MAXIMIZING USER UTILITY IN VIDEO STREAMING APPLICATIONS

Carlos E. Luna and Aggelos K. Katsaggelos

Department of Electrical and Computer Engineering
Northwestern University
2145 North Sheridan Road, Evanston, IL 60208
Phone: (847) 491-3039
Fax: (847) 491-4455
Email: {carlos,aggk}@ece.nwu.edu

ABSTRACT

In this paper, we study the design tradeoffs involved in video streaming in networks with QoS guarantees. We approach this problem by using a utility function to quantify the benefit a user derives from the received video sequence. This benefit is expressed as a function of the total distortion. In addition, we also consider the cost, in network resources, of a video streaming system. The goal of the network user is then to obtain the most benefit for the smallest cost. We formulate this utility maximization problem as a joint constrained optimization problem. The difference between the utility and the network cost is maximized subject to the constraint that the decoder buffer does not underflow. We present a deterministic dynamic programming approach to find the optimal tradeoff for both the Constant Bit Rate (CBR) and Renegotiated Constant Bit Rate ($RCBR$) service classes. Experimental results demonstrate the benefits and the performance of the proposed approach.

1. INTRODUCTION

In this paper, we study the design tradeoffs of video quality and network cost in video streaming applications. There are two major requirements for the video streaming process to be successful. First, the amount of lost packets in the network must not be excessive. Second, the delay experienced by a video frame as it traverses the network must be constant if the display and encoder are to operate at the same frame rate [1, 2]. Both lost packets and packets that do not arrive on time to be displayed result in degradation of the video quality.

Networks that can provide Quality of Service (QoS) guarantees, such as ATM, have been developed in response to the demand for the integrated delivery of new services such as video conferencing and video streaming. In this environment, the user negotiates an agreement with the network where the network guarantees the QoS for the traffic generated by the user as long as this traffic conforms to a specified traffic profile, i.e., the number of bits the user is allowed to transmit at each time instance. The user is charged for the network resources needed to meet the QoS guarantees. Thus, video streaming applications have to be evaluated along two dimensions: received video quality and cost in network resources.

There are two types of traffic profiles: static and dynamic. Static traffic profiles remain in effect for the duration of the connection [3]. Service classes that use static traffic profiles, such as the Constant Bit Rate (CBR) service class, require the user to have accurate *a priori* knowledge of the traffic which will be generated. Dynamic traffic profiles, can be adapted during the connection. Renegotiated services have been introduced in order to accomodate applications in which the traffic characteristics are time varying and cannot be accurately described by a static traffic profile. Variable bit rate video is a good example of this type of traffic. Renegotiated services allow the network to achieve more efficient utilization of network resources by adapting to changing traffic conditions [3].

Renegotiated Constant Bit Rate ($RCBR$) service has been proposed as a simple renegotiated service to accomodate traffic with multiple time scale burstiness, such as VBR video. In this service, the traffic profile is given as a piece-wise CBR profile. That is, the source transmits at a constant rate for a certain period of time and then it switches to a different constant rate. The authors in [4] have proposed a possible implementation of $RCBR$ in ATM networks and in Integrated Services internets.

We use a utility function to quantify the benefit the user derives from a received video sequence. The goal of the user is then to maximize the difference between the utility derived from the received video sequence and the cost in network resources. We present an algorithm to jointly and optimally solve this problem CBR and $RCBR$ classes of service. Two step approaches that consider the network and the codec problem separately have been proposed in the literature [2, 5, 3, 4].

The rest of this paper is organized as follows: In the next section, we present our framework for maximizing user utility. In Section 3 we present the solution to the problem for both the $RCBR$ and CBR streaming situations. In Section 4 we present experimental results illustrating the performance of the algorithm, and in Section 5 we present our conclusions.

2. USER UTILITY MAXIMIZATION

Each user s of the network has a utility function U_s that is used to quantify the satisfaction of the user with the network service [6]. We measure this satisfaction in terms of both the received video quality and the degree to which the real-time display requirements have been met. In an environment with guaranteed QoS, the network can guarantee the delay that traffic generated by the network user will experience as it traverses the network. If the traffic generated by the user meets the constraints imposed by the contract with

This work was supported in part by a grant from SBC.

the network, then only the encoded video quality is a concern.

