

Join-the-Shortest-Queue Prefetching for VBR Video on Demand ¹

Martin Reisslein

Department of Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104
reisslei@seas.upenn.edu
<http://www.seas.upenn.edu/~reisslei>

Keith W. Ross

Department of Systems Engineering
University of Pennsylvania
Philadelphia, PA 19104
ross@seas.upenn.edu
<http://www.seas.upenn.edu/~ross>

January 1997

¹This work has been supported partially by NSF grant NCR93-04601. A shorter version of this paper appeared in the Proceedings of the 1997 IEEE International Conference on Networking Protocols, Atlanta, GA, October 1997.

Abstract

We present a high-performance prefetching protocol for the delivery of prerecorded VBR video from a server across a packet-switched network to a large number of clients. Not only does the protocol give constant perceptual quality and almost 100% link utilization, but it also allows for immediate commencement of the video upon user request and near instantaneous response to pause/resume and temporal jumps. The protocol requires (1) that each client have a small amount of memory dedicated to the application (2) that there is one bottleneck shared link between the server and the clients. The protocol is based on the observation that there are frequent periods of time during which the shared link's bandwidth is under utilized. During these periods the server can prefetch frames from any of the ongoing videos and send the frames to the buffers in the appropriate clients. The server chooses prefetched frames according to a join-the-shortest-queue (JSQ) policy. We present simulation results of our prefetch policy that are based on MPEG encoded traces. Our simulations show that JSQ prefetching performs favorably as compared with the smoothing-based protocols in the existing literature.

1 Introduction

We present a high-performance prefetching protocol for the delivery of video on demand (VoD) from a server across a packet-switched network to a large number of clients. The protocol assumes that the videos resident on the video server are variable-bit-rate (VBR) encoded. Not only does this protocol give constant perceptual quality and almost 100% link utilization, but it also allows for immediate commencement of the video upon user request and near instantaneous response to interactive actions (pause/resume and temporal jumps).

To achieve this high performance our protocol has two requirements. First we require all of the clients to have a small amount of memory dedicated to the VoD application. Second, as shown in Figure 1, we require that there be at most one bottleneck shared link between the video server and the clients. If the clients are connected to an ADSL residential access network, then this second requirement can be achieved by attaching the video server directly to the ADSL central office; in this case the shared link is the link from the server to the ADSL central office. If the clients are connected to cable, then the second requirement can be achieved by attaching the video server directly to the cable headend.

The user's client could be a television with a set-top box capable of performing buffering and decoding, or it could be a household PC. The video server could be a cache containing the week's most popular videos. A central repository could multicast the videos to this and other caches over the Internet using the traditional best-effort service [4].

Our protocols explicitly assume that the videos are VBR encoded with high peak-to-mean ratios. The motivation for our approach is that, for the same perceived video quality, Constant Bit Rate (CBR) encoding produces an output rate significantly higher than the average rate of the corresponding VBR encoding for action movies [3]. CBR traffic allows for nearly 100% link utilization; the number of connections that can be carried over a link of given capacity is roughly the link capacity divided by the CBR rate (assuming homogeneous connections). The number of VBR connections that can be transmitted simultaneously is the achievable link utilization multiplied by the link capacity divided by the average rate of the VBR video stream. Therefore schemes for transmitting VBR encoded video that achieve high average link utilizations while keeping losses at a negligible level, can allow for significantly more video connections than does CBR video.

Our protocol achieves the constant perceptual quality, responsiveness to user interactivity, and high link utilizations by exploiting two special properties of the prerecorded video: (1) for

each video, the traffic in each video frame is known before the video session begins; (2) while the video is being played, some of the video can be prefetched into the client memory. It is this second property – the ability to prefetch a portion of any video – that is particularly central to our high-performance protocol.

Our protocol is based on the observation that, due to the VBR nature of the multiplexed traffic, there will be frequent periods of time during which the shared link’s bandwidth is under utilized. During these periods the server can prefetch video frames from any of the ongoing videos and send the prefetched frames to the buffers in the appropriate clients. With this prefetching, many of the clients will typically have some prefetched reserve in their buffers. Our protocol also specifies the policy for how the server selects the prefetched frames. This policy is the join-the-shortest-queue (JSQ) policy, which can be roughly described as follows: within each frame time the server repeatedly selects frames from the connections that have the smallest number of prefetched frames in their client buffers. The JSQ policy creates a *buffer pooling effect* so that the system behaves as if the individual client buffers are aggregated into one large buffer which is shared by all the clients. Our empirical work with public-domain traces indicates that prefetching combined with the JSQ policy gives dramatic reductions in packet loss. In particular, if each client dedicates a small amount of buffer capacity to the VoD application, this scheme can multiplex a large number of connections over the shared link and have negligible playback starvation. In this paper we also examine several refinements and variations of our JSQ prefetching policy, and we develop schemes for selectively discarding frames for MPEG encoded video when playback starvation is unavoidable.

With JSQ prefetching, through buffer pooling, the connections collaborate in order to minimize the overall packet loss. This collaboration among the connections contributes significantly to the protocol’s outstanding performance. We also present numerical results which show that *Optimal Smoothing*, a non-collaborative prefetching policy, can have packet loss that is several orders of magnitude higher than that of JSQ prefetching for a wide range of buffer sizes.

This paper is organized as follows. In the following subsection we give a brief literature review of traffic management schemes for prerecorded VBR-encoded video. In Section 2 we give a description of our VoD architecture and introduce prefetching. In Section 3 we present our JSQ prefetch policy; Section 3 includes numerical results for public domain traces and a discussion of efficient implementation of the JSQ policy. In Section 4 we show how our VoD protocol allows for user interactivity with minimal delays; numerical results for interactive actions are presented. In Section 5 we compare the performance of JSQ prefetching with that

of Optimal Smoothing [31, 37]. In Section 6 we explore a variation of our JSQ policy – the packet-based JSQ prefetch policy. In Section 7 we examine MPEG compression and propose several policies for selective discard when loss is unavoidable. In Section 8 we discuss how JSQ prefetching can be employed in residential broadband access networks using cable or ADSL. In Section 9 we conclude and discuss future work.

1.1 Literature Review

Work in the area of traffic management of prerecorded video falls into two classes: traffic management issues in the disk subsystem of the video server (e.g., see [12] [13] [26] [34] [36]) and traffic management issues in the network. Although the ideas of this paper may be useful for disk subsystem design, our focus is on management of prerecorded VBR traffic in the network. The papers on the management of prerecorded VBR traffic in the network fall into three categories: deterministic; deterministic with smoothing and/or prefetching; and probabilistic. The principal performance metrics for all of these schemes are average link utilization, initial delay, delays after interactive actions, and client buffer size.

The deterministic schemes send into the network the original VBR traffic, and admission control ensures that the delays never exceed a prespecified limit [35] [19] [16] [31] [23]. For highly variable VBR traffic, these deterministic schemes typically require large initial delays to achieve moderate link utilizations [23].

The deterministic schemes with prefetching and smoothing do not send the original VBR traffic into the network, but instead send some smoothed version of it. Admission control can then be based on the peak-rate of the smoothed trace. CRTT [22] is the extreme case of such a scheme, whereby the server transmits packets into a reserved CBR connection at the average rate of the video; such CBR connections are available from ATM and are being proposed for the Internet [32]. Although CRTT produces close to 100% link utilizations, it has potentially a long initial delay and potentially long delays after temporal jumps. Several independent research teams have proposed schemes whereby the server transmits the video at different constant rates over different intervals; these schemes vary in how the rates and intervals are chosen [14] [31] [24] [8] [9]. As compared with CRTT, these schemes trade off link utilizations and simplicity for lower initial delay and lower interactivity delays. In summary, none of the deterministic schemes (with or without prefetching and smoothing) allows for both high link utilizations and consistently high responsiveness (less than a few seconds) to interactivity.

For the probabilistic approaches, [28] considers sending the original VBR encoded video




Figure 1: Prerecorded videos multiplexed over a link of capacity R packets/sec.

into an unbuffered multiplexer. This scheme allows for responsive interactivity, but introduces packet loss whenever the aggregate transmission rate exceeds the link rate. In [28] we developed an admission control procedure based on a large-deviations approximation. In [14] and [37] related ideas are explored whereby the original traffic is first smoothed before it is statistically multiplexed at an unbuffered link; the statistically multiplexing of the smoothed traffic can substantially increase link utilization at the expense of small packet loss probabilities. All of these hybrid schemes perform non-collaborative prefetching; because JSQ is collaborative, it can have significantly better performance than these hybrid schemes.

