New Compression Techniques for Robust and Scalable Media Communications

1. Research Team

Project Leader: Prof. Antonio Ortega, *Electrical Engineering*

Graduate Students: Zhourong Miao, Phoom Sagetong, Huisheng Wang

Industrial Partner(s): DRS

2. Statement of Project Goals

As international standards (the latest examples being MPEG-4 and JPEG 2000) continue to be developed and provide widely adopted solutions for media communications, some of the most important research problems in multimedia compression are becoming those where pure compression performance is not the main requirement. Instead, in this project we are focusing in a number of system-level issues that will need to be resolved before a full deployment of media communications can take place.

The most popular and widely deployed media communication applications are likely to be those that are accessible over the Internet, both through the currently predominant wired access and through emerging wireless access channels. This requires the design of compression techniques that are robust to channel impairments (variations in bandwidth and delay, packet losses, bit errors, etc). Our work in this area concentrates on: (i) Multiple Description Coding techniques, where compression preserves some redundancy to increase the robustness to packet losses, (ii) caching, where parts of media streams can be stored at intermediate nodes in the network to enable playback to continue even if localized congestion occurs, (iii) streaming, where it is shown that the order in which information is streamed has a significant effect on end-to-end quality, and (iv) Multiple Description Layered Coding (MDLC), where the advantages of layered coding and multiple description coding techniques are combined to provide robust video communication by adapting to the varying channel conditions on the fly.

3. Project Role in Support of IMSC Strategic Plan

While our work is mostly concerned with compression, we are addressing a series of issues that arise when compression technology gets integrated into real network systems. Approaches for robust communication and caching should become useful in applications where video transmission over existing networks, or future wireless links, is required.

We are currently working with other IMSC researchers to integrate our work on caching and streaming within the server systems they are developing (e.g., Yima). As an example of this interaction, the scheduling for streaming applications considered should be integrated with other mechanisms for rate control that are embedded in video servers being studied within IMSC. Our research here complements that work in that we provide mechanisms for changing the transmission rate by modifying the quality of the video streams, while other techniques under consideration limit themselves to modifying the rate at which they transmit data, without actually modifying the data itself. The work presented here will become even more practical as progress is made in developing more efficient scalable video algorithms, e.g., those emerging from studies of Fine Grain Scalability (FGS) within MPEG-4.

Compression is a key component in providing media delivery to the broadest range of terminals and using all types of available transport mechanisms. The new compression technologies described in this project support the IMSC vision by allowing the integration of media delivery into current and future networks (e.g., the Internet) and physical channels (e.g., UWB).

4. Discussion of Methodology Used

Multiple description coding. The specific scenario we consider is that of transmission over a shared network such as the Internet, where there exists no quality of service guarantees. Thus, multimedia data is packetized for transmission but, given that there are no priorities, any packet could be lost during transmission. We consider multiple description coding (MDC) techniques [1], which introduce redundancy in the source coding so that even if some packets are lost, an approximate version of the original signal can be recovered without need for retransmission. Note that MDC techniques can be combined with forward error correction (FEC) to provide special protection to some of the information. MDC has become popular for real time applications as it provides graceful degradation and does not require retransmission. In a MDC system, if only one channel is received then the decoder can reconstruct to an acceptable quality level. However, if both the channels are received, information from one channel augments information from the other, to achieve higher quality than with one channel alone. This makes the system scalable as the quality improves when the packet loss rate is lower. Since MDC does not require prioritized transmission it can be used with current Internet protocols, such as UDP. A fundamental design parameter in an MDC system is the level of redundancy. Higher redundancy provides more error protection, and therefore ideally one would like to match the redundancy to the channel characteristics, and in particular to be able to change it if the channel conditions are time varying.

Video caching. In this work we target the problem of proxy caching for multimedia web objects, and in particular video sequences that are accessed in a streaming mode by users. Since the size of these streaming video objects is very large, caching the whole video sequence at the proxy may not be efficient if memory/disk space is limited. We found that even if only part of the video stream is cached it is still possible to improve the playback quality. Our goal is to define algorithms to select the optimal set of video frames that should be cached, given that the memory available for caching is limited. We found that for the different network environments (namely, QoS networks and best-effort networks) the objectives in optimizing cache performance are different. For example, in QoS networks, since bandwidth reservation is costly, an important objective in designing a caching strategy may be to minimize the bandwidth requirements. Alternatively, in best-effort networks, robustness of continuous playback during periods of network condition is more important. We have developed two caching algorithms corresponding to these network cases.

