Model-Based Delay-Distortion Optimization for Video Streaming Using Packet Interleaving

Yi J. Liang∗†, *John G. Apostolopoulos*◇ *and Bernd Girod*†

◇Streaming Media Systems Group
Hewlett-Packard Labs, Palo Alto, CA 94304

†Information Systems Laboratory
Stanford University, Stanford, CA 94305

Invited Paper

ABSTRACT

Bursty channel losses have been shown to generally produce larger total mean-square error distortion in streaming video than an equivalent number of isolated losses. This paper proposes a simple packet interleaving scheme to combat the effect of bursty losses by dispersing them. The optimal interleaver for minimizing the expected total distortion of the decoded video, subject to a delay constraint, is determined using a model that accurately predicts the expected distortion for different packet loss patterns. Compared to other forms of error-resilience, packet interleaving has the advantages of (1) simplicity and (2) not requiring any extra bitrate. For a simple burst loss channel, where each loss event has 100 ms duration, packet interleaving can provide between .24 and .87 dB gain over conventional non-interleaved streaming, at the expense of a delay increase of 400 ms.

1: Introduction

The unreliable and error-prone nature of the wireless channel is one of the major challenges for wireless video streaming. Applications, such as video streaming to handheld devices with the emerging Third Generation (3G) cellular system, have to cope with this lack of QoS guarantees, including packet loss. Furthermore, wireless channels are afflicted by time-varying fading and interference conditions, which may lead to bursty packet losses [1].

In [2], we find that the pattern of packet loss, including the length of burst loss, has a significant effect on video quality. Specifically, the total distortion produced by a packet loss pattern is *not* proportional to the average packet loss rate, as is often assumed. For example, the total distortion produced by a burst loss is much higher than that by an equal number of isolated losses. The quality degradation increases as the burst length increases. Burst losses are more difficult to conceal due to the amount of information lost.

The problem of error-resilient video communication has received significant attention in recent years, and a variety of techniques have been proposed to combat channel losses and increase the robustness of communication [3]. Examples of recent work in this area include, inter/intra-mode switching [4], [5], [6], [7], the use of Forward Error Correction (FEC) [8], dynamic control of prediction dependency using long-term memory [9], [10], [11], channel-adaptive packet scheduling [12], [13], [14], [15], and the use of multiple description coding and path diversity [16], [17], [18]. These techniques manage and optimize the dependency across predictively coded packets, or alleviate the negative effect of temporal correlation of the channel. All of these schemes improve error-resilience at the expense of increased bitrate.

In this work, we explore a simple packet scheduling scheme, packet interleaving, to convert burst losses into an equivalent number of isolated losses which in general are easier to recover from and produce lower total distortion. Compared to other types of error-resilience techniques, packet interleaving provides the advantages of (1) being simple, and (2) it does not require any increase in bitrate. A particularly appealing scenario is wireless multicast or broadcast, where the channel is afflicted by fading and interference which may lead to burst losses, and these one-to-many applications prohibit the use of retransmission of lost packets. Furthermore, packet interleaving may be used in conjunction with other error-resilient techniques. A potential drawback of packet interleaving is that it requires additional delay. However, the required delay, which depends on the channel burst length characteristics, generally can be relatively short.

It is beneficial to compare this work versus prior work. Interleavers have been used for decades to improve FEC performance in *channel coding* by converting burst losses into isolated (random) losses spread over many FEC codewords (instead of the burst loss afflicting a single codeword) which are then easier to correct using a random-error correcting code. However, in this work we identify that from a *error-resilient source coding perspective* we can reduce the total distortion for video streaming by converting burst losses into isolated losses. Specifically, in our case the total number of losses that afflict the video decoder is the same with or without interleaving, however we are changing the *pattern of the losses* that afflicts the video decoder.

This paper continues in Section 2 by reviewing the model proposed in [2] for accurately estimating the expected total distortion produced by each packet loss pattern. Section 3 presents the proposed packet interleaving scheme, and describes how to determine the optimal interleaver given a delay constraint and knowledge of the channel burst length behavior. Experimental results demonstrating the potential gains are presented in Section 4.