2 Architecture Description

Figure 1 illustrates our basic model for VoD. The video server contains a large number of videos in mass storage. For notational simplicity, assume that each video consists of N frames and has a frame rate of F frames/sec. The videos are VBR encoded using MPEG 1, MPEG 2 or some other video compression algorithm. Let J denote the number of video connections in progress. Although some of the connections might be transmitting the same video, the phases (i.e., the start times) are typically different. The server packetizes the frames of the ongoing connections and then statistically multiplexes and transmits the packets into its link; for simplicity, we assume throughout that the packets are of fixed length. Let $x_n(j)$ denote the number of packets in the n th frame of the j th connection. Because the videos are prerecorded,

$(x_1(j), x_2(j), \dots, x_N(j))$ is fully known at connection establishment.

When a client requests a specific video, the server makes an admission control decision by deciding whether or not to grant the request. If it grants the request, a connection is established and the server immediately begins to transmit the connection's packets into the network. The connection's packets are transmitted in a fixed, predetermined order. When packets arrive at the client, they are placed in the client's prefetch buffer. The video is displayed on the user's monitor as soon as a few frames have arrived at the client.

Under normal circumstances, every $1/F$ seconds the client removes a frame from the prefetch buffer, decompresses it, and displays it. If at one of these epochs there are no complete frames in its prefetch buffer, the client loses the current frame; the client will try to conceal the loss by, for instance, redisplaying the previous frame. At the subsequent epoch the client will attempt to display the next frame of the video.

We denote R (in packets/sec) for the maximum transmission rate of the server. We denote $R(j)$ for the maximum reception rate of the j th client. Although our protocol allows for pause and temporal jumps, we will initially exclude these interactive features in order to simplify the discussion; thus N/F seconds elapse from when the user begins to watch the video until when the video ends. We shall also initially assume that all prefetch buffers are infinite and that $R(j) \geq R$ for $j = 1, \dots, J$.

To make these ideas a little more precise, we divide time into slots of length $1/F$. Let $p_l(j)$ be the number of frames in the prefetch buffer for connection j at the beginning of slot l . Let $\Delta_l(j)$ be the number of frames of connection j that arrive to the prefetch buffer during the the l th slot. At the end of each slot, one frame is removed from each prefetch buffer that has one or more frames. Thus

$$p_{l+1}(j) = [p_l(j) + \Delta_l(j) - 1]^+, \quad (1)$$

where $[x]^+ = \max(x, 0)$. Denote $\theta_l(j)$ for the frame number of connection j that is supposed to be removed at the end of slot l ; thus $\theta_{l+1}(j) = \theta_l(j) + 1$.

During each slot of length $1/F$ seconds the server must decide which frames to transmit from the J ongoing videos. The *prefetch policy* is the rule that determines which frames are transmitted in each slot. The maximum number of packets that can be transmitted in a slot is R/F (which for simplicity we assume to be an integer).

For each ongoing connection j the server keeps track of the the prefetch buffer contents $p_l(j)$; this can be done through the recursion (1) without communicating with the client.

Initially, we require the server to skip the transmission of the entire frame $\theta_l(j)$ whenever $p_l(j) + \Delta_l(j) = 0$. Thus, the server does not transmit a frame that will not meet its deadline at the client.

Before defining our prefetch policy, it is useful to introduce some more notation. Let $b_l(j)$ be the number of packets in the prefetch buffer for connection j at the beginning of slot l . Let $\Delta'_l(j)$ denote the number of packets of connection j that arrive to the prefetch buffer during the l th slot. These definitions imply

$$b_{l+1}(j) = [b_l(j) + \Delta'_l(j) - x_{\theta_l(j)}(j)]^+. \quad (2)$$

3 The JSQ Prefetch Policy

Our Join-the-Shortest-Queue (JSQ) prefetch policy attempts to balance the number of frames across all of the prefetch buffers. In describing this policy, we drop the subscript l from all notations. At the beginning of each slot the server determines the j^* with the smallest $p(j)$, transmits one frame from connection j^* and increments $p(j^*)$. Within this same slot the server repeats this procedure over and over again, at each iteration finding a new j^* that minimizes $p(j)$, transmitting a frame from connection j^* and incrementing $p(j^*)$.

Due to the finite transmission rate of the server, at some point the server must stop transmitting frames within the slot. To this end, let z be a variable that keeps track of the total number of packets sent within the slot; z is reinitialized to zero at the beginning of every slot. The stopping rule works as follows. Before transmitting a frame from connection j^* we check to see if

$$z + x_{\sigma(j^*)}(j^*) \leq R/F, \quad (3)$$

where $\sigma(j^*)$ is the frame of connection j^* that is being considered for transmission. If this condition holds, then we transmit the frame and update z ; otherwise, we do not transmit the frame, set $p(j) = [p(j) - 1]^+$ for $j = 1, \dots, J$ and recommence the procedure for the subsequent slot. This is our *basic stopping rule*; later we shall discuss a slightly more complicated stopping rule.

With prefetching, it is possible that all of a connection's frames have been transmitted but not all of its frames have been displayed. When a connection reaches this state, we no longer consider it in the above JSQ prefetching policy. From the server's perspective, it is as if the connection has been terminated.

3.1 Refinements of the JSQ policy

We now discuss a few important refinements of the JSQ policy. First, we introduce a refined stopping rule. Recall that during each slot the server transmits a sequence of frames until condition (3) is violated; once (3) is violated, the server does not transmit any more frames in the slot. An alternative stopping rule is to try to transmit more frames in the slot by removing from consideration the connection that violates (3) and finding a new j^* that minimizes $p(j)$. If the condition (3) holds with the frame from the new connection j^* , we transmit the frame, update $p(j^*)$, and continue the procedure of transmitting frames from the connections that minimize the $p(j)$'s. Whenever a frame violates condition (3), we skip the corresponding connection and find a new j^* . When we have skipped over all of the connections, we set $p(j) = [p(j) - 1]^+$ for $j = 1, \dots, J$ and move on to the next slot. This is our *refined stopping rule*. To reduce the online computational effort we can also, of course, consider rules which fall between the basic and refined stopping rules. For example we could use a rule which stops when condition (3) has been violated K times where $1 < K < J$.

The next refinement of the JSQ policy limits the number of packets an ongoing connection may have in its client's prefetch buffer. This important refinement is useful when the client for connection j , $j = 1, \dots, J$, has finite buffer capacity $B(j)$. This refinement works as follows. Suppose that the server is considering transmitting frame $\sigma(j^*)$ from connection j^* . Let $b(j^*)$ be the current number of packets in the prefetch buffer for connection j^* . It transmits this frame in the current slot only if condition (3) and the condition

$$b(j^*) + x_{\sigma(j^*)}(j^*) \leq B(j) \tag{4}$$

are satisfied. Condition (4) ensures that the server does not overflow the prefetch buffer for connection j^* . With this additional condition, we extend the definitions of the stopping rules in the obvious way.

The final refinement we consider in this subsection is applicable when $R(j) < R$ for at least one connection j . This situation occurs when the client's access link has limited bandwidth (as with ADSL) or when the the client has a limited packet processing capability. In this case, within each slot the server would keep track of $z(j)$, the total number of packets from connection j sent in the slot. Suppose the server is considering transmitting the frame $\sigma(j^*)$ from connections j^* . It transmits the frame in the current slot if and only if conditions (3), (4) and

$$z(j^*) + x_{\sigma(j^*)}(j^*) \leq R(j^*)/F \tag{5}$$

are satisfied. Condition (5) ensures that the server does not violate the reception constraint for connection j^* . Once again, we can extend the definitions of the stopping rules in the obvious manner.

