# LAYERED MULTIPLE DESCRIPTION CODING


*Philip A. Chou, Helen J. Wang, and Venkata N. Padmanabhan*

Microsoft Corporation
One Microsoft Way, Redmond, WA 98052-6399
{pachou,helenw,padmanab}@microsoft.com



## ABSTRACT

Multiple description codes address the problem of unreliable channels by means of independent descriptions, while layered codes address the problems of heterogeneous client bandwidths and dynamic network congestion by means of sequences of layers. With the goal of achieving both robustness to unreliable channels and adaptivity to client bandwidth heterogeneity and network congestion, we explore constructions for *layered multiple description codes*, where base layer descriptions can be transmitted to low bandwidth clients, while both base and enhancement layer descriptions can be transmitted to high bandwidth clients. The low bandwidth client quality is an increasing function of the number of base layer descriptions received, while the high bandwidth client quality is a bivariate increasing function of both the number of base layer descriptions and the number of enhancement-layer descriptions received. By optimizing the base layer descriptions for the low bandwidth client, in our construction, the high bandwidth client pays a penalty of 1.4 dB relative to a non-layered multiple description code optimized to the high bandwidth client.


## 1. INTRODUCTION

Multiple description coding (MDC) has been proposed for use in packet audio and video transmission systems as a means of combatting both packet loss and component failure, in a variety of application scenarios [1]. In this paper we are motivated by the multicast scenario, as explored in [2, 3]. In [2], multiple descriptions are striped across multiple packets, and are transmitted to a collection of clients over IP multicast, thereby ameliorating the loss of packets due to congestion. In [3], multiple descriptions are striped across multiple distribution trees, and are transmitted to a collection of clients over application-level multicast in a peer-to-peer setting, to ameliorate the failure of unreliable hosts. Both of these works assume that the client population is homogeneous in bandwidth, so that a fixed bit rate is transmitted to each client. In contrast, several previous works on multicast of audio and video, most notably Receiver driven Layered Multicast, exploited the properties of layered coding to efficiently transmit to clients at different bit rates [4, 5]. Low-bandwidth clients would receive only a base layer, for example, while high-bandwidth clients would receive both a base layer and an enhancement layer. In this paper we explore *layered multiple description codes*, which have the advantages of both layered codes and multiple description codes, by permitting low-bandwidth clients to receive a base MDC layer, for example, while high-bandwidth clients receive both a base and an enhancement MDC layer.

In Section 2 we review a packetization technique called priority encoding transmission and its optimization, on which our layered MDC codes are based, and in Section 3 we present several constructions of layered MDC codes. In Section 4 we present our results, and in Section 5 we present our conclusions.

## 2. MDC BY PRIORITY ENCODING TRANSMISSION

Many methods of multiple description coding have been developed over the years. One particularly efficient and practical method is based on the Priority Encoding Transmission (PET) technique of Albanese et al. [6]. PET is a packetization scheme that combines layered source coding with unequal erasure protection. As illustrated in Figure 1, PET partitions the source into groups of frames; each group of frames is independently source coded into layers of differing importance; each layer $n$ is blocked into source blocks of length $K_n$ bytes; and each source block is expanded into channel codewords of length $N \geq K_n$ bytes using an $(N, K_n)$ Reed-Solomon code (or other minimum distance separable code) such that the channel code rate $K_n/N$ provides erasure protection commensurate with the importance of the layer. Finally, for each group of frames, the length-$N$ channel codewords for all the source blocks in all the layers are packetized into $N$ packets by putting the $i$th byte in each channel codeword into the $i$th packet, $i = 1, \ldots, N$.