Streaming media delivery. We study the problem of scalable layered streaming media delivery over a lossy channel. The goal is to find an optimal transmission policy to achieve the best playback quality at the client end. The problem involves some trade-offs such as time-constrained delivery and data dependencies. For example, a layer should be dropped before transmission if it already has a delay such that it cannot be played before its scheduled time. Moreover, less important layers with near-playback-time may also be dropped or delayed for delivery in order to save bandwidth for other layers with a high priority.

Multiple description layered coding. By combining the advantages of Layered Coding (LC) and Multiple Description Coding (MDC), this approach can provide reliable video communication over a wider range of network scenarios and application requirements. MDLC improves LC in that it introduces redundancy in each layer so that the chance of receiving at least one description of base layer is greatly enhanced. Though LC and MDC are each good in limit cases (e.g., long end-to-end delay for LC vs. short delay for MDC), the proposed MDLC system can address intermediate cases as well. Same as a LC system with retransmission, the MDLC system can have a feedback channel to indicate which descriptions have been correctly received. Thus a low redundancy MDLC system can be implemented with our proposed runtime packet scheduling system based on the feedback information. The goal of our scheduling algorithm is to find a proper on-line packet scheduling policy to maximize the playback quality at the decoder. Previous work on scheduling algorithms has not considered multiple decoding choices due to the redundancy between data units, because of the increase in complexity involved in considering alternate decoding paths. In this work, we introduce a new model of Directed Acyclic HyperGraph (DAHG) to represent the data dependencies among frames and layers, as well as the data correlation between descriptions. The impact of each data unit to others is represented by messages passing along the graph with updates based on new information received.

5. Short Description of Achievements in Previous Years

Previously proposed MDC schemes do not allow the level of redundancy to be easily changed. For example, this may require changing the transform at both encoder and decoder. Clearly, these approaches are limited in their ability to adapt to changing transmission conditions. In our work [10], we have proposed a simple approach for MDC that uses a polyphase transform and deterministic redundancy (e.g., each sample of input data is transmitted several times, with different coding rates). This approach is useful in that it greatly simplifies the design of a MDC scheme, since the rate allocation determines the amount of redundancy. Moreover, it provides a great deal of flexibility as it enables the choice of redundancy to be almost arbitrary. We have introduced an optimal bit allocation algorithm that allows us to select the amount of redundancy to be introduced in the signal that best matches a given target packet loss rate [11]. Our results show significant differences between optimal and suboptimal choices of redundancy for a given packet loss rate. Moreover, given that the decoder remains unchanged when the bit allocation changes, it is possible to adapt very simply to the changes in channel behavior without requiring a change in the packet sizes, or the structure of the decoder. Our results are summarized in Figure 1, where, as was to be expected, the optimal number of polyphase components in a description (i.e., the best level of redundancy) increases as the probability of packet loss increases.

In addition we studied the use of MDC for correcting errors in DPCM and other memory based compression techniques [12,17]. For typical image/audio sources, which have memory, it is necessary to use these kinds of methods to achieve reasonable compression performance. However, approaches such as DPCM have the drawback of catastrophic error propagation, i.e., a 1 bit channel erasure in the prediction error will cause all samples after this erasure to be decoded incorrectly. In our work we propose to use a low rate description to correct erasures in the high rate description. This is a novel approach for erasure recovery in memory coders and results are very encouraging.

In the future, we propose to extend our work to motion prediction scheme in MPEG coders. Experimental results have shown that the proposed system can estimate data lost due to large burst of erasures.

We have also compared MDC to Layered Coding under various network conditions using a network simulator [13]. We find that for applications that require short latency, e.g. video conferencing, and for networks with large round-trip-time, MDC outperforms Layered Coding in case of even small packet losses over the network.

Video caching. Previously, we have developed a caching algorithm for video delivery over best-effort networks. This work has been extended to QoS networks, and another caching algorithm has been proposed especially for QoS network environment.