We also mention that if the switch has buffer capacity for each connection, then it may be possible to ignore (5) even when $R(j) < R$. The idea is to store the traffic exceeding the capacity of the j th access link in the output buffer of connection j , and to feed the access link at rate $R(j)$. An upper bound on the switch buffer capacity necessary for connection j is $[R - R(j)]B(j)/R$. This expression is based on the worst case scenario that the prefetch policy dedicates the entire link capacity R temporarily to connection j in order to fill its prefetch buffer completely. This scenario occurs if connection j starts with empty prefetch buffer while all other connections have full prefetch buffers.

To simplify the discussion, for the remainder of this paper we shall assume that either (1) $R(j) \geq R$ for all $j = 1, \dots, J$, or (2) the switch has sufficient buffer capacity so that condition (5) can be ignored.

3.2 System Dynamics and Pooling

We now crudely describe the dynamics of the prefetch buffer contents. Let us make the assumption that whenever all N frames of a connection are displayed, the same user immediately requests a new connection; thus, with this assumption there are always J videos in progress. Let us make the realistic assumption that

$$\sum_{j=1}^J x_{avg}(j) < R/F, \tag{6}$$

where $x_{avg}(j)$ is the average number of packets in a frame in the j th connection. The condition (6) says that the long-run average aggregate consumption rate of the ongoing videos is less than the maximum server supply rate.

The JSQ prefetch policy will make most of the $p(j)$'s nearly equal to each other. Moreover, because of (6) there will be a general tendency for each of the $p(j)$'s to increase. In other words, there is typically a "pack" of $p(j)$'s with near equal values, with each $p(j)$ in the pack drifting towards $B(j)$. A $p(j)$ can separate itself from the pack if it hits $B(j)$ while the pack continues to gain higher occupancies. It will also break from the pack if the corresponding connection reaches the state of having all its frames transmitted. In this case, the $p(j)$ will descend by one in each slot until it hits zero; when it hits zero, $p(j)$ will quickly return to the pack.

Trace	Mean bits	Peak/ Mean	Avg. Frame Size			% of info		
			I	P	B	I	P	B
lambs	7,312	18.4	38,025	7,436	3,424	43	25	32
bond	24,308	10.1	83,293	41,427	10,513	29	42	29
terminator	10,904	7.3	37,387	14,120	6,387	29	32	39
mr.bean	17,647	13.0	75,146	18,275	10,215	35	26	39

Table 1: Statistics of MPEG-1 traces.

One important feature of the JSQ prefetching policy is that it creates a *pooling effect*, namely, the individual buffers act as a large collected buffer of capacity $B(1) + \dots + B(J)$. To see this, suppose that a large number of connections start simultaneously (therefore having few frames in their prefetch buffers) and the remaining connections have a large number of frames in their prefetch buffers. Also assume that the aggregate consumption rate temporarily exceeds R packets/sec. Then if the consumption rate across the connections that have just started is less than R , the JSQ policy will prevent frame loss as long as the high aggregate consumption rate does not persist for too long. This is because the prefetch policy will feed the connections that have just started until their prefetch buffers catch up with rest of the pack. Thus, the likelihood of cell loss in the near future typically depends on the $p(j)$'s only through the pooled buffer content $p(1) + p(2) + \dots + p(J)$.

3.3 Experimental Results

In this subsection we present the results of a simulation study for the JSQ prefetch policy described in the previous section. Throughout we use the refined stopping rule discussed in Section 3.1. Our simulation study makes use of MPEG 1 encodings of the four movies in Table 1. The associated traces, obtained from the public domain [29], give the number of bits in each frame. (We are aware that these are low resolution traces and some critical frames are dropped; however, the traces are extremely bursty. We have obtained similar results, not reported here, for *Star Wars* and *Oz*.) Each of the movies was compressed with the GOP pattern IBBPBBPBBPBB at a frame rate of $F = 24$ frames/sec. Each of the traces has 40,000 frames, corresponding to about 28 minutes. The mean number of bits per frame and the peak-to-mean ratio are given in Table 1. Table 1 also gives the average size of I, P and B frames and the percentage of encoding information carried in I, P and B frames. It can be argued that the average rate in bits/sec is lower than what we would expect for digital compressed video

(e.g., MPEG-2 video); for this reason, we have chosen a relatively small server transmission rate of 45 Mbps. We expect VoD systems in the future to have videos with larger average rates and a proportionally larger server transmission rate. In this scaling, the number of videos that can be multiplexed will be approximately constant. We furthermore assume that each packet consists of 512 bytes of payload and 40 bytes of overhead; therefore, $R = 81,521$ packets/sec.

We define the link utilization as the average number of packets per second, summed over all connections in progress, divided by R . In our experiments we use various mixes of the the four movies. Each of the mixes has a different link utilization. Our 90% link utilization consists of 52 lambs connections, 16 bond connections, 35 terminator connections, and 22 mr.bean connections. Our 95% link utilization consists of 55 lambs connections, 17 bond connections, 37 terminator connections, and 23 mr.bean connections. With these numbers, each of the four movies accounts for roughly one fourth of the link load.

In each realization of our simulation, we generate a random starting frame $\theta(j)$ for each of the J ongoing connections. The value $\theta(j)$ is the frame that is removed from the j th prefetch buffer at the end of slot 1. The $\theta(j)$'s are independent and uniformly distributed over $[1, N]$. All connections start with empty prefetch buffers at the beginning of slot 1. When the N th frame of a video is removed from a prefetch buffer, we assume that the corresponding user immediately requests to see the entire movie again. Thus, there are always J connections in progress.

In Figure 2 we set the buffer capacity of each client to 1 Mbyte and the link utilization to 95%. The figure shows the prefetch buffer contents in bytes for one of the lamb's connections over 120,000 simulated frame periods; the initial starting random phase for this connection is $\theta = 2,689$. At $l = 36,079$ frame periods all of the movie's frames have been transmitted to the prefetch buffer; this is why the buffer content drops to zero at $l = 37,311$, when the movie ends and starts over for the user. Note that when the movie restarts at $l = 37,311$ the JSQ prefetch policy fills the prefetch buffer for the movie in just a few frame periods. Also note that the buffer contents typically hug the 1 Mbyte buffer limit. Note that the buffer occupancy is nearly periodic, with period equal to the length of the movies (40,000 frames).

In Figure 3 we plot the 90 % confidence intervals of $P_{\text{loss}}^{\text{time}}$ for the prefetch buffer sizes ranging from 2 KBytes to 256 KBytes. $P_{\text{loss}}^{\text{time}}$ is the long-run fraction of frame periods for which loss occurs for at least one of the J ongoing videos; see [28] for more details. Each of the confidence intervals is based on 1000 replications. On this figure $P_{\text{loss}}^{\text{time}}$ is plotted for both 95% and 90% utilizations.




Figure 2: Prefetch buffer contents in a 1 Mbyte buffer in bits for a lamb's connection.




Figure 3: Confidence intervals of $P_{\text{loss}}^{\text{time}}$ for link utilizations of 90 % and 95 %. $P_{\text{loss}}^{\text{time}}$ is the fraction of frame times during which loss occurs.

Figure 3 shows the dramatic improvement in performance that comes from prefetching and the JSQ policy. Without any prefetching we have $P_{\text{loss}}^{\text{time}} = 0.29$ for 95 % utilization. By increasing the capacity only to 256 KBytes, the loss probability is reduced to $P_{\text{loss}}^{\text{time}} \approx 4.7 \times 10^{-6}$. By further increasing the buffer to 512 KBytes, no loss was observed for any of 8000 replications, corresponding to 3.2×10^8 frame periods! Similarly, for 90 % utilization we did not observe any loss for 1000 replications, corresponding to 4×10^7 frame periods, for 64 KByte prefetch buffer. Whenever loss occurred, the buffer contents at each of the clients was nearly zero, confirming the existence of a strong pooling effect. (The existence of pooling is further justified by the analytical work of Reiman [27]).

Figure 3 demonstrates that with a small prefetch buffer at each client the JSQ prefetching policy will allow for almost 100% utilization with negligible packet loss. Also note that the scheme allows for near immediate playback of the video upon user request. It can be argued

that for MPEG-2 traces with an order of magnitude larger average rate, the prefetch buffer will have to be an order of magnitude larger to achieve the same loss probabilities. But even with this order of magnitude increase, only 5 Mbytes of prefetch buffer is required to give negligible packet loss.