The PET packetization scheme has the property that if any $K$ out of the $N$ packets are received, then all layers

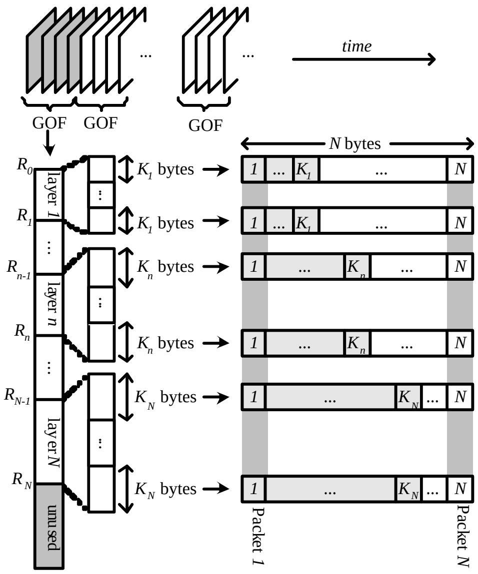


**Fig. 1**. Packetization for priority encoding transmission.

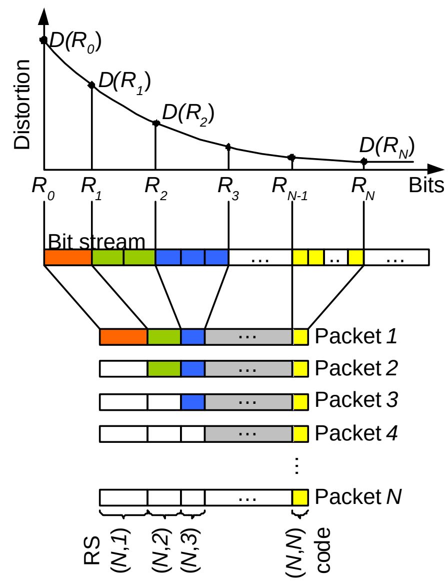


**Fig. 2**. Packetization for priority encoding transmission.

$n$ with $K_n \leq K$ are recovered, while all layers $n$ with $K_n > K$ are lost. Without loss of generality (because a layer may be empty) it can be assumed that there are exactly $N$ layers $n = 1, 2, \ldots, N$, indexed in order of decreasing importance such that layer $n$ is protected with an $(N, K_n)$ code, $K_n = n$. We denote the boundaries of layer $n$ in the encoded bit string for a group of frames (GOF) by bits $R_{n-1}$ and $R_n$, such that $0 = R_0 \leq R_1 \leq \cdots \leq R_N$. This arrangement is illustrated in Figure 2. Thus, regarding the layers as the constituents of an embedded bit string for a GOF, if any $n \leq N$ packets are received, then the initial $R_n$ bits from the embedded bit string for the GOF can be recovered, resulting in distortion $D(R_n)$, where $D(R_0) \geq D(R_1) \geq \cdots \geq D(R_N)$. In this sense all $N$ packets are equally important; only the *number* of packets received determines the reconstruction quality of the GOF. In this way, the PET packetization scheme is a form of multiple description code. The $n$th packet constitutes the $n$th description for a GOF; the sequence of $n$th packets for GOFs constitute the $n$th description for a media stream.

Originally, Albanese et al. applied the PET packetization scheme to the I, P, and B layers of MPEG video, and they did not optimize the code rates $\{K_n/N\}$ to minimize the end-to-end distortion for a given overall transmission rate. Davis and Danskin [7] showed how to perform this optimization for any number of layers, using a simple slope-matching algorithm. Mohr, Riskin, and Ladner [8] showed, in addition, how to adjust the breakpoints $\{R_n\}$ between layers when each GOF is encoded with a finely embedded source coder. However, they used a greedy search algorithm for this purpose. Puri and Ramchandran [9] showed how to solve this latter problem optimally (to within a relaxation of the integral alignment constraints on the number of bits in each layer) using a fast algorithm, which we use for the experiments in this paper. Mohr, Ladner, and Riskin [10] later showed another efficient optimization procedure, using a different relaxation. And recently Stanković, Hamzaoui, and Xiong [11] presented an extremely efficient algorithm for greedy search from a near-optimal initial condition.