Streaming media delivery. During the previous years, we have proposed a packet scheduling algorithm for scalable encoded streaming media. The simulation has been done with simple scalable audio streaming. This year, the algorithm has been extended to scalable video stream with more complicated dependencies.

Since our previous work [11] demonstrated that bit allocation could be used to adjust the level of redundancy to match channel conditions we have studied this bit allocation problem in the context of a progressive wavelet coder such as SPIHT (Set Partitioning in Hierarchical Trees) [14,15]. Note that the bit allocation problem we have solved is useful not only for MDC but also other applications such as Region of Interest coding (ROI). In both cases, wavelet coefficients are divided by different factors before coding in order to enable different bit allocation to different regions. The basic idea in SPIHT (and in similar algorithms) is to perform bit plane encoding of the wavelet coefficients, so that larger coefficients start being transmitted first. By selecting appropriate weighting factors one can ensure that the coefficients in one region tend to be larger so that they are transmitted first and refined to higher resolution than those in other regions. Essentially this is equivalent to providing unequal bit allocation to each of the regions. While this technique has been used for ROI coding, we are the first to propose using it for MDC as well. In this work, we call ***priority-scaling factor*** (psf) the dividing factor. Our main contribution is to provide an analytical technique to determine what the psf value should be, given criteria such as relative importance of the regions or degree of redundancy in an MDC coder. Our approach is based on a novel Rate-Distortion (RD) model, which is an extension of Mallat's model to incorporate the specific characteristics of our proposed psf technique. In other words, our approach enables us to model analytically the RD values for any psf values, so that the optimal psf values can be calculated.

Standard-compatible MDC video coding: Another important contribution of our work is to demonstrate that it is possible to achieve some of the robustness of MDC while being fully compatible with existing standard video coders. We propose an unbalanced MDC system, where one of the descriptions has a high resolution and the other a low, but acceptable resolution (quality). The low-resolution description (LR) is explicitly used as redundancy, to be decoded in cases of losses in the high-resolution description (HR). Each of the descriptions are encoded and decoded by a standard compression system, in our case the H.263 video encoder. There has been limited work on MDC based video systems mainly because of the *prediction loop mismatch* problem. In a prediction-based decoder, if the prediction is based on information unavailable at the decoder, there is a prediction loop mismatch leading to poor MSE performance. Our system avoids the prediction loop mismatch problem because the central decoder is the same as the high-resolution decoder, with the LR description being used only when there are no losses in the HR description. Thus, the LR signal, unlike the LR signals in layered coding, is not meant to be usable by itself.

In our approach we obtain HR first and then for each block in each HR coded frame we decide to select only a certain number of those transform coefficients to be included in LR. This is done according to a criterion where given a total amount of allowable redundancy we choose the coefficients that minimize the distortion in case of loss. The motion vectors are duplicated in both the descriptions and it is assumed that there are no correlated losses, i.e., erasures do not occur at the same location in both the descriptions.

Video caching. We have developed two caching algorithms for QoS and best-effort networks. For caching in QoS networks, we first find the minimum bandwidth that has to be reserved for the video playback, based on the video data characteristics. Then we analyze how caching different parts of the video sequence will affect this bandwidth requirement. The proposed caching algorithm is therefore designed to maximally reduce the bandwidth requirement with only limited caching space by selecting proper frames of the video sequence to be cached.

For caching in best-effort networks, we have identified several cost functions for this optimization, along with optimal and near optimal algorithms to minimize the cost. As an example, we consider the "decoder buffer content trace" and use it to locate frames that correspond to times when the decoder buffer "risks" to be in underflow if network congestion occurs. Such risks occur when the decoder buffer contains only a few frames, so that the client application can only play from the client buffer for a limited time. Thus our algorithms are based on determining the times at which maximum risk of decoder buffer underflow and then caching frames such that the number of frames in the decoder buffer at those times is increased. By doing this we ensure that should congestion occur there will be more frames to play from buffer while no data arrives from the network. We have developed an algorithm to select frames to be cache within certain cache space budget to maximize the robustness. More details can be found in [20].