3.4 Efficient Implementation of the JSQ Prefetching Policy

In this subsection we describe a list based approach that allows for efficient implementation of the JSQ prefetch policy. We found in our experiments that it takes only 2.5 msec to execute the administrative steps necessary in each frame period for 132 video connections on a SPARCstation 2. We applied the refined stopping rule in these experiments and took also the buffer capacity test (4) into account.

The underlying data structure in our implementation is a singly linked list. Each list element is a record storing the data pertaining to an active connection; in particular, we keep track of $\theta(j)$, $p(j)$ and $b(j)$. We also keep track of an index indicating which video is being transmitted. Finally, each record contains a pointer to the next list element. The list is maintained such that the index $p(j)$ is ascending as one moves down the list, that is, the connection with the smallest number of prefetched frames is on the top. In each frame period we start by considering the connections at the top of the list. We first check whether conditions (3) and (4) (and also (5) when applicable) are satisfied. If these conditions are satisfied, we transmit the frame, increment $p(j)$ and adjust $b(j)$. Next, we check whether the successor in the list has a smaller $p(j)$. If not, we try to prefetch the next frame. This is repeated until $p(j)$ is larger than that of the next connection in the list. At that point we have to rearrange the pointer references in order to maintain the order of the list. After the order has been restored we start over, trying to prefetch for the connection on top of the list. If we find at any point that (3) or (4) (or (5) when applicable) is violated, we skip the connection and try to prefetch for the successor in the list, that is, we prefetch no longer for the connection on top of the list, but move down the list as we skip over connections.

4 Interactivity

In this section we adapt the JSQ prefetch policy to account for pauses as well as forward and backward temporal jumps. Our protocol allows these interactive actions to be performed with minimal delay. We assume that whenever a user invokes a interactive action, the client sends

a message indicating the interactive action to the server.

Suppose that the user for connection j pauses the movie. Upon receiving notification of the action, the server can simply remove connection j from consideration until it receives a resume message from the client; while the connection is in the paused state, its prefetch buffer remains at a constant level. A slightly more complex policy would be to fill the corresponding buffer with frames once all the other prefetch buffers are full or reach a prespecified level.

Suppose that the user for connection j makes a temporal jump of δ frames into the future. If $\delta < p(j)$, we discard δ frames from the head of the prefetch buffer and set $p(j) = p(j) - \delta$. If $\delta \geq p(j)$ we set $p(j) = 0$ and discard all the frames in the prefetch buffer. Finally, suppose that the user for connection j makes a backward temporal jump. In this case we set $p(j) = 0$ and discard all frames in the prefetch buffer.

In terms of frame loss, pauses actually improve performance because the movies which remain active have more bandwidth to share. Frequent temporal jumps, however, can degrade performance since prefetch buffers would be frequently set to zero. Whenever we set a prefetch buffer to zero, the pool loses some of its total savings, thereby increasing the likelihood of loss.

We now present some numerical results for interactive actions. We consider only forward and backward temporal jumps and ignore pauses as pauses can only improve performance; we furthermore assume that $\delta > p(j)$ for all forward temporal jumps. Our results give therefore a conservative estimate of the actual performance. In our simulation, we assume that each user performs temporal jumps repeatedly, with the time between two successive jumps being exponentially distributed with constant rate. When a user performs such an action, her prefetch buffer is set to zero. Figure 4 shows the 90 % confidence interval for $P_{\text{loss}}^{\text{time}}$ for 11, 22, 32 and 43 temporal jumps per hour (on average). This experiment was conducted with $B(j) = 256$ KBytes and a 95 % link utilization. As we would expect, loss probabilities increase as the rate of temporal jumps increase; however, the increase is not significant for a sensible number of temporal jumps.

5 Comparison between JSQ Prefetching and Optimal Smoothing

As mentioned in the Introduction, there are other protocols in the literature for the transport of VBR-encoded video which exploit prefetching and client buffers [14] [22] [24] [31] [8] [9]. In all of these other schemes, the transmission schedule for a given video is determined off-line




Figure 4: $P_{\text{loss}}^{\text{time}}$ as a function of the average number of temporal jumps per hour for 256 KByte buffers and 95 % link utilization.

and depends only on the traffic characteristics of that video; thus, the transmission schedule of a connection does not take into account the traffic requirements of the other connections in progress. For this reason, we refer to all the schemes cited just above as *non-collaborative* prefetching schemes.

On the other hand, the JSQ prefetching protocol determines the transmission schedule of a connection on-line, as a function of the buffer contents at *all* of the clients. For this reason, we refer to JSQ as a *collaborative prefetching scheme*. The purpose of this section is to show that the JSQ buffer fill policy is responsible for the outstanding performance reported in Section 3.3. To this end we shall compare JSQ prefetching with one of the non-collaborative schemes in the existing literature, namely, Optimal Smoothing [31] [37]. We apply the Optimal Smoothing algorithm to the traces used for the JSQ experiments (see Section 3.3). We then statistically multiplex the smoothed traces on a bufferless 45 Mbps link and calculate $P_{\text{loss}}^{\text{time}}$ using the Large Deviation (LD) approximation described in [28] and [37]. The LD approximation is known to be very accurate [1, 15, 6, 7, 28]. We do this for two versions of optimal smoothing: no initiation delay and a 10 frame initiation delay [31] [37] [5]. We chose Optimal Smoothing for our comparison, as optimal smoothing minimizes the peak rate and variability of the smoothed traffic for a given client buffer.

The results are given in Figure 5; only the point estimates of the simulation results are plotted here. The figure shows $P_{\text{loss}}^{\text{time}}$ as a function of buffer size at the client for 95% average link utilization. The figure shows that JSQ prefetching gives a $P_{\text{loss}}^{\text{time}}$ that is more than two orders of magnitude smaller than Optimal Smoothing for a 128 Kbyte buffer, and more than three orders of magnitude smaller for a 256 Kbyte buffer. Furthermore, for this case of 95% utilization, Optimal Smoothing has unacceptably high loss probabilities for all buffer sizes shown, whereas JSQ prefetching gave no loss in 4×10^8 simulated frame periods for a buffer size of 300 Kbytes. (To account for this point on the graph, we suppose that loss occurs for one of the 4×10^8 frame periods at 300 Kbytes.) We are thus led to the conclusion that the collaborative nature of JSQ prefetching contributes significantly to its high performance. JSQ prefetching is inspired by the least-loaded routing algorithm for circuit-switched loss networks, which is known to give excellent performance and to be extremely robust over a wide range of traffic conditions [30].

In Figure 6 we compare the admission region of JSQ prefetching, Optimal Smoothing and GOP smoothing. We use the lambs video, the burstiest of the mix of movies (see Table 1), for this experiment. We fix the client buffer size at 1 MByte. A client buffer of 1 MByte can




Figure 5: $P_{\text{loss}}^{\text{time}}$ as a function of buffer size for JSQ prefetching and statistical multiplexing of optimally smoothed traces for 95% average link utilization.




Figure 6: Number of lambs connections as a function of normalized bandwidth for JSQ prefetching, Optimal Smoothing and GOP smoothing. The client buffer of 1 MByte holds on average 41 seconds of the lambs video.

store on average 41 seconds of the lambs video. We set the QoS criterion to $P_{\text{loss}}^{\text{time}} \leq 10^{-5}$. The plot gives the number of lambs connections that can be supported as a function of the normalized bandwidth of the bottleneck link. We define normalized bandwidth as the link bandwidth divided by the average rate of the lambs video. The JSQ prefetching results are obtained from simulation. The Optimal Smoothing and GOP smoothing results are computed using the very accurate LD approximation. With GOP smoothing the video is smoothed over each Group of Pictures (GOP), that is, 12 frames in the case of the lambs movie (see Section 3.3). The plot shows that JSQ prefetching clearly outperforms the other smoothing schemes. With a normalized bandwidth of 21.4, corresponding to a link rate of 4 Mbps, JSQ prefetching can support 20 lambs connections (95 % average link utilization) while Optimal Smoothing can support 15 connections (70 % average link utilization). GOP smoothing can support 10 connections (48 % average link utilization).