2: Modeling the Effect of Packet Losses

Many of the works on packet scheduling, such as [12], [13], [15], use RD optimization techniques to improve the performance of

∗This work was performed during a summer internship at HP Labs.

video communication over lossy channels. The goal of these optimization algorithms is to minimize the expected distortion subject to a bitrate constraint. The performance of these RD optimizations crucially depend on an accurate estimate of the distortion that results for different packet loss events. In this work, the proposed packet interleaving scheme also depends on an accurate estimate of the distortion to perform the optimization.

In [2], a model was proposed which accurately estimates the expected mean-squared error distortion for different packet loss patterns. To estimate the distortion, the model explicitly considers the effect of different loss patterns, including burst losses and separated (non-consecutive) losses spaced apart by a lag, and accounts for inter-frame error propagation and the correlation between error frames.

In the model, and also in this paper, we assume that each predictively coded frame (P-frame) is coded into a single packet, so that the loss of a packet corresponds to the loss of an entire frame. The results can also be extended to the case when each frame is coded into multiple packets where the loss of one packet does not result in the loss of an entire frame. A simple loss concealment scheme is assumed where the lost frame is replaced by the previous frame at the decoder output.

This model differs from prior models in that prior models generally assumed that the total distortion is proportional to the number of lost packets, and therefore prior models did not consider the effect of the specific packet loss pattern. However, it was shown that the total distortion produced by a packet loss pattern is not proportional to the number of lost packets. For example, the total distortion produced by a burst loss is generally much higher than that produced by an equal number of isolated losses. For a burst loss of length 2, it was shown in [2] that the resulting total distortion is not equal to the sum of two independent losses occurring at the same locations, but it also includes a cross-correlation term between the two error frames. Due to the correlation term, which is positive in most cases, the distortion resulting from the burst loss is generally greater than the sum of the distortions of two single and independent losses. The distortion as a function of burst length is illustrated in Figure 1, which plots (1) the actual measured distortion, (2) an estimate of the distortion assuming that the total distortion is proportional to the number of lost packets (this linear relationship corresponds to a straight line in the log-log plot), and (3) the estimate of the distortion using the model proposed in [2]. There are two key observations from this figure. First, the actual distortion produced by a burst loss is generally greater than an equal number of isolated losses, and hence it is not proportional to the packet loss rate. This difference is approximately 1.5 dB for a burst loss of length two, and increases with burst length. Second, the model of [2] is able to accurately estimate the total distortion for different loss events, e.g. within $\pm.25$ dB for a burst loss of length two.

A loss model for two losses separated by a small lag is also derived in [2]. In that case the total distortion is determined by two independent components, plus a correlation term between the propagated first error and the second error. The total distortion is a function of the lag between the two losses. With the models above, the distortion for more general loss patterns can be obtained by using those models concatenated and combined.

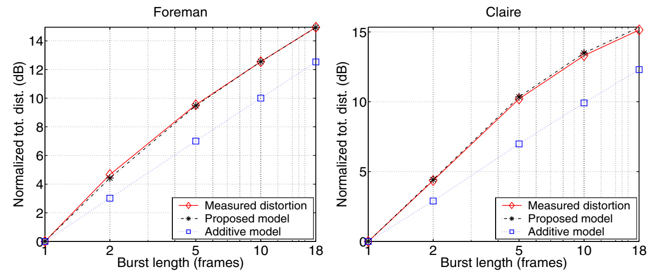

Fig. 1. Measured versus estimated total distortion using two models [2], as a function of burst loss length, normalized by total distortion for a single loss. The total distortion is the sum of MSEs of all frames affected by channel errors in the recovery period, not including quantization error.

3: Delay-Distortion Optimized Packet Interleaving

The previous section described that a burst loss generally produces greater total distortion than an equivalent number of isolated losses. This suggests that when communicating over a channel that exhibits burst losses, it would be beneficial to use interleaving to convert the burst losses into an equal number of isolated losses. In this section we use the loss model from [2] to design the optimal packet interleaving scheme that maximizes the performance (minimizes total distortion) given knowledge of the burst loss characteristics of the channel. Since there exist many approaches for interleaving, we begin by introducing the specific packet interleaving strategy used in this paper.