Streaming media delivery. We propose a framework for scalable streaming media delivery that involves a novel scheduling algorithm called *Expected runtime Distortion Based Scheduling*, EDBS. This algorithm decides the order in which packets should be transmitted in order to improve client playback quality in the presence of channel losses.

The *Expected runtime Distortion* considers all the important elements that will affect the playback quality at client end, including, (i) the information contained in each data packet, which determines the reconstructed quality of the video, (ii) the playback end-to-end delay, which determines the duration of packet lifetime, (iii) the channel conditions, which affect the packet loss, and (iv) the packet dependencies, which produced by the scalable compression codec. All these factors are evaluated in our proposed transmission algorithm, which identifies the best schedule for packet delivery based on the result of this evaluation so that the playback quality is maximized. The detail of this algorithm and evaluation procedure can be found in [28]. We also provide a fast greedy search algorithm that achieves almost the same performance as an exhaustive search technique (98% of the time it results in the same schedule) with very low complexity and is applicable for real-time applications.

5a. Detail of Accomplishments During the Past Year

Multiple description layered coding. Our general approach for multiple description layered coding (MDLC) [32] uses an MDC encoder to generate two base layer descriptions BL_1 and BL_2. For each base layer description BL_i, we create its corresponding enhancement layer EL_i. An extra enhancement layer description EL_0 is created as the difference between the original video signal and the combined signal of both base layer descriptions. The MDLC decoder selects which *EL* stream to decode given what base layer was received, and finally reproduces the signal by combining the base layer and enhancement layer information. The key advantage of our MDLC scheme is that it combines the hierarchical scalability provided by LC with the reliability introduced by adding redundancy into base layer with multiple descriptions. It provides the flexibility for the scheduling algorithm to choose the right base layer and enhancement layer descriptions based on the current network conditions and feedback information of previous transmission history.