We furthermore note that the JSQ policy allows for instantaneous interactive actions, making the scheme amenable to highly interactive multimedia applications. These interactions cause only a modest drop in performance (see Section 4). An indepth study of interactivity for Optimal Smoothing is given in [5]. In general, no smoothing scheme can always give instantaneous playback after all temporal jumps, because for some playback points a build-up delay is required to ensure no starvation in the future. In [5] an algorithm is given for which the build-up delay is typically very small for the traces considered. However, this build-up delay typically increases with the client buffer size, and at playback points before long high-bandwidth scenes the delay can be large (an example is given for which the delay after a temporal jump can be as much as 63 seconds for a 4Mbyte buffer). Thus, with Optimal Smoothing, as the client buffers become larger, the link utilizations increase but interactivity performance decreases. In contrast, with JSQ prefetching, both link utilization and interactivity performance improve with increasing client buffer.

We now provide an intuitive explanation of why packet loss is less of JSQ prefetching. The JSQ policy strives to (1) equalize the number of prefetched frames over all ongoing connections and (2) keep the buffers always filled. Optimal Smoothing, on the other hand, first fills the buffer when a connection starts up. The buffer is then emptied, then filled again and so on; the buffer content thus oscillates between empty and full.

Now consider a situation where a number of connections have been in progress for a while and a new connection becomes active. The JSQ policy dedicates the entire link capacity temporarily to the new connection so as to bring the number of prefetched frames up to the

level of the other connections. These other connections are fed from their prefetch buffers during this period. For loss to occur, the prefetch buffers of the older connections need to be drained completely and the aggregate rate of the ongoing connections has then to exceed the link capacity. This is a highly unlikely scenario given that the prefetch buffers typically hold a large number of prefetched frames. With optimal smoothing, however, the new connection gets only a small fraction of the link capacity since the other connections have to transmit according to their fixed schedule. Losses are therefore more likely, particularly for movies requiring a large amount of bandwidth in the beginning.

Next, consider a situation where a number of connections have been active for a while and high action scenes in the movies collude and thus cause temporarily an excessive demand for bandwidth. Since JSQ prefetching keeps the buffers always full while the optimal smoothing policy drives the buffer contents periodically up and down, the aggregate buffer contents of all the ongoing connections is larger with JSQ prefetching. The JSQ policy can therefore support longer periods of excessive demand for bandwidth.

We conclude this section by mentioning that the non-collaborative prefetching schemes do have an important advantage over our JSQ prefetching protocol – namely, they can be implemented over multiple, decentralized servers, and they can be easily adapted to run over a network with multiple bottleneck links. In particular, the server does not need to be attached to a cable headend or an ADSL switch, but can instead be distributed and placed deeper into the network. In our current research, we are adapting the JSQ prefetching protocol so that the server can be distributed and placed deeper into the network.

6 A Packet-Based JSQ Prefetch Policy

Up to now we have measured the length of the queue in a prefetch buffer by the number of frames present in the buffer. An alternative measure for length is the number of packets in the buffer. With this measure, the JSQ policy strives to equalize the number of packets in each of the prefetch buffers. In defining this policy, we drop the subscript l from all notations. At the beginning of each slot the server determines the j^* with the smallest $b(j) - x_{\theta(j)}(j)$, transmits one packet from connection j^* and increments $b(j^*)$. Within this same slot the server repeats this procedure over and over again, at each iteration finding a new j^* that minimizes $b(j) - x_{\theta(j)}(j)$, transmitting a packet from connection j^* and incrementing $b(j^*)$. For infinite prefetch buffers, the procedure stops when z (the total number of packets transmitted within

the slot) equals R/F . For finite prefetch buffers, the procedure stops when $z = R/F$ or when all the prefetch buffers are full.

The packet-based JSQ policy is easier to implement than the frame-based JSQ policy. For brevity, we only describe this implementation for the case of infinite prefetch buffers. At the beginning of each slot we order the prefetch buffers according to their $y[j] = b(j) - x_{\theta(j)}(j)$ values. Let $\alpha(j)$, $j = 1, \dots, J$, be this ordering. We fill buffer $\alpha(1)$ with $y[\alpha(2)] - y[\alpha(1)]$ packets. We then fill buffers $\alpha(1)$ and $\alpha(2)$ with $y[\alpha(3)] - y[\alpha(2)]$ packets. The procedure continues as long as $z \leq R/F$. If at any iteration we will have $z > R/F$, we spread the remaining packets evenly over the buffers in question.

The drawback of this policy is that it can produce frame levels in the prefetch buffers which are highly unbalanced; this can occur when one set of connections has a large number of packets per frame and a second set has a small number of packets per frame. An advantage of this policy is that it can fill all the prefetch buffers to the brim when drain rate is far below the link rate.

Throughout the remainder of this paper we use the original frame-based JSQ policy.

7 Selective Discard Policies for MPEG Compression

There is cell loss in a slot l if and only if

$$\sum_{j \in \mathcal{J}} x_{\theta_l(j)}(j) > R/F, \quad (7)$$

where \mathcal{J} is the set of connections that have no frames in their buffers at the beginning of the slot. However, when (4) holds, the server still has the freedom to select which part of the aggregate traffic,

$$\sum_{j \in \mathcal{J}} x_{\theta_l(j)}(j),$$

to transmit and which part to discard. A rule for choosing which part of the aggregate traffic to discard is a *selective discard policy*. A good selective discard policy will discard the parts of the aggregate traffic that can most easily be estimated at the client. The estimation of lost video traffic at the client is called *error concealment*, which, for MPEG encoding, has been thoroughly discussed in the literature [20, 2, 38].

7.1 Selective Discard Strategies for MPEG-1

Recall that for MPEG-1 encoding the frames have a GOP pattern, e.g., IBBPBBPBB. Because B and P frames are encoded with respect to the I frames, it is important to minimize loss for each I frame, as this loss propagates through the entire GOP. Similarly, loss for P frames is more damaging than loss for B frames. These observations lead to the following simple selective discard policy. Let $\mathcal{J}(I)$, $\mathcal{J}(P)$ and $\mathcal{J}(B)$, be the set of connections in \mathcal{J} of the current slot that are to transmit I, P and B frames. Note that the sets \mathcal{J} , $\mathcal{J}(I)$, $\mathcal{J}(P)$ and $\mathcal{J}(B)$ change from slot to slot. If

$$\sum_{j \in \mathcal{J}(I)} x_{\theta_I(j)}(j) + \sum_{j \in \mathcal{J}(P)} x_{\theta_I(j)}(j) \leq R/F, \quad (8)$$

then we transmit all the frames in $\mathcal{J}(I) \cup \mathcal{J}(P)$ and discard “part” of the traffic from $\mathcal{J}(B)$. If (8) fails, but

$$\sum_{j \in \mathcal{J}(I)} x_{\theta_I(j)}(j) \leq R/F, \quad (9)$$

then we transmit all the frames in $\mathcal{J}(I)$, discard all the frames in $\mathcal{J}(B)$ and discard “part” of the traffic from $\mathcal{J}(P)$. Finally, if (9) fails, we discard all of the frames in $\mathcal{J}(B) \cup \mathcal{J}(P)$ and “part” of the traffic in $\mathcal{J}(I)$. The above discard policy has not been fully defined, as we have not yet specified the “part” of traffic that is discarded in each of the cases. There are several natural strategies here:

1. Slice discarding strategy: Suppose it suffices to discard traffic only from $\mathcal{J}(B)$. Then we discard slices from the B frames, spreading the discarded frames over the connections in $\mathcal{J}(B)$; furthermore, if we must discard multiple slices in a frame, we attempt to discard non-adjacent slices in order to facilitate estimation and concealment at the client. If discarding all the B frames in $\mathcal{J}(B)$ does not suffice, then we also discard slices from P frames; finally, if discarding all the B and P frames from $\mathcal{J}(B) \cup \mathcal{J}(P)$ does not suffice, then we also discard slices from I frames.
2. Frame discarding strategy: Discard entire frames in a round robin fashion across all the videos in \mathcal{J} . With this strategy, loss of consecutive frames of a specific connection should be rare.
3. Do not discard the frames or any “parts” of the frames, but instead to re-encode the frames on-the-fly with a coarser quantization. A related strategy is to include in the server’s storage a second version of each video which has been off-line encoded at a lower




Figure 7: The fraction of lost P and B frames as a function of the prefetch buffer size for 95 % utilization.

rate; when \mathcal{J} is nonempty and loss is unavoidable, then the server switches to the low-rate versions for all videos in \mathcal{J} .