3.1: Packet Interleaver

A simple block interleaver is used at the sender to interleave the packets before transmission. Packets are first read into the interleaver in rows, with each row corresponding to a block of n packets. Packets are transmitted as soon as d rows of packets fill up, and are transmitted by columns. Here n is referred to as the *block size* and d is the *interleaving depth* of the interleaver. Fig. 2 shows an (n, d) interleaver. For example, consider the case when the transmission channel is afflicted by a burst loss of length three. If no packet re-scheduling (interleaving) is performed then packets 1, 2 and 3 would be lost. Using the interleaver shown in Fig. 2, packets 1, 5 and 8 would be lost if afflicted by the same loss event. Since the same burst loss now affects separated packets instead of successive packets, according to the loss model discussed in Section 2 we expect the resulting distortion to be lower. To illustrate this effect, Fig. 3 plots the PSNRs of the *Claire* sequence transmitted over a bursty channel, with and without using a $(9, 3)$ interleaver, when the same packet loss realization affects both transmissions. The locations of the lost frames that result for using and not using interleaving, for the same packet loss pattern, is also illustrated in the figure. It is observed that the simple interleaver leads to lower total distortion by converting the burst losses into isolated losses, and the quality of the reconstructed video is more uniform and also recovers faster from packet losses. Note that the total number of losses is the same in both cases, the difference is the pattern of the losses.

A simple packet interleaver permutes the locations of losses in

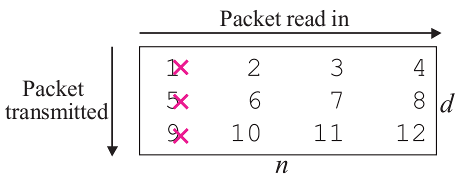

Fig. 2. A block interleaver with block size $n = 4$ and interleaving depth $d = 3$.

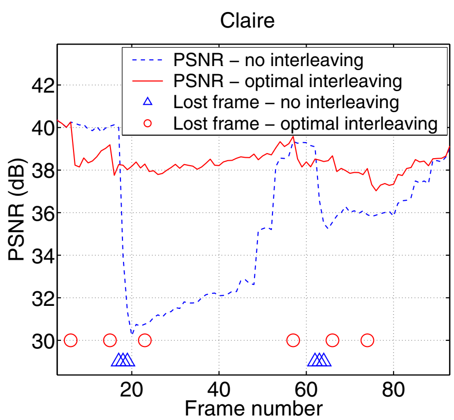

Fig. 3. PSNRs of *Claire* sequence transmitted, with and without a $(9, 3)$ interleaver, over the channel with the same loss realization. Slices are intra updated periodically to facilitate loss recovery.

order to convert burst losses into isolated losses. The effectiveness of the interleaver depends on the block size and the interleaving depth of the interleaver, and the loss characteristics of the channel. With an interleaving depth of d, a burst loss of length B can be converted into a shorter burst with a maximum length of $\lceil B/d \rceil$, where $\lceil x \rceil$ denotes the smallest integer not smaller than x. In an ideal case, when $d \geq B$, the burst loss is converted into isolated losses. In this case, the separation between any two losses is either n or $n - 1$.

A larger interleaver is more effective in that it can convert a longer burst loss into isolated losses or increase the separation of the converted isolated losses. However, this is at the cost of higher latency. At the client, an interleaved packet received cannot be used until all the packets it depends on are received. For a (n, d) interleaver, the n-th packet in the original order suffers from the highest delay, which has to be transmitted in the $((n-1) \cdot d + 1)$-th place. Hence, the decoding delay corresponding to a (n, d) interleaver is $(n-1) \times (d-1)$, and a trade-off exists between the effectiveness in permuting the packets and the latency. It should be noted that generally a large delay is not required since, as shown in Section 4, as $n \times d$ increases beyond a certain point, further increase in n or d does not necessarily improve the performance, i.e. a larger interleaver is not always better. Also note that the total delay here is not the typical $n \times d$ which arises in channel coding situations, since we do not have the delay of applying FEC across the entire interleaved data. In the next subsection, we determine the optimal interleaver (n, d) under certain delay constraints.