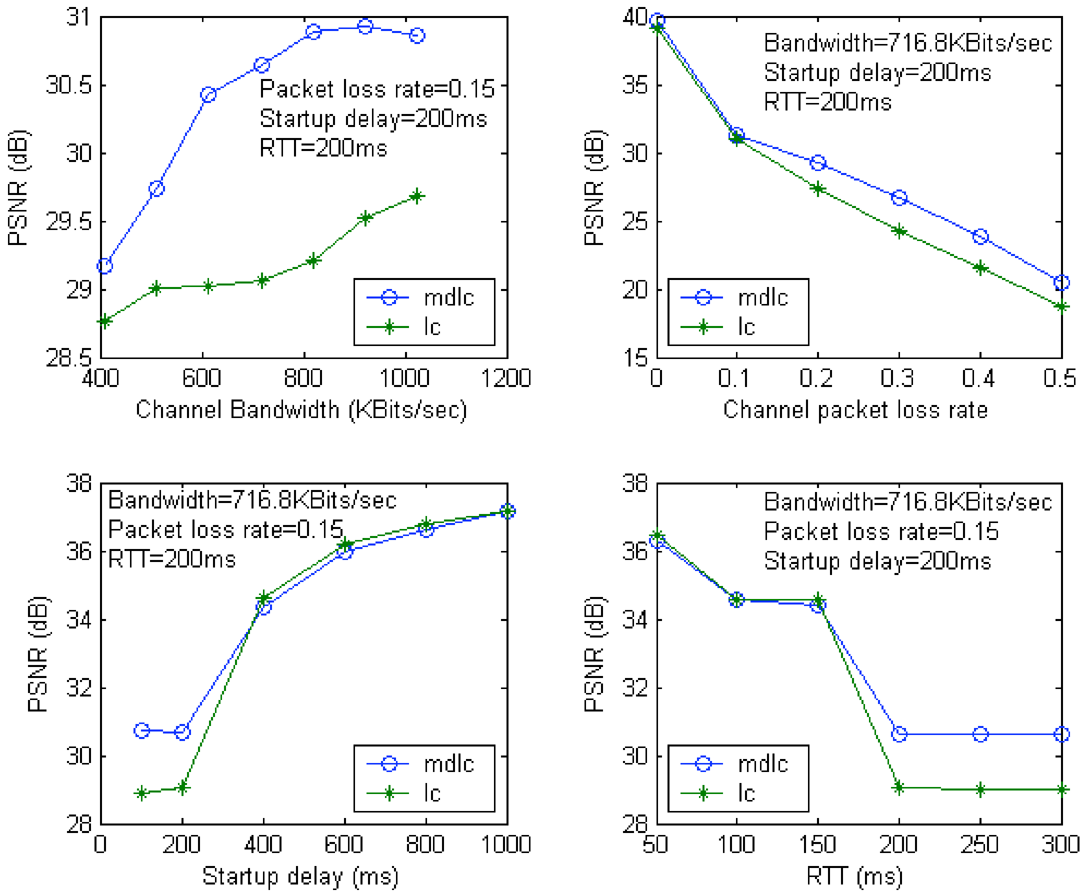

Figure 9. Comparison of PSNR for *Stefan* sequence between MDLC and LC with various parameters. The results show that by using the proposed MDLC system the playback quality improves about 2dB compared to the traditional LC system in the scenarios when a channel has large packet loss rate, long RTT or an application has stringent delay requirement.

For these experiments we have assumed that the transmitter would incorporate a run-time scheduler and a time window control. In the experiments shown here an observation of the channel and an *a priori* model, are used by the transmitter to select the most important packets. Our double time window control (DTWC) scheme, is different from the traditional method [33] in that it introduces separate windows of transmission opportunities (WoP) w_1 and w_2 for different descriptions. The DAG model used in previous research [33, 28] cannot be directly applied to a MDLC system because it can only consider one decoding choice. Here we introduce a new model called Directed Acyclic HyperGraph (DAHG). A DAHG is like a normal DAG, but each vertex is composed of a clique which contains a set of nodes and every pair of nodes are connected by an undirected edge to indicate their particular redundancy relationship. An *expected runtime packet distortion* is calculated for each candidate packet in the transmission buffer based on the channel conditions, data dependencies and correlation between packets, and the previous transmission history. At a given time, the packet with the greatest expected distortion is selected to be sent.

Figure 9 shows an experimental result on comparing the performance of LC and MDLC both of which use on-line scheduling. The results show that our proposed MDLC scheme provides more efficient and robust video communication in a wider range of network scenarios and application requirements.

6. Other Relevant Work Being Conducted and How this Project is Different

Multiple Description Coding. Research work on robust communication using multiple description coding (MDC) techniques [1] has been actively conducted in recent years in both industry, e.g., Bell Labs [2,3] or AT&T Research [4,5,6], as well as in academia, e.g., the University of Illinois, at Urbana-Champaign [6,7], the University of California, Berkeley [8], the University of Washington [9] and Princeton University [5]. In an MDC scheme the source coder generates two or more compressed versions (or descriptions) of the signal, such that if *either* one is lost a reasonably good reproduction of the signal at the decoder can still be achieved. Conversely, high quality decoding is achieved when both descriptions are received. Obviously, the added robustness of sending two descriptions comes at the cost of redundancy and therefore worse coding performance. Some recent work has focused on transform-based MDC techniques [2], where the transform itself preserves some redundancy, while other researchers have proposed MDC approaches based on scalar quantization [5,6]. More recent approaches have demonstrated the benefits of combining MDC with forward error correction (FEC) to increase the robustness and enable unequal error protection to the various components of the signal [7,8,9].

The main novelty in our MDC approach is that we use explicit redundancy to generate the descriptions, that is, each sample of the signal can be coded and transmitted more than once. As will be discussed, this approach greatly reduces the complexity of designing the system. Previously proposed techniques tend to be optimized for an expected level of packet losses, and will not perform well in situations where there is a mismatch between expected and actual packet losses. In our proposed approach switching between different levels of redundancy is very straightforward so that a system can be made adaptive to changes in packet loss rates without requiring a complete redesign of encoder/decoder. We also introduce the model-based analytical bit allocation based on a novel rate-distortion model inside each description where we use the divide-and-multiply method to vary the quality of different parts of an image. Another major contribution of our MDC work is towards robust transmission of predictive encoded data. We have developed an algorithm for estimating data lost due to channel noise for a DPCM encoder and a MDC based video encoder that is compliant with the H.263 standard video codec. Our work on MDC has been published in [10-17].