Figure 7 illustrates the performance of this discard policy. In the simulation we simplify the policy by never transmitting any “parts” of the traffic; specifically, we drop all frames from $\mathcal{J}(B)$ when (8) holds, all frames from $\mathcal{J}(B) \cup \mathcal{J}(P)$ when (8) fails and (9) holds; and all frames from \mathcal{J} when (9) fails. We conducted this experiment with a link utilization of 95% and ran 1000 replications for buffer sizes ranging from 4 to 64 KByte. The figure plots the point estimates of the fraction of P and B frames that are lost (we ignore confidence intervals to avoid visual clutter). In all our simulations we never observed any loss of I frames.

7.2 An Admission Policy for MPEG-1

We now present a connection admission policy which, when combined with our selective discard policy, ensures that I frames experience practically no loss. To this end, observe that with our selective discard policy the loss probability for I frames will be less than that when the clients have no buffer and the traffic in the P and B frames is equal to zero. For this conservative model, we can adapt the theory in [28] to construct a conservative large-deviation approximation for the loss of I frames. Using this approximation, we only admit a new connection if (1) the fraction of frame periods which have I-frame loss remains less than some minute number, say, 10^{-9} , and (2) the link utilization remains below, say, 95 %. This admission policy essentially guarantees no loss of I frames while simultaneously ensuring a low probability of loss for the B and P frames.

We give now a brief outline of how the large-deviation approximation applies to our scheme. First, for each video in the server we construct a modified trace $y_n(j)$, $j = 1, \dots, J$, with $y_n(j) = x_n(j)$ if the n th frame is an I frame and $y_n(j) = 0$ if the n th frame is either a P or B frame. We then calculate the frame size distribution for the modified trace:

$$\pi_j(l) := \frac{1}{N} \sum_{n=1}^N 1(y_n(j) = l),$$

The logarithmic moment generating function of the aggregate amount of traffic in a frame for the modified trace is:

$$\mu_Y(s) := \ln E[e^{sY}],$$

where

$$Y = \sum_{j=1}^J Y(j)$$

and $Y(j)$ is distributed according to $\pi_j(\cdot)$. The probability that at least one I frame is lost in an arbitrary frame period is approximated as

$$\begin{aligned} P_{\text{loss}}^{\text{time}}(I) &= P(\text{"all I frames cannot be transmitted in a frame period"}) \quad (10) \\ &\approx \frac{1}{s^* \sqrt{2\pi \mu_Y''(s^*)}} e^{-s^* R/F + \mu_Y(s^*)}, \end{aligned}$$

where s^* satisfies

$$\mu_Y'(s^*) = R/F.$$




Figure 8: $P_{\text{loss}}^{\text{time}}(I)$ as a function of $P_{\text{loss}}^{\text{time}}$ for 16 KByte prefetch buffer.

The admission control test is thus satisfied if

$$\frac{1}{s^* \sqrt{2\pi\mu_Y''(s^*)}} e^{-s^*R/F + \mu_Y(s^*)} \leq \epsilon,$$

where ϵ is some minute number such as 10^{-9} .

Figure 8 shows the relationship between the $P_{\text{loss}}^{\text{time}}$ (the fraction of frame periods during which loss occurs for I,P, or B frames) and $P_{\text{loss}}^{\text{time}}(I)$ (the conservative estimate for the fraction of frame periods during which loss of an I frame occurs when using selective discard). In this figure we have used a 16 KByte prefetch buffer at each of the clients. The series of points was obtained by varying the link utilization. (For visual simplicity we ignore the confidence intervals for $P_{\text{loss}}^{\text{time}}$ and only report the mid-point in the figure.) From this curve we see, for example, that when $P_{\text{loss}}^{\text{time}}$ is 2×10^{-5} then the fraction of frame periods during which I frames are lost is 3×10^{-8} . The values of $P_{\text{loss}}^{\text{time}}(I)$ reported in this figure come from the large-deviation approximation. This approximation is conservative, as it assumes that the prefetch buffers are

zero; the actual loss probability for I frames is much less. In fact, for the last point (as well as the other points) on the graph, corresponding to $P_{\text{loss}}^{\text{time}} \approx 0.1$ and 99.5% utilization, we observed absolutely no loss of I frames in the simulation!

7.3 MPEG-2

The MPEG-2 standard is similar to MPEG-1 but includes extensions and refinements to cover a wider range of applications. While MPEG-1 was primarily introduced for CD-ROM and has a typical bitrate of 1.5 Mbps, MPEG-2 is intended for the digital transmission of broadcast TV quality video with bitrates between 4 and 9 Mbps.

The most important extension of MPEG-2 over MPEG-1 for our discussion is scalability. Scalability offers the possibility to partition the video stream into two different layers: the base and enhancement layers. This allows for transmission with different priorities. The enhancement layer conveys refined image information at a lower priority. Transmitting a basic quality image at the base layer with high priority ensures that a minimum quality image can be decoded at the client. If both layers are received, a high quality image can be displayed. If the low priority enhancement layer is lost, the client will at least be able to display a basic quality image (provided the base layer is received intact).

Although there are four scalability modes we shall focus on SNR scalability as it offers a flexible way to partition the video data while keeping the encoder and decoder architecture simple. The basic idea of SNR scalability is to first generate the base layer by quantizing the DCT coefficients of a frame with a coarse base quantization matrix. The difference between the original DCT coefficients and the DCT coefficients reconstructed from the base layer is quantized with a finer quantization matrix and transmitted in the enhancement layer. If the low priority enhancement layer is lost, the frame is decoded using only the base layer.

An admission control policy for MPEG-2 video may proceed similar to the one outlined for MPEG-1 video in subsection 7.2. Again, we conduct 2 tests:

1. The admission control test for a bufferless link outlined under 7.2 with a very small QoS parameter ϵ such as 10^{-9} . Consider only the base layer in this test.
2. Taking base and enhancement layer into account, check whether the average link utilization is less than 0.95.

The first test practically guarantees the delivery of the base-layer image.

8 Video Delivery to the Home

8.1 ADSL

In this section we discuss how JSQ prefetching ties into the Asymmetric Digital Subscriber Line (ADSL) technology. First, we give a brief overview of the ADSL technology; for more details see [11, 21]. ADSL exploits advances in digital signal processing to provide a high speed downstream channel from the central office to the home (up to 9 Mbps total capacity) and a lower rate upstream channel (up to 640 kbps total capacity) over a single twisted pair copper loop while leaving the POTS (Plain Old Telephone Service) unaffected. The high speed downstream channel provides a switched point-to-point link for delivering video from the central office to the home. The lower rate upstream channel can be used to control the downstream video stream; in particular, the upstream channel can be used to convey the request for a video and VCR actions to the video server. The downstream channel capacity provided by the ADSL technology, however, depends on the length of the copper loop, wire gauge and interference. The maximum downstream rate of 9 Mbps can be supported over approximately 3,000 meter of regular copper wire. As the wire length increases to 6,000 meter the available downstream capacity drops to 1.5 Mbps due to increased attenuation. The bandwidth requirements for VBR transmission of video can readily be accommodated by ADSL for homes within 3,000 meter of the central office since the peakrate of VBR video is typically well below 9 Mbps.

Figure 9 shows the architecture of a VoD system utilizing ADSL. The video server multiplexes the video streams onto an essentially bufferless link. The videos may be transported in ATM cells or TCP/IP packets. The packet switch in the central office forwards the cells/packets to the appropriate output buffer. The cells/packets are then transmitted to the individual homes over the point-to-point ADSL operated copper lines. The ADSL Forum is developing interface and protocol guidelines for three basic distribution protocols [10] [33]:

1. The ADSL modem is configured for continuous bit-rate synchronous transmission; thus, the ADSL modem works as a bitpump.
2. The modem is configured for packet transport. The ADSL modem is configured for Ethernet packets in this case.
3. The ADSL modem is configured for an ATM interface.