3.2: Optimal Interleaving

We use the set $K_{orig} = \{k_1, k_2, ...\}$ to denote the indices of the original lost packets when transmitted over the channel with no interleaving. With interleaving, the losses are redistributed across packets, and the loss indices are a function of the interleaver parameters. We use $\mathcal{K} = I(n, d, K_{orig})$ to denote the indices of the lost packets when a (n, d) interleaver is used, where $I(\cdot)$ is the functional representation of the interleaver (n, d), and K_{orig} denotes the indices of the lost packets before interleaving. In the example in Fig. 3, where two bursts of losses occur, $K_{orig} = \{17, 18, 19, 62, 63, 64\}$, while $\mathcal{K} = I(9, 3, K_{orig}) = \{6, 15, 23, 57, 66, 74\}$.

The total distortion D of the decoded video sequence, which depends on the loss pattern, is a function of the lost packets $\mathcal{K}$, and hence a function of the interleaver used, (n, d). If the channel loss statistics are known (for instance, the distribution of B is known) we are able to design the optimal interleaver (n_{opt}, d_{opt}) that achieves the lowest distortion given a delay constraint. The problem is formally stated as follows: given the channel loss characteristics, and the delay constraint C_{delay}, determine the optimal interleaver (n_{opt}, d_{opt}), such that the total distortion of the decoded video sequence $D[I(n, d, K_{orig})]$ is minimized, i.e.,

$$(n_{opt}, d_{opt}) = \arg\min_{n,\, d:\, (n-1)\times(d-1) \leq C_{delay}} D[I(n, d, K_{orig})].$$

This is a delay-distortion optimization problem. To solve for the optimal n and d, we need to estimate the distortion that results for different loss patterns $\mathcal{K}$. This is achieved using our proposed loss model discussed in Section 2.

When the characteristics of a channel are known, e.g., the probability distribution for burst loss length B, the distortion used above for optimal interleaver selection is the expected distortion. The optimal interleaver is selected to minimize the expected distortion.

3.3: Algorithm

Given an estimate of the channel loss characteristics, we can estimate the probability of different loss patterns and hence the associated loss events K_{orig}. For a given delay constraint C_{delay}, we determine all factorizations of n and d, such that $(n-1) \times (d-1) \leq C_{delay}$. For each set of interleaver parameters (n, d), we calculate the indices $I(n, d, K_{orig})$ of the redistributed losses. For a particular loss event $\mathcal{K} = I(n, d, K_{orig})$, we are able to estimate the corresponding total distortion, $D[\mathcal{K}]$, using the loss model discussed in Section 2. The estimated distortion for a particular loss event $\mathcal{K}$, and for a particular video sequence, can also be stored at the sender or streaming server for future use.

In Fig. 4, we present an example of determining the optimal interleaver (n, d) for a delay constraint of 13 frames, and given a simple channel loss model where each loss event corresponds to a burst loss of length 3 packets (frames). The detailed experimental conditions are described in Section 4. There are 34 eligible interleavers for this delay, including $(14, 2)$, which has a delay of 13, and $(2, 13)$, $(3, 7)$, $(4, 5)$, $(5, 4)$, $(7, 3)$, $(13, 2)$, which all have delays of 12, and many other factorizations (not shown in example) which have lower delays. For each of these eligible interleavers (n, d), the total distortion is calculated using the loss model, and averaged result over multiple loss realizations.