Video caching. Previous work on video caching has focused on three approaches. Caching complete video streams is the simplest technique, but has the disadvantage that it is inefficient in terms of storage, especially when the caching space is limited. As an alternative, video prefix caching [13] calls for storing complete video streams first but then, as disk space becomes scarcer, only parts of the video streams (namely, the initial parts of the stream or prefixes) are stored. This has the advantage of reducing initial latency, since the data that the client needs before starting to decode is all in the cache. This technique can also be used to smooth the transmission of VBR video data as in [14] or [15]. The third approach, referred as “video staging", splits the video frames into two parts and caches one part of them into the proxy. This approach can improve the streaming performance without caching the entire video objects. However, in the case of network congestion in the server-client channel, the cached video data may not be useful for playback, since only partial frame data is available there.

Our proposed technique [17] is a frame-wise selective caching, which selects proper frames to be cached according to the caching storage, video object and client buffer constraints. Furthermore, in designing our caching algorithms, we consider two important network environments, namely, QoS networks (e.g., ATM, where it is possible to reserve network resources) and best-effort networks (e.g., the Internet). This is more general than the work in [18] or [19], which only consider the QoS network environment. The caching performances are measured differently in those networks, and thus, we propose two different caching approaches for corresponding network environment to improve the overall performance.

Finally, approaches based on layered video caching [16] are also becoming popular. Here only certain layers are stored for each frame. A drawback of these approaches is that they rely on the availability of scalable video codecs, which have not been broadly used to date. However, it is likely that efficient scalable codecs will become available in the near future, and we are investigating the extension of our techniques to the scalable video case.

Streaming media delivery. The existing transmission protocols for data packets over the Internet are not very suitable for real-time media delivery. Although there have been developments in special protocols for streaming applications, such as RTP/RTCP/RTSP, these protocols still have drawbacks in transmitting media compressed with complex data dependencies. In particular, as scalable, or layered, media representations become more available, there is a need to develop specific transmission mechanisms for these types of data that can exploit their inherent advantages over non-scalable formats. Retransmission strategies under delay constraints have been studied in [24][25] but have not considered scalable media. In [26] a rate control algorithm for delivering MPEG-4 video over the Internet was proposed with a priority re-transmission strategy for the recovery of lost packets, which considers the constraint to prevent the receiver buffer underflow. This is achieved by giving priority to retransmission of base (lower) layers, which are necessary for decoding the corresponding enhancement layers. However this work did not address the problem of the delivery order of new packets in the sender's buffer.

Our proposed framework solves the problem of selecting the delivery order of both the *new* and *lost packets* (i.e., those that have to be retransmitted) and can therefore complement other rate control schemes, which choose the bit rate for each frames in the video. Our experimental results show that this approach can improve the playback quality by 2 to 4 dB. Podolsky *et al* [23] have studied policies for scheduling the delivery of scalable media data. In their work, a Markov chain analysis is used to estimate the average distortion of each candidate policy. For each set of system parameters an optimal steady state policy is found. This design is performed off-line. On-line adaptation between different policies is achieved by estimating the system parameters and switching to the policy that is most suitable. The main difference in our work is that no attempt is made to define deterministic policies and rather a simple algorithm is used to determine on-line which packet should be sent next. The fact that policies do not have to be defined in advance allows more flexibility, e.g., in our approach we can use the exact distortion values for each frame, instead of making the assumption these are always the same as in [23].

Multiple description layered coding. Our work is an extension of recent work at USC and elsewhere on scheduling algorithms for layered coded video. Chou and Miao [33] provide a rate-distortion optimized approach of packet scheduling over a lossy packet network. Based on a simple source model of a Directed Acyclic Graph (DAG), the proposed approach can allocate bandwidth among packets and schedule them in such a way that a Lagrangian cost function of expected distortion and rate is minimized. To reduce the scheduling complexity,