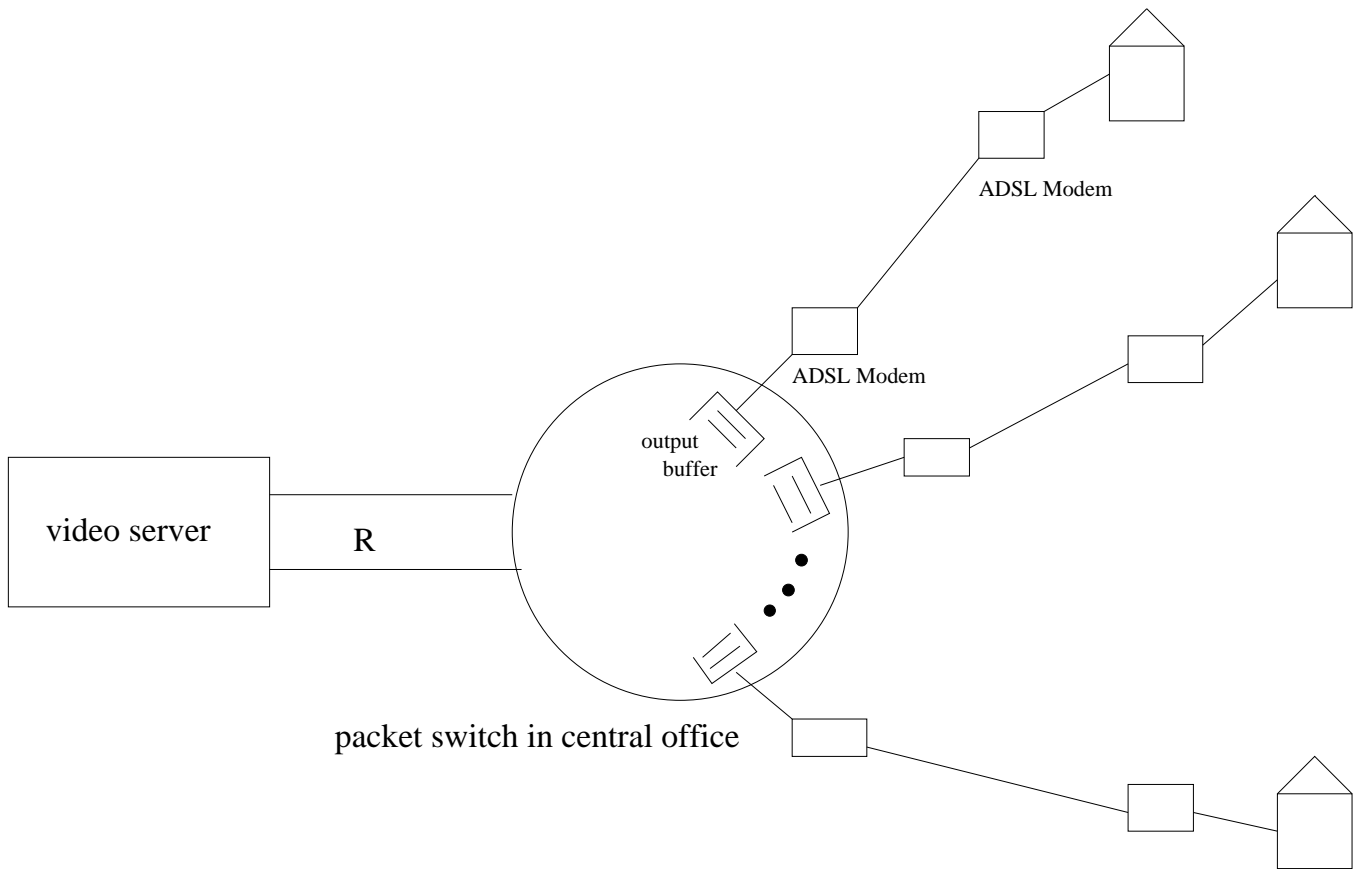


Figure 9: VoD architecture with ADSL

In this paper we are focusing on packet or ATM cell based transmission. Hence, the following transport modes are possible on the link joining the video server and the central office and on the ADSL link joining the central office and the home:

1. ATM cells from server through switch to home
2. IP datagrams from server through switch. After switch datagrams are encapsulated in Ethernet frames.
3. ATM cells from server through switch. After switch, cells are converted to IP packets.

This last option has the advantage of using powerful ATM switching technology in the central office while allowing the user to have inexpensive Ethernet equipment.

8.2 Cable

In this section we discuss how the JSQ prefetching protocol for VoD proposed in this paper ties into the cable modem technology. First, we give a brief overview of the cable modem technology; for more details see [17, 18, 25]. The coaxial cable was installed for broadcast of one-way analog TV. Cable is a shared medium; many homes are attached to the same coaxial cable. Medium access control for the downstream video traffic is particularly simple as there is only one sender, the headend. The upstream control traffic, however, poses a problem. Carrier sensing fails for cable plants with tree-and-branch structure, where only the headend hears every source. Remedies for this problem are currently being developed [17]. As of the writing of this paper, there are no fixed standards that specify how upstream and downstream bandwidth are allocated to homes. Typically, the upstream traffic is transmitted in the 5–40 MHz range. The downstream bandwidth from 40 – 750 MHz is split into 6 MHz channels for analog TV. Each of these channels yields approximately 25 Mbps when 64 Quadrature Amplitude Modulation (QAM) is employed and could thus carry a couple of video streams.

Figure 10 shows a possible VoD architecture with cable. The video server is attached directly to the cable headend as are multiple cable trunks. We assume in this architecture that the link connecting the video server to the headend has infinite bandwidth. The bottleneck link is the cable trunk connecting the homes to the headend. Homes are attached to cable trunks via cable modems. The request for a video is relayed from the viewers home to the headend via the upstream channels. The headend forwards the request to the video server. The video server immediately starts transmitting the video frames. The video server runs the JSQ algorithm

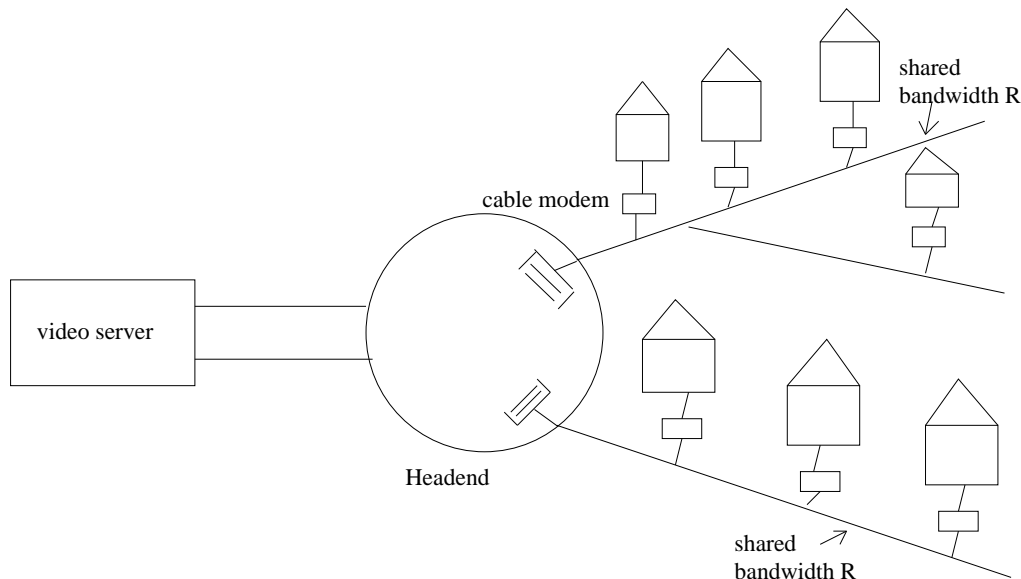


Figure 10: VoD architecture for cable residential access

for each cable trunk. The switch in the headend forwards the frames to the appropriate output queue. All the videos requested by viewers connected to the same cable trunk are multiplexed onto the shared channel of capacity, say R bps. Our JSQ prefetching protocol allows for the efficient use of the valuable trunk bandwidth, R . We achieve transmission with negligible losses and thus constant high video quality for average trunk bandwidth utilizations of 95%.