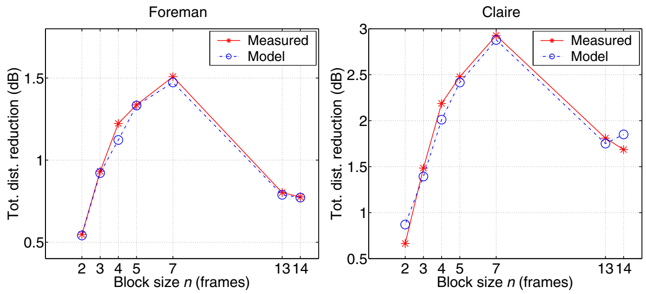

Fig. 4. Determining the optimal interleaver given a delay constraint of 13 frames. Candidate interleavers include $(2, 13)$, $(3, 7)$, $(4, 5)$, $(5, 4)$, $(7, 3)$, $(13, 2)$, and $(14, 2)$, with only the block size n labeled on the horizontal axis. The distortion reduction is calculated based on the conventional case when no interleaving is used.

The reduction in total distortion by using different interleavers (only the best 7 out of 34 are shown) is plotted in Fig. 4 for the $Foreman$ and $Claire$ sequences. Results shown are given both by actual measurements and model estimation. Note that the model predicts that the $(7, 3)$ interleaver provides the maximum improvement (minimum distortion), which agrees with the measured results. Intuitively, since the burst loss length is 3, it is sufficient to have the interleaving depth $d = 3$ to convert the burst loss into separated losses, and a large block size n provides the advantage of further separating the losses. However as n further increases, d has to drop to satisfy the delay constraint. When d drops to below 3, the performance degrades since a burst loss can no longer be converted into isolated losses. It is also important to note that an interleaver with higher delay does not necessarily give better performance. In this example, the $(14, 2)$ interleaver, with a delay of 13, does not outperform some of the other interleavers with a delay of 12, due to the limited factorizations and hence limited choices of (n, d). Also note that the *same interleaver* was identified as optimal by both the model and the measurements for both the $Foreman$ and $Claire$ sequences. This demonstrates the high accuracy of the model-based distortion estimation and its effectiveness in selecting the optimal interleaver to minimize the total distortion. This example also suggests that while the optimal interleaver depends on the channel burst loss characteristics, it may not strongly depend on the specific video sequence to be transmitted.

4: Experimental Results

This section presents experimental results to illustrate the potential performance gain that may be achieved by using the proposed simple interleaving scheme for a channel that exhibits a significant amount of burst loss. In addition, we investigate the trade-off between performance gain from larger interleavers and the corresponding delay. We use a simple bursty channel model to illustrate the effects. We simulate that time is divided into 100 ms intervals, with each interval corresponding to 3 packets (frames) for a frame rate of 30 fps. Each interval may be in either a good state or a bad state. In a good state, 3 consecutive packets are received; while in a bad state, 3 are lost. Each time interval is assumed to be independent and identically distributed (Bernoulli), with the probability that a time interval is in the bad state is 0.10. The average packet loss rate is therefore also 0.10, i.e. 10 % of the packets are lost. Our primary reason for choosing this simple channel model is that it simplifies interpretation of the results.

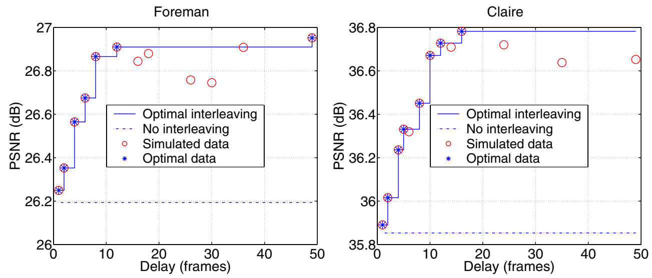

Fig. 5. Optimal PSNR versus delay constraint. Experimental data points (eligible interleavers) are marked with circles, and only those corresponding to optimal interleavers are marked with stars.

Video sequences are coded using JM 2.0 of the emerging JVT/H.26L video compression standard. Four standard test sequences in QCIF format are used, *Foreman, Mother-Daughter, Salesman* and *Claire.* Each has 280 frames at 30 fps, and each is coded with a constant quantization level of 16 which produces an average PSNR of about 36 dB. The first frame of each sequence is intra-coded, and all subsequent frames are coded as P-frames. Every 4 frames a slice is intra updated to improve error-resilience, corresponding to an intra-frame update period of $N = 4 \times 9 = 36$ frames. The distortions are obtained by averaging the results for 6 random channel loss realizations shifted across the whole sequence, or, a total of 280×6 loss realizations.

