel shared by all the video streams. The controller guarantees that the amount of video data sent does not exceed the pre-defined bandwidth requirement. Here, we define a time-slot as the period of time when one frame should be displayed in each receiver, and the data limit as the number of bytes the transmission link can support in a time-slot. The data limit divided by a time-slot is equal to the pre-defined bandwidth. In each time-slot, the controller subtracts the frame size of the frame to be sent from the data limit. If the data limit is not exceeded, this frame can be sent successfully. The controller then tries to send another frame of another video stream in a round-robin fashion until the data limit becomes zero. The sequence of sending

frames in the frame-skipping aggregation controller is depicted in Fig. 2.

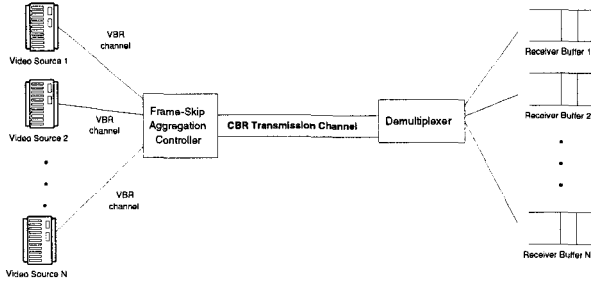


Figure 1: The Frame-Skipping Aggregation System

2.2 Underflow safety margin on receiver buffers

In order to prevent underflow at receiver buffers, we monitor the interval between the currently displayed frame number and the latest received frame number. Let us refer to this as the underflow safety margin (USM). When this interval drops below a pre-defined threshold (USMT), we start to skip transmitting B frames at the controller in the next time-slot. When a B frame is skipped, the receiver will repeat the display of the previous frame in place of the B frame. Setting this safety margin enables us to skip B frames at the transmitting side in advance before the onset of underflow which in turn minimizes image degradation and preserves continuous video display. If we were to react only after underflow has occurred, we would simply have no choice but to skip whatever frames (I, P or B) that are next in line for transmission.

A frame-skipping scenario at the controller is explained with the help of Fig. 2. Assume that the USM drops below USMT in previous time-slot. The controller starts to skip sending the B frames in this time-slot. For example, suppose that the I, P and B frames are of size 50000, 20000 and 7000 bytes respectively. If 100000 bytes are the data-limit, I₁ and P₂ can be sent successfully. After this, only 30000 bytes are left for the remaining frames. Since B₅ is skipped¹, the next frame sent is P₁. Similarly, B₃, B₆, B₁ and B₄ are skipped accordingly. Finally, the remaining 10000 bytes are used to send half of P₄. Note that if we do not adopt the frame-skipping scheme, we can only send the first four frames and about half of B₃. By virtue of the frame-skipping scheme, the USM can be built up quickly. The frame-skipping mechanism will be in effect until the USM goes above the USMT.

¹For simplicity, the bandwidth taken by the necessary headers and DC terms are ignored

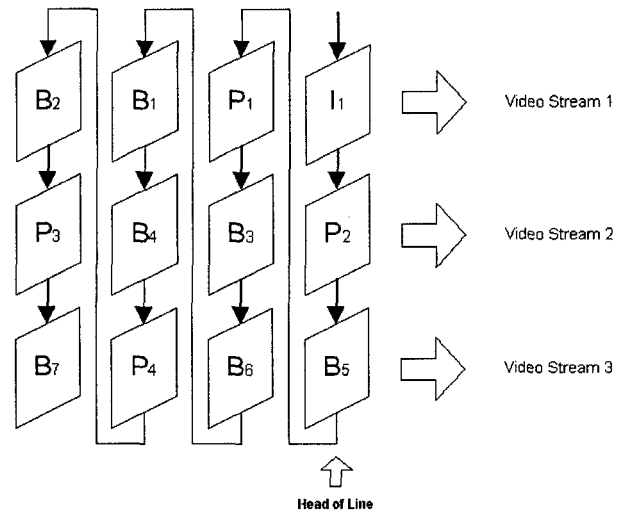


Figure 2: The frame sending and frame skipping sequence

2.3 Frame skipping

The synchronization of video and audio display may be lost if we skip a B frame entirely. To solve this problem, several strategies are being investigated. One possible solution is to send the necessary headers and the DC-terms of the DCT coefficients in all coded blocks, of the B frame being skipped, as well as the audio data, to the receiver for decoding. The audio synchronization is retained and the decoder just treats these 'skipped' frames as ordinary frames. A subjective assessment test² of the video sequences using frame skipping technique described above is conducted and the result is very satisfactory.