9 Conclusion

Prerecorded video has two special properties: (1) for each video, the traffic in each video frame is known before the video session begins; (2) while the video is being played, some of the video can be prefetched into the client memory. In this paper we have shown how these two properties can be exploited to achieve high performance when there is one shared link between the server and the clients. We have also shown how selective discard can enhance the performance when all of the video connections originate from the same server. The results should be useful for designing VoD systems that connect servers to residential broadband networks using cable or ADSL, or for VoD systems that connect the server to its clients through a LAN. Our client-server scheme can be part of a larger Internet solution to VoD, whereby the prerecorded videos are multicast to local servers at off-peak hours with the best effort service. We are currently working on a variation of the JSQ prefetch protocol so that the server can be distributed and

placed deeper into the network.

References

- [1] Methods for the performance evaluation and design of broadband multiservice networks. In *The COST 242 Final Report*, 1996.
- [2] L. T. Chia., D. J. Parish, and JWR Griffiths. On the treatment of video cell loss in the transmission of motion-JPEG and JPEG images. *Computers & Graphics*, 18(1):11–19, January 1994.
- [3] I. Dalgic and F. A. Tobagi. Characterization of quality and traffic for various video encoding schemes and various encoder control schemes. Technical Report CSL–TR–96–701, Stanford University, Departments of Electrical Engineering and Computer Science, August 1996.
- [4] S. E. Deering. *Multicast Routing in Datagram Internetwork*. PhD thesis, Stanford University, December 1991.
- [5] J. Dey, S. Sen, J. Kurose, D. Towsley, and J. Salehi. Playback restart in interactive streaming video applications. In *To appear in Proceedings of IEEE Multimedia*, Ottawa, Canada, 1997.
- [6] A. Elwalid, D. Heyman, T. Lakshman, D. Mitra, and A. Weiss. Fundamental bounds and approximations for ATM multiplexers with application to video teleconferencing. *IEEE Journal on Selected Areas in Communications*, 13(6):1004–1016, August 1995.
- [7] A. Elwalid, D. Mitra, and R. H. Wentworth. A new approach for allocating buffers and bandwidth to heterogeneous regulated traffic in an ATM node. *IEEE Journal on Selected Areas in Communications*, 13(6):1115–1127, August 1995.
- [8] W. Feng, F. Jahanian, and S. Sechrest. Providing VCR functionality in a constant quality video-on-demand transportation service. In *IEEE Multimedia*, Hiroshima, Japan, June 1996.
- [9] W. Feng and J. Rexford. A comparison of bandwidth smoothing techniques for the transmission of prerecorded compressed video. In *Proceedings of IEEE Infocom*, Kobe, Japan, April 1997.
- [10] ADSL Forum. ADSL forum system reference model. Technical Report TR–001. http://www.adsl.com/adsl_reference_model.html.
- [11] ADSL Forum. *ADSL Tutorial: Twisted Pair Access to the Information Highway*. http://www.adsl.com/adsl_tut.html.
- [12] D.J. Gemmell and J. Han. Multimedia network file servers: multichannel delay sensitive data retrieval. *Multimedia Systems*, 1:240–252, 1994.
- [13] D.J. Gemmell and H.M. Vin. Multimedia storage servers: A tutorial. *IEEE Computer*, pages 40–49, 1995.

- [14] M. Grossglauser, S. Keshav, and D. Tse. RCBR: A simple and efficient service for multiple time-scale traffic. In *ACM SIGCOMM*, 1995.
- [15] I. Hsu and J. Walrand. Admission control for ATM networks. In *IMA Workshop on Stochastic Networks*, Minneapolis, Minnesota, March 1994.
- [16] E. W. Knightly and H. Zhang. Traffic characterization and switch utilization using a deterministic bounding interval dependent traffic model. In *Proceedings of IEEE Infocom '95*, Boston, MA, April 1995.
- [17] M. Laubach. Cable modem basics: An introduction to cable modem technology, October 1996. Presentation at the University of Pennsylvania
Slides available at <ftp.com21.com/pub/laubach/>.
- [18] M. Laubach. To foster residential area broadband internet technology: IP datagrams keep going and going and going. . . . *ConneXions*, 10(2), February 1996.
- [19] J. Liebeherr and D. Wrege. Video characterization for multimedia networks with a deterministic service. In *Proceedings of IEEE Infocom '96*, San Francisco, CA, March 1996.
- [20] W. Luo and M. El Zarki. Analysis of error concealment schemes for MPEG-2 video transmission over ATM based networks. In *Proceedings of SPIE Visual Communications and Image Processing 1995*, Taiwan, May 1995.
- [21] K. Maxwell. Asymmetric digital subscriber line: Interim technology for the next forty years. *IEEE Communications Magazine*, pages 100–106, October 1996.
- [22] J. M. McManus and K. W. Ross. Video on demand over ATM: Constant-rate transmission and transport. *IEEE JSAC*, 14(6):1087–1098, August 1996.
- [23] J.M. McManus and K.W. Ross. A comparison of traffic management schemes for prerecorded video with constant quality service. Available at <http://www.seas.upenn.edu/~ross>.
- [24] J.M. McManus and K.W. Ross. Prerecorded VBR sources in ATM networks: Piecewise-constant rate transmission and transport. In *Proceedings of SPIE*, Dallas, TX, October 1997. Available at <http://www.seas.upenn.edu/~ross/>.
- [25] P. Mockapetris. @HOME network overview, October 1996. Presentation at the University of Pennsylvania.
- [26] A.L.N. Reddy and J. Wyllie. Disk scheduling in multimedia I/O system. In *Proceedings of ACM Multimedia Conference*, 1992.
- [27] M.I. Reiman. Some diffusion approximations with state space collapse. In F. Baccelli and G. Fayolle, editors, *Lecture Notes in Control and Informational Sciences 60*, pages 209–240. Springer-Verlag, 1983.
- [28] M. Reisslein and K. W. Ross. Call admission for prerecorded sources with packet loss. *IEEE Journal on Selected Areas in Communications*, 15(6):1167–1180, August 1997.

- [29] O. Rose. Statistical properties of MPEG video traffic and their impact on traffic modelling in ATM systems. Technical Report 101, University of Wuerzburg, Insitute of Computer Science, Am Hubland, 97074 Wuerzburg, Germany, February 1995.
ftp address and directory of the used video traces:
ftp-info3.informatik.uni-wuerzburg.de /pub/MPEG/.
- [30] K. W. Ross. *Multiservice Loss Models for Broadband Telecommunication Networks*. Springer Verlag, 1995.
- [31] J. Salehi, Z.-L. Zhang, Kurose J, and D. Towsley. Supporting stored video: Reducing rate variability and end-to-end resource requirements through optimal smoothing. In *Proceedings of ACM SIGMETRICS*, May 1996. Philadelphia, PA.
- [32] S. Shenker and C. Partridge. Specification of guaranteed quality of service. Technical report. Internet Draft; December 1995.
- [33] D. Veeneman and R. Olshansky. ADSL for video and data services. In *IEEE International Conference on Communications*, June 1995. <http://www.gte.com/Adsl/News/Docs/9506.html>.
- [34] H. Vin and P. Venkat Rangan. Chapter 4: Multimedia storage systems. In *Multimedia Systems and Techniques*, ed. B. Furht, 1996.
- [35] D. Wrege, E. Knightly, H. Zhang, and J. Liebeherr. Deterministic delay bounds for VBR video in packet-switching networks: Fundamental limits and tradeoffs. *IEEE/ACM Transactions on Networking*, 4(3):352–362, June 1996.
- [36] P. Yu, M.S. Chen, and D.D. Kandlur. Design and analysis of a grouped sweeping scheme for multimedia storage management. In *Third International Workshop on Network and Operating System Support for Digital Audio and Video*, 1993.
- [37] Z. Zhang, J. Kurose, J. Salehi, and D. Towsley. Smoothing, statistical multiplexing and call admission control for stored video. *IEEE Journal on Selected Areas in Communications*, 13(6):1148–1166, August 1997.
- [38] Q. Zhu, Y. Wang, and L. Shaw. Coding and cell-loss recovery in DCT-Based packet video. *IEEE Transactions on Circuits and Systems for Video Technology*, 3(3):248–258, June 1993.