10. References

[1] A. A. El-Gamal and T. M. Cover, “Acehievable rates for multiple descriptions”, IEEE Trans. Information Theory, vol. IT-28, no. 6, pp. 851-857, Nov. 1982.
[2] V. K. Goyal and J. Kovacevic, "Optimal multiple description transform coding of Gaussian vectors", in Proc. of IEEE Data Compression Conference, 1998.
[3] V. Goyal and J. Kovacevic, "Method description for transform coding using optimal transforms of arbitrary dimensions," US Patent, February 25, 1998. Filed.
[4] V. Vaishampayan, "Design of Multiple description Scalar Quantizers," IEEE Trans. Comm., vol. 39, no. 3, pp. 821-834, Nov. 1993.
[5] M. T. Orchard, Y. Wang, V. Vaishampayan, and A. R. Reibman, "Redundancy rate-distortion analysis of multiple description codingusing pairwise correlating transforms", in Proc. ICIP'97.
[6] S. D. Servetto, K. Ramchandran, V. Vaishampayan, and K. Nahrstedt, "Multiple descriptions wavelet based image coding", in Proc. ICIP'98.
[7] D.G. Sachs, R. Anand and K. Ramchandran, “Wireless image transmission using multiple-description based concatenated codes,” in Proc. IVCP 2000
[8] R. Puri and K. Ramchandran, “Multiple Description coding using forward error correction codes,” Proc. Asilomar Conf. on Signals and Sys., Pacific Grove, CA, Oct. 1999.
[9] A. Mohr, E. Riskin and R. Ladner, “Graceful degradation over packet erasure channels through forward error correction,” in Proc. DC 1999.
[10] W. Jiang and A. Ortega, “Multiple Description Coding via Polyphase Transform and Selective Quantization”, in Proc. VCIP’ 99.
[11] P. Sagetong and A. Ortega, “Optimal bit allocation for channel-adaptive multiple description coding”, in Proc. of IVCP 2000.
[12] R. Singh, and A. Ortega, “Multiple Description Coding Bases Erasure Recovery for State based Compression Scheme”, in Proc. MMSP’99.
[13] R. Singh, A. Ortega, L. Perret, and W. Jiang, “Comparison of Multiple Description Coding with Layered Coding using Network Simulator”, in Proc. IVCP 2000.
[14] P. Sagetong and A. Ortega, “Analytical model-based bit allocation for wavelet coding with applications to Multiple Description Coding and Region Of Interest coding,” IEEE International Conference on Multimedia and Expo (ICME 2001), Aug 2001.
[15] P. Sagetong and A. Ortega, “Rate-distortion model and analytical bit allocation for wavelet-based region of interest coding,” submitted to IEEE International Conference on Image Processing (ICIP 2002).
[16] D. Comas, R. Singh and A. Ortega “Rate-distortion optimization in a robust video transmission based on multiple description coding”, in Proc. MMSP 2001.
[17] R. Singh and A. Ortega, “A Robust DPCM scheme based on Multiple Description Coding”, accepted for publication in special issue of Journal of VLSI Signal Processing.
[18] S. Acharya, “Techniques for improving multimedia communication over wide area networks”, Ph.D. Thesis Cornell University, 1999.
[19] S.Sen and J. Rexford and D. Towsley, “Proxy Prefix Caching for Multimedia Streams”, in proc. of IEEE Infocom99, Mar., 1999, New York, USA.
[20] Y. Wang, Z.-L. Zhang, D. Du and D. Su, “A network conscious approach to end-to-end video delivery over wide area networks using proxy servers”, in Proc. of IEEE Infocom, Apr., 1998.

[21] Z. Miao and A. Ortega, "Proxy caching for efficient video services over the Internet", in proc. of 9th International Packet Video Workshop (PVW '99), April, 1999, New York.
[22] R. Rejaie, M. Handley, H. Yu and D. Estrin, "Proxy Caching Mechanism for Multimedia Playback Streams in the Internet", in Proc. of the 4th Web Caching Workshop, San Diego, CA, Mar.1999.
[23] Real Time Streaming Protocol (RTSP), proposed standard, RFC 2326, April, 1998.
[24] C. Papadopoulos and G. Parulkar, "Retransmission-Based Error Control for Continuous Media Applications", in *Proc. NOSSDAV*, April, 1996.
[25] M. Lucas, B. Dempsey and A. Weaver, "MESH: distributed error recovery for multimedia streams in wide-area multicast networks", in *Proc. IEEE Int. Conf.*