As an example, we will use a logarithmic utility function given by,

$$U_s(D) = U_{max} - \alpha \log(D), \tag{1}$$

where D is the total distortion and U_{max} is the utility the user derives when the original video sequence is viewed. The parameter α can be interpreted as the cost per dB of distortion introduced at the encoder.

Network cost depends on the class of service that is used and on the type of traffic that the user generates. The network only guarantees the QoS level if the user generates traffic that conforms to the negotiated traffic profile. Given the traffic profile and the class of service, the network has to allocate resources, such as bandwidth and buffer space in order to meet the QoS guarantees. The user is then charged a cost C_s for the ability to use these resources.

The goal of the network user is then to obtain the most utility for the smallest possible price. This is achieved when we maximize the difference between U_s and C_s. We therefore can formulate the problem we consider here as follows. Given a set of desired QoS guarantees, find an encoded video sequence and a traffic profile that result in,

$$\max\{U_s - C_s\}, \tag{2}$$

subject to the real time constraints mentioned in the introduction. In the next section we will study the implications of these constraints.

3. PROBLEM FORMULATION

We consider the simplified system shown in figure 1. We will treat this system as a discrete time system. Raw video frames, $X(i)$, are fed into the encoder which selects a quantizer Q_i from a finite set of quantizers Q. Encoding frame i with quantizer Q_i results in an encoded frame $\hat{X}(i)$ which is of size $R(i)$ bits and results in distortion $D(i)$. The encoded frames are fed into the encoder buffer which is drained at variable rate $C(i)$.

At the receiver, the arriving data is stored in the decoder buffer. The decoder waits a fixed amount of time, Δ, and then begins decoding the received video sequence. We assume here that large enough buffers are available at the encoder and decoder and thus we only need to worry about the effects of buffer underflow. We can guarantee that this does not occur as long as we meet the following constraint

$$B_e(i) \leq \sum_{j=i+\Delta+1}^{i+2\Delta} C(j), \tag{3}$$

where $B_e(i)$ is the encoder buffer occupancy at time i and the upper bound is known as the effective buffer size [2]. This effective buffer size is the maximum buffer occupancy allowed at the encoder buffer such that all the bits can be delivered at the receiver without violating the end-to-end delay constraint.

The cost for a $RCBR$ traffic profile is given as,

$$C_s = \gamma \sum_{i=1}^{T} C(i) + \phi \sum_{i=2}^{T} (1 - \delta(C(i-1), C(i))), \tag{4}$$

where γ is the cost per unit flow of reserving bandwidth for one frame time, ϕ is the cost per renegotitation and $\delta(.)$ denotes the delta function. The duration of the connection T is given by $T = N + \lceil \frac{B_e(N)}{C(N)} \rceil$, where N is the number of frames to be encoded, and the second term represents the number of times the channel will be used to empty the encoder buffer once the last video frame, N, has been encoded.

In practice, renegotiations are limited to occur every M frame intervals. If the renegotiation interval, M, is equal to the length of the sequence, N, then we have the problem of video streaming using a CBR service class. Thus, the solution of the CBR problem is a special case of the $RCBR$ solution we present next.

If we consider the utility function given in Eq.(1) and the $RCBR$ class of service, we can pose the utility maximization problem as a minimization problem. We need to find an $RCBR$ traffic profile and $\{Q_i\}_{i=1}^{N}$ that minimize the cost function

$$\alpha \log(D) + \gamma \sum_{i=1}^{T} C(i) + \phi \sum_{i=2}^{T} (1 - \delta(C(i-1), C(i))), \tag{5}$$

subject to the constraint

$$B_e(i) \leq \Delta C(i). \tag{6}$$

Note that this constraint is not the same as the more general constraint of Eq.(3) which incorporates knowledge of future time instants. The constraint of the minimization problem of Eq.(5) can be interpreted as follows: Frames cannot be encoded to violate the current parameters of the traffic contract. This means that if we were to continue transmitting at rate $C(i)$, frame i would meet its delay requirement. Thus, we are restricting ourselves to the situation where we have no knowledge of the outcome of future renegotiations. Therefore, our solution can provide a fair comparison with practical systems in which the encoding and renegotiation process are done online.