3 Simulation Results

To demonstrate the merit of the frame-skipping aggregation technique, we have performed experiments and gathered some simulation results based on a video distribution system described in section 2. The performance metric is the maximum number of supportable video streams. A number of video streams is said to be supportable if 1) the receiver buffer does not underflow and 2) the frame-skipping rate is not more than some pre-defined percentage. The video stream used is taken from the film Star War [3, 7], and it lasts about 2 hours long and is coded as MPEG-1 with 24 frames per second and GOP IBBPBBPBBPBB. The bandwidth of the channel is represented as the number of bytes allowed to be sent in one time-slot (1/24 sec.). The mean frame size of the video stream is found to

²Owing to limited space, it is not included here.

be 15614 bytes. For bandwidth of 100000 bytes/time-slot,

$$\text{Ideal number of streams supportable} = \frac{100000}{15614} = 6.4$$

The ideal number of streams supportable is used as a reference of performance in the experiment. In the first experiment, for the frame-skipping systems, we restrict the maximum percentage of the frames skipped in each individual video streams to 1% and 2%. The results of the first experiment are shown in Fig. 3. For all the three systems (No Frame-skip, 1% Frame-skip, 2% Frame-skip), we set the receiver buffer size to 10 frames. For the two frame-skipping systems, the USMT of the receiver below which frames will be skipped is 2 frames time. It is observed that the number of streams that can be supported by a system that skips a maximum of 1% or 2% B frames is greater than or equal to that of a regular system with no frame-skip for all bandwidth. Note that the percentage of frame-skip corresponds to quality of service. As shown, the number of streams supportable when 2% frame-skip is allowed is quite close to the ideal one.




Figure 3: Number of streams supportable for different systems for different bandwidth (display time safety margin=2 frames time)

The second set of experiment explores the characteristics of frame-skipping aggregation itself. Specifically we investigate the relationship of three parameters of the system, namely, number of streams supported, USMT of receiver buffer for frame-skipping to take place and percentage of frames skipped. The video streams used in this experiment are the same as those of the previous ones. For bandwidth of 1Mbytes

(per time-slot), the results are shown in Fig. 4. No frames are skipped when there are 57 or fewer streams for all USMTs since the bandwidth is quite sufficient. As the number of streams increases, the frame-skip percentage increases generally. This is because we have to accommodate more data in the channel in each time-slot, and the USMT is reached more frequently which leads to more frames being skipped. It can also be observed that for a fixed frame-skipping requirement, we can support more streams by decreasing the USMT. However, it should be understood that the trade-off is that there will then be a greater risk of buffer underflow.




Figure 4: Frame-Skipping Aggregation characteristic for bandwidth = 1M bytes/time-slot

4 Conclusion

We have proposed a new multiplexing method for video traffic using the frame-skipping aggregation technique. It is a lossy multiplexing scheme which discards B frames data in MPEG video streams just before the onset of traffic congestion; the reason is that after traffic congestion has occurred, we may not selectively skip B frames anymore. We have investigated the performance of this technique relative to that of the regular video transmission system without the frame-skipping mechanism on the basis of the number of supportable video streams. Our result shows that more streams can be supported if we are willing to tolerate a small portion of frames being skipped. In practice, regardless of the higher level of efficiency, some scheme to wisely select the data to be skipped will be needed anyway in many situation. Different video streams may have drastically different traffic charac-

teristics, and it is generally difficult to prevent data loss altogether. Our scheme makes it viable to adopt a simple call admission strategy in which requests are granted based on mean rate because it minimizes the detrimental effect of data loss which cannot be prevented with such a simple call admission strategy. We have also looked into the characteristics of the frame-skipping mechanism by varying its associated parameters. The result obtained is useful for setting the correct parameter values of the frame-skipping system in the determination of these parameters. This is important for optimal system performance in terms of the number of streams supportable and the frame-skip percentage.

References

- [1] Chi-yin Tse and Soung C. Liew, "An Integrated Video Compression and Multiplexing Scheme for Broadband Networks", *IEEE Infocom 95*, pp.439-446
- [2] Simon S. Lam, Simon Chow, and David K. Y. Yau, "An Algorithm for Lossless Smoothing of MPEG Video"
- [3] M. W. Garrett, W. Willinger, "Analysis, Modeling and Generation of Self-Similar VBR Video Traffic", *Proc. ACM SigComm, London*, Sept. 1994
- [4] Hanford H. Chan and Soung C. Liew, "Lossless Aggregation: A Scheme for Transmitting Multiple Stored VBR Video Streams over a Shared Communications Channel without Loss of Image Quality", *IEEE Journal of Selected Areas on Communications*, Aug. 1997
- [5] Didier Le Gall, "MPEG: A Video Compression Standard for Multimedia Applications", *Communications of the ACM*, April 1991
- [6] Chi-yin Tse and Soung C. Liew, "Video Aggregation: Adapting Video Traffic for Transport over Broadband Networks by Integrating Data Compression and Statistical Multiplexing", *IEEE Journal on Selected Areas in Communications*, Aug. 1996
- [7] MPEG1 *Star War* trace available via anonymous ftp site - <ftp://thumper.bellcore.com/pub/vbr.video.trace>