The precise optimization problem can be stated as follows. For each GOF, let $N$ be the desired number of descriptions, let $p_n$ be the probability that $n$ of the $N$ descriptions are received, and let $D(R)$ be the distortion if the first $R$ bytes of the embedded code for the GOF are recovered. That is, $D(R)$ is the operational distortion-rate function for the GOF. Let $\mathbf{R} = (R_0, R_1, \ldots, R_N)$ be the vector of breakpoints (expressed in bytes) describing the PET packetization for the GOF. Then the expected distortion for the PET packetization is

$$D(\mathbf{R}) = \sum_{n=0}^{N} p_n D(R_n),$$

while the rate (in bytes per packet) is

$$R(\mathbf{R}) = \sum_{n=1}^{N} (R_n - R_{n-1})/n,$$

due to the fact that the $(R_n - R_{n-1})$ source bytes of layer $n$ are coded with a total of $(R_n - R_{n-1})N/n$ source plus parity bytes. Note that the rate can be expressed as $R(\mathbf{R}) = \sum_{n=1}^{N} \alpha_n R_n$, where $\alpha_n = 1/(n(n+1))$ for $n = 1, \ldots, N-1$, and $\alpha_N = 1$.

The objective is to find the breakpoints $\mathbf{R}$ that minimize the expected distortion $D(\mathbf{R})$ subject to the constraints $R(\mathbf{R}) \leq R^*$ and $0 = R_0 \leq R_1 \leq \cdots \leq R_N$, where $R^*$ is a target rate in bytes per packet. This can be accomplished by minimizing

$$D(\mathbf{R}) + \lambda R(\mathbf{R}) = p_0 D(R_0) + \sum_{n=1}^{N} p_n D(R_n) + \lambda \alpha_n R_n$$

for some positive Lagrange multiplier $\lambda$, subject to the constraints $0 = R_0 \leq R_1 \leq \cdots \leq R_N$. Finding the appropriate Lagrange multiplier can be performed by a binary search and is not discussed here.

Puri and Ramchandran[9] have shown that if $\alpha_{n+1}/p_{n+1} \geq \alpha_n/p_n$, then the optimal solution must have $R_{n+1} = R_n$. Hence if $\alpha_{n+1}/p_{n+1} \geq \alpha_n/p_n$, then the problem can be reduced to finding the reduced-dimensional vector $\mathbf{R}' = (R_0, \ldots, R_n, R_{n+2}, \ldots, R_N)$ minimizing

$$D(\mathbf{R}') + \lambda R(\mathbf{R}') = p_0' D(R_0') + \sum_{n=1}^{N'} p_n' D(R_n') + \lambda \alpha_n' R_n'$$

subject to the constraints $0 = R_0' \leq R_1' \leq \cdots \leq R_{N'}'$, where

$$\begin{aligned} \mathbf{p}' &= (p_0, \ldots, p_{n-1}, p_n + p_{n+1}, p_{n+2}, \ldots, p_N), \\ \alpha' &= (\alpha_0, \ldots, \alpha_{n-1}, \alpha_n + \alpha_{n+1}, \alpha_{n+2}, \ldots, \alpha_N), \end{aligned}$$

and $N' = N - 1$. By repeatedly performing this reduction, it can be ensured that the sequence $\alpha_n/p_n$ is decreasing, which we henceforth assume.

Now, if we ignore the constraints $0 = R_0 \leq R_1 \leq \cdots \leq R_N$, it is clear that

$$\begin{aligned} &\min_{\mathbf{R}} D(\mathbf{R}) + \lambda R(\mathbf{R}) \\ &\quad = \quad p_0 D(R_0) + \sum_{n=1}^{N} \min_{R_n} \{p_n D(R_n) + \lambda \alpha_n R_n\}, \end{aligned}$$

that is, the minimization can be performed pointwise. Note, however, that since the sequence $\alpha_n/p_n$ is decreasing, the sequence $\{R_n\}$ minimizing $\{p_n D(R_n) + \lambda \alpha_n R_n\}$ is increasing, and hence happens to also satisfy the constraint $0 = R_0 \leq R_1 \leq \cdots \leq R_N$. Thus, after the above reduction ensures that $\alpha_n/p_n$ is decreasing, the constrained problem can be solved by finding the $R_n$ minimizing $p_n D(R_n) + \lambda \alpha_n R_n$ for each $n$.