3.1. Algorithm

We consider here the case of intra frame encoding only. Our approach will be to find for each allowable traffic profile, the quantizer sequence that produces the minimum cost given by Eq. (5). Let Γ be the set of allowable channel rates. We can represent the set of allowable traffic profiles using a tree as shown in figure 2.

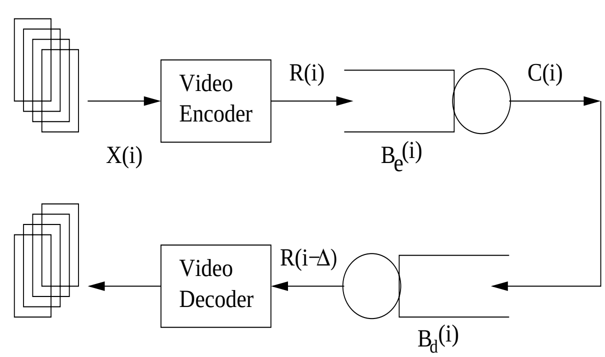

Fig. 1. Simplified Video Streaming System

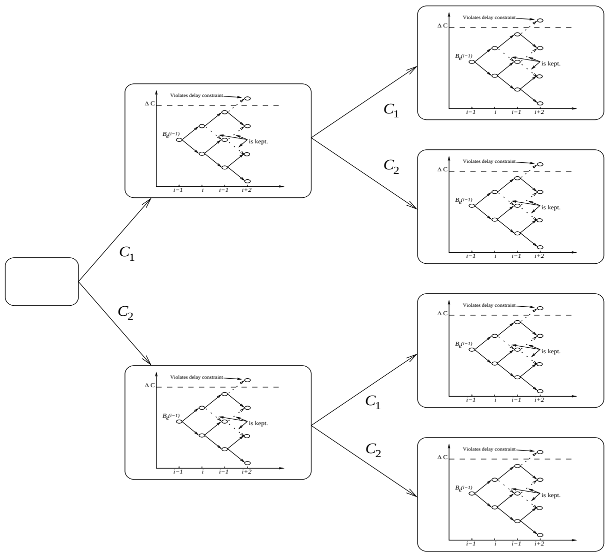

Fig. 2. Trellis Tree

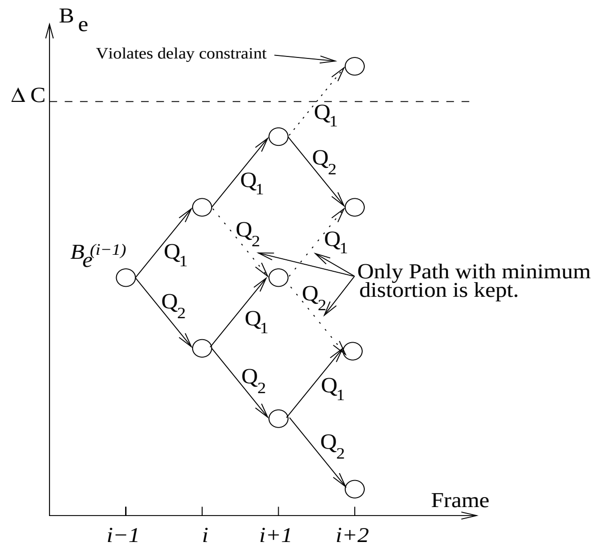

Fig. 3. Node in the Trellis Tree

The number of levels in the tree is $\lceil N/M \rceil$. Each node of the tree represents a trellis that is used to optimally encode the corresponding M frames of the video sequence for transmission at a constant rate (refer to figure 3). The trellis tree is formed by using the following procedure:

1. Let $j = 0$. Start with a root node that has a one node trellis. This trellis has a single node with $B_e(0) = 0$. We initialize the distortion to node this to zero.
2. For each node in level j of the tree, create a child in level $j + 1$ of the tree for each channel rate $C_j \in \Gamma$.
3. For each node in level $j + 1$, create a trellis of M stages (or finish the encoding process) according to the procedure presented below.
4. Let $j = j + 1$. If there are frames left to code then go back to $step$ 2.

