Proceedings, Fifth IFIP Int'l Workshop on Quality of Service, May 1997.

Adaptive video applications for non-QoS networks

S. Jacobs and A. Eleftheriadis
Department of Electrical Engineering
Columbia University, New York, NY 10027, USA
{sej,eleft}@ee.columbia.edu

Abstract

We propose an architecture which can support video services in networks which have no Quality of Services (QoS) guarantees. We demonstrate that there is a need for such services today and also in the future, when QoS will be built into most networks. We then describe our current Internet-based system, which combines both image processing and networking techniques in order to provide good quality video without harming the performance of the Internet. These techniques are shown to be superior to previous work for the environment of video services operating in the Internet.

Keywords
Internet video, adaptive video, non-QoS networks

1 INTRODUCTION

The underlying technologies of today's Internet are not sufficient to support Quality of Service (QoS) guarantees. The evolution of the base technologies of the Internet could result in any number of possibilities, including ATM backbones, IP switching, or fully deployed ATM networks. Regardless of the specifics, the general consensus is that the network infrastructure of the future will have QoS and that users will be able to request connections with or without QoS.

Certainly connections which reserve resources will demand a higher cost and users may not always want to pay this extra cost for a particular video service. If a user is watching a movie, perhaps it would be worthwhile. However, just browsing a video database may not warrant the extra cost since the user may browse many video streams before deciding on the one that is being sought. Thus, some video services may not warrant the extra cost of QoS.

Another motivation for being able to provide good quality video services without QoS is that some networks may never be able to provide QoS guarantees. Wireless networks fit into this category because there may be no way to guarantee a certain bandwidth when so many other variables exist that affect the throughput on the wireless link.

The technique for developing video services without QoS involves an explicit attempt at avoiding network congestion. Clearly, network congestion hurts the performance of all users of the network. The goal is to send only the data that can fit into the network at a particular time. This requires both a networking and an image processing approach. From the networking perspective, an estimate of the available bandwidth in the network must be found. Then, from the image processing perspective, a technique for shaping the compressed video into that available bandwidth is necessary. In the next sections we discuss previous work as well as our proposal for both the networking and image processing techniques.

2 BANDWIDTH ESTIMATION

As real-time services have become more prevalent over the last few years, there has been a growing concern that the current Internet infrastructure may not be able to support them, leading perhaps, to a ‘‘congestion collapse.” This is a valid concern in general since many real-time services send their bits through the network without concern for congestion control or avoidance.

Network bandwidth estimation has been explored extensively. Many algorithms provide a somewhat different view of the state of the network (Brakmo 1994, Busse 1995, Ramakrishnan 1990, Sakatani 1994). Despite the different techniques

for finding the available bandwidth, probably the most important factor in designing a congestion avoidance algorithm is that it does not degrade the quality of other traffic on the network. For this reason, it is important to ensure that the technique used provides fairness among the different traffic types, i.e. ftp, http, audio, video, etc.

Our architecture uses the TCP Congestion Control (TCP-CC) algorithm as a congestion indicator. TCP is not used for transport and retransmissions are performed selectively. An algorithm which always retransmits increases the delay, which in general, is unacceptable for real-time applications and certainly for video.

The reason for using TCP-CC as an indication of available bandwidth is that TCP streams work well together; today's Internet is proof of that. They operate according to a greedy but "socially-minded" and cooperative algorithm which attempts to get as much bandwidth as possible, but backs off substantially during congestion. Using TCP-CC can make real-time traffic look as harmless as a file transfer to the network and still maintain relatively low delay (Jacobs 1996).

The Internet is only one type of network which has no QoS guarantees. ATM-ABR provides a guaranteed minimum QoS, but informs the server when more bandwidth is available. In this case, the bandwidth estimate is explicit in that the network furnishes the exact bandwidth which is available. Also, as mentioned earlier, wireless may never have QoS guarantees.

3 ADAPTABLE MEDIA

Just as the Internet is one type of network which has no QoS guarantees, video is only one type of media which can be adapted on the fly. For environments where minimal delay is more important than perfect quality, the rates of both audio and still images can be adapted. Thus, although this paper focuses on video and the Internet, some of the concepts also apply to any combination of non-QoS network and adaptable media, i.e. wireless and audio.

Techniques for adapting compressed video on the fly have been limited in the past. One technique is to adjust the quantization parameters at the encoder based on the state of the network (Ortega 1995). Although this works quite well for live video, it cannot work for precompressed streams. Another commonly used technique for changing the bit rate of video is to drop frames (Chen 1996). However, frame dropping alone is a crude technique which provides only a coarse approximation to the available bandwidth since the smallest unit of data which can be removed is an entire frame. Although subjective tests have not been completed, it seems intuitive that very low frame rate video is perceived as less valuable for many applications.