Figure 3 shows the PET packetization resulting from the above optimization procedure for $N = 32$ descriptions, target rate $R^* = 1250$ bytes per packet, packet loss probability $\epsilon = 10\%$ (yielding $p_n = \binom{N}{n}(1-\epsilon)^n \epsilon^{(N-n)}$) and an operational distortion-rate function $D(R)$ obtained by encoding the first one second of the standard MPEG test sequence *foreman* using a fine-grain scalable (FGS) video codec. Apparently, since the probability of receiving fewer than 18 and more than 26 descriptions is negligible, the optimal breakpoints in this example satisfy $0 = R_0 = \cdots = R_{18} < R_{19} < \cdots < R_{25} < R_{26} = \cdots = R_{32}$.

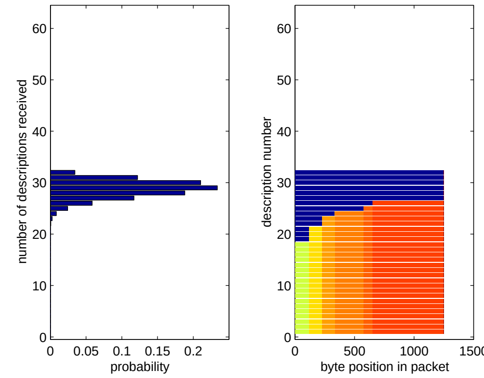


**Fig. 3**. Binomial$(32, 0.10)$ distribution and matched PET packetization. The light-colored bands correspond to source bytes and the dark ones to FEC.

## 3. LAYERED MULTIPLE DESCRIPTION CODING

We now turn to the problem of constructing layered multiple description codes. For simplicity we consider only two layers. In all of our constructions, the first or *base* MDC layer consists of $N_1$ packets per GOF, while the second or *enhancement* MDC layer consists of $N_2$ packets per GOF, such that each packet has a fixed length of $R^*$ bytes. We assume that the base MDC layer is transmitted to each low-bandwidth client over an iid packet erasure channel with probability of packet loss $\epsilon_1$, while both the base and enhancement MDC layers are transmitted to each high-bandwidth client over an iid packet erasure channel with probability of packet loss $\epsilon_2$. Thus, what distinguishes a layered multiple description code from two independent multiple description codes is that the MDC base layer descriptions are shared by the two codes.

### 3.1. Layered MDC by Splitting a Single MDC

One obvious way to construct a layered multiple description code is to optimize a single multiple description code for the high-bandwidth client, and split it into base and enhancement layers by transmitting only a fraction of the descriptions to the low-bandwidth client. This results in the minimum possible distortion for the high-bandwidth client, but a potentially large distortion for the low-bandwidth client. Figure 4 shows the optimal PET packetization for the high-bandwidth client, with $N_1 + N_2 = 64$ and $\epsilon_2 = 10\%$. It is clear from the figure that if the low-bandwidth client receives only $N_1 = 32$ of the 64 descriptions, it will not be