Each node in the tree of figure 2 is a trellis which is grown according to the following procedure (See figure 3):

1. Let $i = 0$. For each node in the first stage with buffer occupancy $B_e(i)$ set the distortion equal to the corresponding node in the last stage of the parent node in the trellis tree.
2. For each node in stage i with buffer occupancy $B_e(i)$, create a branch for each possible quantizer Q_k. Each branch connects a node in stage i to a node in stage $i + 1$ with buffer occupancy $B_e(i + 1) = B_e(i) + R^{Q_k}(i + 1) - C_j$, if $0 \leq B_e(i+1) \leq \Delta C_j$, with $R^{Q_k}(i+1)$ the bit rate generated by encoding frame $i + 1$ using quantizer Q_k. The distortion associated with this branch is $D^{Q_k}(i + 1)$.
3. If two or more branches arrive at the same node, only keep the branch resulting in the smallest cumulative distortion.
4. $i = i + 1$. Go to $step$ 2 if $i < M$.

4. EXPERIMENTAL RESULTS

In this section we present experimental results to illustrate some important aspects of the problem and the provided solution. We use an H.261 coder [7] operating in intra frame mode to encode the first 100 frames of the “foreman” sequence in $QCIF$ format. The time scale corresponds to the time needed to display a frame of video. The set of four possible quantizer step sizes, $\{8, 10, 12, 31\}$, was used in our experiments.

We consider two experiments with an $RCBR$ service. We let $\Gamma = \{10000, 15000, 20000, 25000\}$ bits per frame, the renegotiation interval is set to $M = 50$ frames, and $\gamma = 0.0002$, and $\phi = 5$ in the cost function of Eq. (4). We use a logarithmic utility function as in Eq. (1) with $U_{max} = 500log(255^2N)$. Figure 4(a) shows the cost function for this experiment. In figure 4(b) we can see the $(U_s - C_s)$ function for this experiment. We can see here that the best profile corresponds to streaming in CBR at 20000 bits per frame.

In the second example, we set $\Gamma = \{10000, 20000, 25000\}$ and the renegotiation interval to $M = 30$ frames. The cost and utility functions are the same as before. Figure 5(a) shows the results of this experiment. In the vertical axis we show the resulting $(U_s - C_s)$. The horizontal axis corresponds to each of the 81 allowable traffic profiles. The traffic profiles are ordered arbitrarily in this figure.The optimal solution corresponds to the traffic profile numbered 43, shown in figure 5(b). This example illustrates the benefits of renegotiations.

Clearly, the results depend on he parameters γ and ϕ in Eq.(4), as well as the renegotiation interval M. Some parameter values were chosen in the experiments for demonstration purposes only, and without loss of generality. Actual parameters values clearly depend on the specific application.

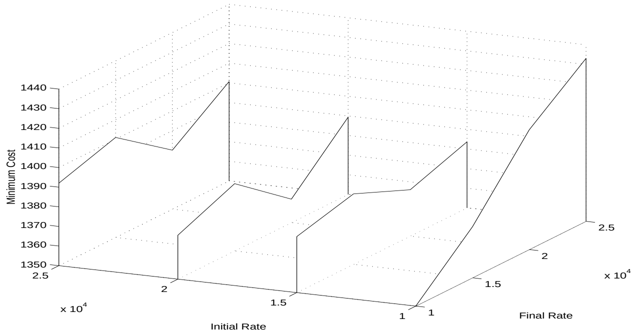

(a)

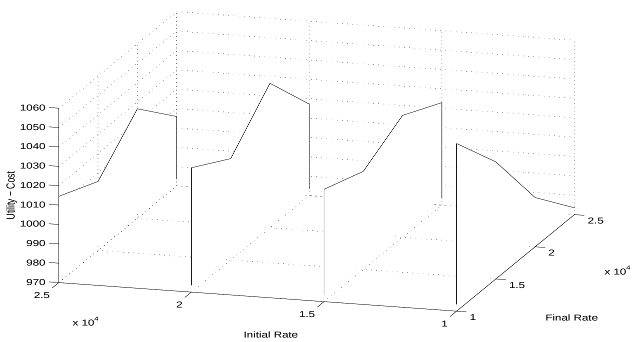

(b)