For different delay constraints, all of the eligible interleavers are identified and their performances are then estimated. The PSNRs for *Foreman* and *Claire* with different interleavers as a function of delay constraint are shown in Fig. 5. Note that the PSNRs shown in Fig. 5 are the averaged results for all frames, including both good and error-afflicted frames, in a sequence; while the quantities shown in Fig. 1 and 4 are the normalized total distortion of the error-afflicted frames only. For a particular delay constraint C_{delay}, an optimal interleaver (n_{opt}, d_{opt}) is found using the algorithm in Subsection 3.3. Although many eligible interleavers are tested and marked in the plots, only those that provide optimal performance are marked with stars. For example, for $C_{delay} = 12$ frames, $(n_{opt}, d_{opt}) = (7, 3)$ is found. For $12 < C_{delay} < 49$, although larger interleavers are eligible under the delay constraint, none of them is optimal, as shown by the circled data points in Fig. 5, since the optimal interleaver is still (7, 3), which has a delay of 12. For this reason, the PSNR curve in the plots is stair-cased, which is the outer bound of all data points tested. This figure also illustrates that larger interleavers are not necessarily more effective. For a given burst loss behavior, increasing the interleaver size beyond a certain point does not improve the effectiveness. In particular, for short burst lengths, a small interleaver with low latency is sufficient to provide most of the gain.

It is observed from Fig. 5 that using interleaver $(5, 3)$ with a delay of 8 frames (267 ms) provides a gain of 0.67 dB over the case of no interleaving for *Foreman*. Using interleaver $(7, 3)$ with a delay of 12 frames (400 ms), increases the gain to 0.72 dB.

For *Claire* sequence, gains of 0.81 dB and 0.93 dB are achieved for delays of 333 and 533 ms, respectively. It is also

Table 1. Gain in PSNR (dB) provided by the optimal interleaver for different delay constaints.

Delay (frame/ms)	Foreman	Mother	Salesman	Claire
8/267	0.67	0.16	0.32	0.60
12/400	0.72	0.24	0.36	0.87
16/533	0.72	0.24	0.36	0.93

observed that as the delay constraint exceeds 16 frames, the performance of eligible interleavers does not further increase. This is because as the interleaver becomes larger, each burst loss is separated farther apart, and the isolated losses of one burst loss come closer to the isolated losses corresponding to the next burst loss in the sequence. Also note that as the distances between isolated losses increases (and approach the intra-period) the losses begin to act as independent losses and any further increase in the spacing between losses does not lead to further reduction in total distortion.

The gains in PSNR for all four video test sequences examined in the experiments are listed in Table 1, for different delay constraints and corresponding optimal interleavers. Note that these gains are obtained without requiring any increase in bitrate. The optimal interleavers with delay of 8 and 12 frames are $(5, 3)$ and $(7, 3)$, respectively, for all sequences, which indicates the optimal interleaver's weak dependence on the sequence.

5: Conclusions

Burst losses in packetized compressed video produce greater total distortion than an equivalent number of isolated losses. This paper proposed a simple packet interleaving (or packet scheduling) scheme to combat bursty channel losses by converting the burst losses into isolated losses. Given knowledge of the channel burst loss characteristics, the optimal interleaver is determined for a given delay constraint, in order to minimize the total distortion as seen by the receiver. Specifically, the optimal interleaver is determined by using an accurate model for predicting the distortion that results for different packet loss patterns. The proposed approach of packet interleaving is simple, provides an improvement without requiring any additional bit rate, and can be used in conjunction with other forms of error-resilient coding and communication.

References

[1] Theodore S. Rappaport, *Wireless Communications: Principles and Practice*, Prentice Hall, 2nd edition, 2001.