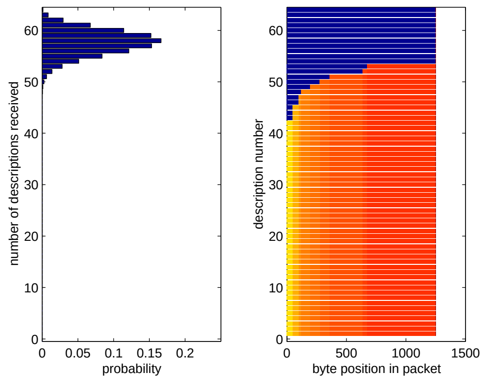


**Fig. 4**. Binomial$(64, 0.10)$ distribution and matched PET packetization.

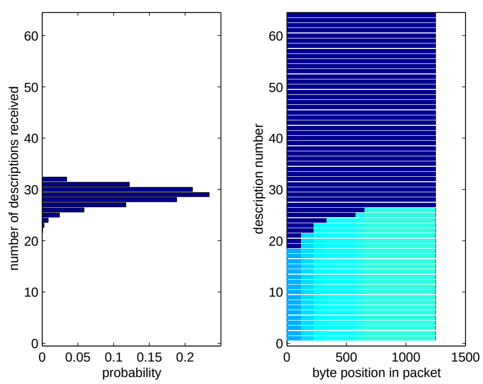


**Fig. 5**. Packetization matched to low-bandwidth clients.

able to recover any source data layer in its entirety (where a "layer" here refers to the source bytes that lie between two consecutive breakpoints), resulting in a large distortion.

At the other extreme it is possible to optimize a single multiple description code for the low-bandwidth client, and transmit all of it, plus $N_2$ additional parity packets, to the high-bandwidth client, as shown in Figure 5 for $N_1 = N_2 = 32$ and $\epsilon_2 = 10\%$. This results in the minimum possible distortion for the low-bandwidth client, but a potentially large distortion for the high-bandwidth client, typically only slightly better than that of the low-bandwidth client.

Between these two extremes, it is possible to optimize the multiple description code for a mixture of low- and high-bandwidth clients. Chou and Ramchandran [2] have shown that if $\Lambda$ is a population of clients such that client $\theta \in \Lambda$ receives exactly $n$ out of $N$ descriptions with probability $p_{\theta,n}$, then for a given set of breakpoints $\mathbf{R} = (R_0, R_1, \ldots, R_N)$, the expected distortion at the client is

$$D_\theta(\mathbf{R}) = \sum_{n=0}^{N} p_{\theta,n} D(R_n),$$

and hence the expected distortion averaged over the client

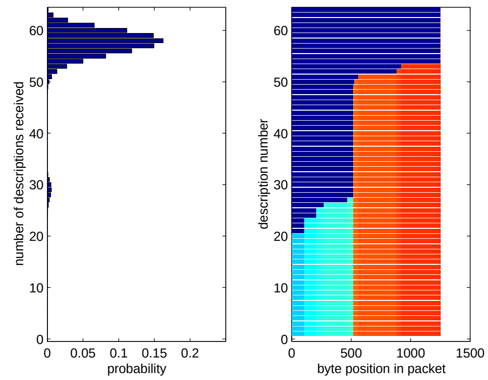


**Fig. 6**. Packetization matched to mixture of (mostly) high- and (some) low-bandwidth clients.

population (with respect to distribution $\nu$) is

$$D(\mathbf{R}) = \int D_\theta(\mathbf{R}) d\nu(\theta) = \sum_{n=0}^{N} p_n D(R_n),$$

where $p_n = \int p_{\theta,n} d\nu(\theta)$ is the probability that exactly $n$ out of $N$ descriptions are received after transmission to a client chosen randomly according to $\nu$. In a manner of speaking, the PET packetization that minimizes the expected distortion averaged over the client population is the PET packetization that minimizes the expected distortion for the average client. Thus, in a client population where fraction $\beta$ of the clients are high-bandwidth, the PET packetization that minimizes the average expected distortion is optimized for the distribution $\mathbf{p} = (1-\beta)\mathbf{p}_1 + \beta\mathbf{p}_2$, where

$$p_{1,n} = \binom{N_1}{n}(1-\epsilon_1)^n \epsilon_1^{(N_1-n)}$$

for $n = 0, 1, \ldots, N_1$, $p_{1,n} = 0$ for $n = N_1+1, \ldots, N_1 + N_2$, and

$$p_{2,n} = \binom{N_1+N_2}{n}(1-\epsilon_2)^n \epsilon_2^{(N_1+N_2-n)}$$

for $n = 0, 1, \ldots, N_1 + N_2$. The extremes, $\beta = 1$ and $\beta = 0$, result in the PET packetizations shown in Figures 4 and 5, respectively. Figure 6 shows the PET packetization optimized for a mixture of low- and high-bandwidth clients with $\beta = 97.5\%$. Note that only a very small fraction of low-bandwidth clients is required to radically alter the optimal packetization.