Fig. 4. Results for Experiment 1

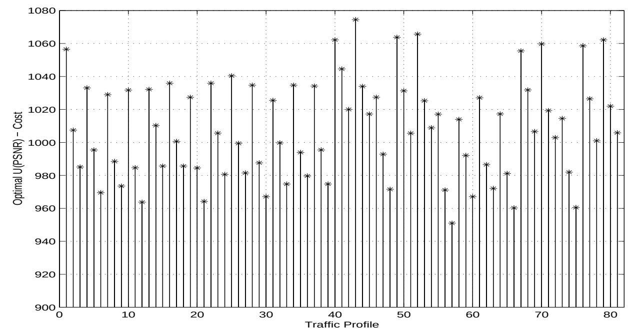

(a)

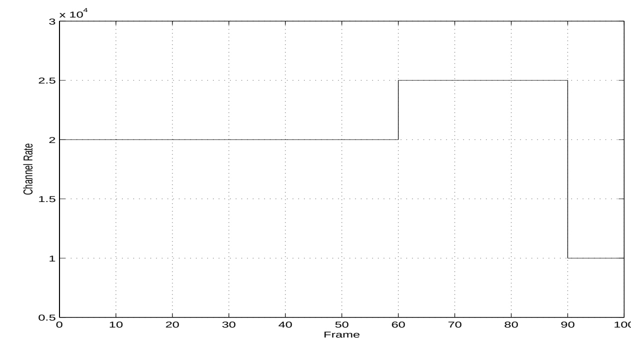

(b)

Fig. 5. Results for Experiment 2

5. CONCLUSIONS

In this paper, we have studied some of the design tradeoffs of video streaming systems in networks with QoS guarantees. We have approached this problem by using an utility function to quantify the user benefit derived from the video quality of the received video sequence. We have measured the performance of a video streaming system in terms of the difference between the user benefit and the network cost. Our goal has been to maximize this difference.

We have formulated this utility maximization problem as a joint problem where we maximize the video quality of the received sequence for a given traffic profile.

We have solved this problem for classes of service that use both static and dynamic traffic profiles. Specifically, we have considered CBR and $RCBR$ service classes. Our experimental results suggest that when renegotiations are taken into account by the video encoder, the received video quality can be improved. This can prove particularly useful when we want to transmit video encoded using a predictive encoder and when the application does not have an accurate representation of the source.

6. REFERENCES

[1] T. V. Lakshman, A. Ortega, and A. Reibman, "VBR Video: Tradeoffs and Potentials," *Proceedings of IEEE*, vol. 86, no. 5, pp. 952–973, May 1998.

[2] C. Hsu, A. Ortega, and A. Reibman, "Joint Selection of Source and Channel Rate for VBR Video Transmission Under ATM Policing Constraints," *IEEE Journal on Selected Areas in Communications*, vol. 15, no. 6, pp. 1016–1028, Aug 1997.

[3] E. W. Zegura, S. McFarland, and O. Parekh, "A Survey and New Results in Renegotiated Service," *Journal of High Speed Networks*, vol. 6, no. 3, pp. 197–206, 1997.

[4] M. Grossglauser, S. Keshave, and D. Tse, "RCBR: A simple and Efficient Service for Multiple Time Scale Traffic," *IEEE/ACM Transactions on Networking*, vol. 5, pp. 741–754, December 1997.

[5] J. J. Chen and D. W. Lin, "Optimal Bit Allocation for Coding of Video Signals over ATM networks," *IEEE Journals on Selected Areas in Communications*, vol. 15, no. 6, pp. 1002–1015, August 1997.

[6] Frank Cowell, *Microeconomic Principles*, Oxford University Press, 1986.

[7] The Portable Video Research Group (PVRG), "PVRG-64 CODEC v. 1.2," *ftp://havefun.stanford.edu/pub/p64v1.2.tar.Z*, 1991.