To solve this problem, we developed a technique called Dynamic Rate Shaping (DRS) (Eleftheriadis 1995A). DRS provides the ability to dynamically change the bit rate of a precompressed stream. In its simplest form, DRS selectively drops

coefficients from the MPEG-1 or MPEG-2 bit stream which are least important in terms of image quality.

DRS operates on the compressed video bit stream, eliminating DCT coefficients run-lengths. The coefficients to be eliminated are determined using Lagrangian optimization, resulting from an operational rate-distortion formulation. The problem is complicated by the fact that MPEG coding utilizes predictive and interpolative modes for motion compensation, and thus any modification of the bit stream will result in error propagation. We have experimentally shown in (Eleftheriadis 1995B) that, if decisions within each frame are optimal, then ignoring the accumulated error does not impact the quality in any way. Thus the algorithm can operate in a memoryless mode which has significantly less complexity, while achieving essentially optimal (within 0.3 dB) performance. Also, DRS is much less complex than a complete decoder, making it feasible to run on a video server.

DRS can meet any reasonable bandwidth estimate exactly. This means that the bandwidth estimate is more fully utilized. It also maintains the original frame rate of the video. Lastly, DRS decouples the adaptable media from the encoder so that it can be used for both live and stored video streams. A combination of DRS and frame dropping will probably yield better perceptual results than either technique working alone, especially for large rate reductions.

4 CONCLUDING REMARKS

We have presented a novel architecture for supporting video services in a non-QoS network, namely the Internet. Several experiments have been performed to verify performance in both a controlled environment and in external wide area connections in the Internet. Our results show that our proposed architecture provides a stable system with good quality video which does not degrade the quality of other Internet traffic.

5 REFERENCES

Brakmo, L. S. and O'Malley, S. W. (1994) TCP Vegas: New techniques for congestion detection and avoidance. *SIGCOMM Symposium on Communications Architectures and Protocols*, London, United Kingdom, 34-35.

Busse, I. Deffner, B., and Schulzrinne, H. (1995) Dynamic QoS control of multimedia applications based on RTP, *First International Workshop on High Speed Networks and Open Distributed Platforms*, St. Petersburg, Russia.

Chen, Z., Tan, S. M. Campbell, R. H. and Li Y. (1996) Real Time Video and Audio in the World Wide Web, *World Wide Web Journal*, Vol 1.

Eleftheriadis, A. and Anastassiou, D. (1995) Constrained and General Dynamic Rate Shaping of Compressed Digital Video, Proceedings, *2nd IEEE International Conference on Image Processing*, Washington, DC, III.396-399.

Eleftheriadis, A. and Anastassiou, D. (1995) Meeting Arbitrary QoS Constraints Using Dynamic Rate Shaping of Coded Digital Video. *Proceedings, 5th International Workshop on Network and Operating System Support for Digital Audio and Video, Durham*, New Hampshire, 95-106.

Jacobs, S. and Eleftheriadis, A. (1996) Providing Video Services over Networks without Quality of Service Guarantees, *World Wide Web Consortium Workshop on Real-Time Multimedia and the Web*, Sophia-Antipolis, France.

Ortega, A. and Khansari, M. (1995) Rate Control for Video Coding over Variable Bit Rate Channels with Applications to Wireless Transmission, *Proceedings of the 2nd IEEE International Conference on Image Processing*, Washington, DC.

Ramakrishnan, K. K. and Jain, R., (1990) A binary feedback scheme for congestion avoidance in computer networks, *ACM Transactions on Computer Systems*, vol. 8, 158-181.

Sakatani, T. (1994) Congestion avoidance for video over IP networks, *Multimedia Transport and Teleservices Proceedings*, 256-273.

6 BIOGRAPHY

Stephen Jacobs received both the B.S. and M.S. in electrical engineering from Columbia University in 1994 and 1995, respectively. He also holds a B.A. in Physics from Bard College. He is currently a Ph.D. candidate in the Electrical Engineering Department at Columbia University and a Graduate Research Assistant in the Image and Advanced Television Laboratory. His research interests include adaptive multimedia protocols and applications for networks without quality of service guarantees. In 1996, Mr. Jacobs was awarded the Kodak Fellowship. He is a student member of the IEEE.

Alexandros Eleftheriadis received the Diploma in EECS from the National Technical University of Athens, Greece, in 1990, and the M.S., M.Phil., and Ph.D. degrees in Electrical Engineering from Columbia University, New York, in 1992, 1994, and 1995 respectively. Since 1995 he has been an Assistant Professor in the Department of Electrical Engineering at Columbia University. His research interests are in the areas of visual information representation and compression, video communication systems, and the fundamentals of compression. Dr. Eleftheriadis is a member of the IEEE, ACM, ANSI, and he serves as the Editor of the forthcoming MPEG-4 Systems standard (ISO 14496-1).