[2] Y.J. Liang, J.G. Apostolopoulos, and B. Girod, "Analysis of packet loss for compressed video: Does burst-length matter?," *IEEE International Conference on Acoustics, Speech, and Signal Processing, ICASSP '03*, submitted.

[3] S. Wenger, G.D. Knorr, J. Ott, and F. Kossentini, "Error resilience support in H.263+," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 8, no. 7, pp. 867–877, Nov. 1998.

[4] J.Y. Liao and J.D. Villasenor, "Adaptive intra update for video coding over noisy channels," in *Proceedings IEEE International Conference on Image Processing*, Lausanne, Switzerland, Sept. 1996, vol. 3, pp. 763–6.

[5] R.O. Hinds, T.N. Pappas, and J.S. Lim, "Joint block-based video source/channel coding for packet-switched networks," in *Proceedings of the SPIE VCIP 98*, San Jose, CA, Oct. 1998, vol. 3309, pp. 124–33.

[6] K. Stuhlmüller, N. Färber, M. Link, and B. Girod, "Analysis of video transmission over lossy channels," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 6, pp. 1012–32, June 2000.

[7] R. Zhang, S.L. Regunathan, and K. Rose, "Video coding with optimal inter/intra-mode switching for packet loss resilience," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 6, pp. 966 –976, June 2000.

[8] W. Tan and A. Zakhor, "Video multicast using layered FEC and scalable compression," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 11, no. 3, pp. 373–87, Mar. 2001.

[9] T. Wiegand, N. Färber, and B. Girod, "Error-resilient video transmission using long-term memory motion-compensated prediction," *IEEE Journal on Selected Areas in Communications*, vol. 18, no. 6, pp. 1050–1062, June 2000.

[10] M. Budagavi and J.D Gibson, "Multiframe video coding for improved performance over wireless channels," *IEEE Transactions on Image Processing*, vol. 10, no. 2, pp. 252 – 265, Feb. 2001.

[11] Yi J. Liang and B. Girod, "Low-latency streaming of pre-encoded video using channel-adaptive bitstream assembly," in *Proc. of the IEEE International Conference on Multimedia and Expo (ICME)*, Lausanne, Switzerland, Aug. 2002, vol. 1, pp. 873 –876.

[12] P. A. Chou and Z. Miao, "Rate-distortion optimized streaming of packetized media," *IEEE Trans. Multimedia*, Feb. 2001, submitted. http://research.microsoft.com/~pachou.

[13] M. Kalman, E. Steinbach, and B. Girod, "R-D optimized media streaming enhanced with adaptive media playout," in *Proc. IEEE International Conference on Multimedia, ICME-2002*, Lausanne, Switzerland, Aug. 2002, vol. 1, pp. 869 –872.

[14] S. J. Wee, W. Tan, J.G. Apostolopoulos, and M. Etoh, "Optimized video streaming for networks with varying delay," in *Proc. of the IEEE International Conference on Multimedia and Expo (ICME)*, Lausanne, Switzerland, Aug. 2002, vol. 2, pp. 89 –92.

[15] J. Chakareski, P.A. Chou, and B Aazhang, "Computing rate-distortion optimized policies for streaming media to wireless clients," in *Proc. Data Compression Conference*, Snowbird, UT, Apr. 2002, pp. 53–62.

[16] John G. Apostolopoulos, "Reliable video communication over lossy packet networks using multiple state encoding and path diversity," in *Proceedings Visual Communication and Image Processing*, Jan. 2001, pp. 392–409.

[17] S. Lin, S. Mao, Y. Wang, and S. Panwar, "A reference picture selection scheme for video transmission over ad-hoc networks using multiple paths," in *Proc. of the IEEE International Conference on Multimedia and Expo (ICME)*, Aug. 2001.

[18] Y.J. Liang, E. Setton, and B. Girod, "Channel-adaptive video streaming using packet path diversity and rate-distortion optimized reference picture selection," in *Proceedings of IEEE 5th Workshop on Multimedia Signal Processing*, St. Thomas, Virgin Island, Dec. 2002.