Unfortunately, this method of constructing a layered MDC by splitting a single MDC optimized for a mixture of low- and high-bandwidth clients does not offer a good tradeoff between the distortions seen by the different clients. In the next two subsections, therefore, we investigate two additional methods. Both methods use as the MDC base layer the set of $N_1$ packets optimized for the low-bandwidth client; hence they achieve the minimum possible distortion for the low-bandwidth client. The two methods then optimize the $N_2$ packets in the MDC enhancement layer to minimize the distortion for the high-bandwidth client, using two different constructions.

### 3.2. Layered MDC by Unequal Erasure Protection

This first construction is inspired by the use of unequal erasure protection for layered codecs, in which the base layer is typically more heavily protected than the enhancement layer, since for these codecs the enhancement layer is useless without the base layer. Thus, in this construction, herein called *layered MDC by unequal erasure protection*, some number $q$ of the $N_2$ packets in the MDC enhancement layer are allocated as additional parity packets to protect the MDC base layer, while the remaining $N_2' = N_2 - q$ packets in the MDC enhancement layer remain as PET packets for new source data not already present in the MDC base layer. (When $q = N_2$ this is the special case already seen in Figure 5 of the previous section.) In this way, the first $N_1' = N_1 + q$ packets contain source data up through source byte $R_{N_1}$, where $0 = R_0 \leq R_1 \leq \cdots \leq R_{N_1}$ are the breakpoints in the MDC base layer. Receiving any $n$ of these packets is sufficient to recover up through source byte $R_n$. Furthermore, the last $N_2'$ packets contain source data between source bytes $R_{N_1}$ and $R'_{N_2'}$, where $R_{N_1} = R_0' \leq R_1' \leq \cdots \leq R'_{N_2'}$ are the breakpoints in the MDC enhancement layer. Receiving any $n$ of these packets is sufficient to recover up through source byte $R_n'$ *provided* $R_{N_1} = R_0'$ *source bytes have already been recovered from the first* $N_1'$ *packets*; otherwise the $n$ packets received out of the last $N_2'$ are useless.

To see how to optimize the breakpoints $R_0' \leq R_1' \leq \cdots \leq R'_{N_2'}$, we derive an expression for the expected distortion seen by the high-bandwidth client as a function of these breakpoints. Let the probability that the high-bandwidth client receives exactly $n$ of the first $N_1'$ packets be $p'_{1,n} = \binom{N_1'}{n}(1-\epsilon_2)^n \epsilon_2^{(N_1'-n)}$, and similarly let the probability that the high-bandwidth client receives exactly $n$ of the last $N_2'$ packets be $p'_{2,n} = \binom{N_2'}{n}(1-\epsilon_2)^n \epsilon_2^{(N_2'-n)}$. Let $k$ be the minimum number of the first $N_1'$ packets necessary to recover up through source byte $R_{N_1}$. Then the expected distortion seen by the high-bandwidth client can be expressed

$$D(\mathbf{R}') = \sum_{n=0}^{k-1} p'_{1,n} D(R_n) + \left( \sum_{n=k}^{N_1'} p'_{1,n} \right) \left( \sum_{n=0}^{N_2'} p'_{2,n} D(R_n') \right).$$

Clearly, the last factor, and hence the overall distortion, can be minimized by the usual PET optimization procedures.

Figure 7 shows on the left the $N_1 = 32$ packets in the MDC base layer (optimized for $\epsilon_1 = 10\%$ packet loss) plus the $q = 2$ additional parity packets taken from the $N_2 = 32$ packets in the MDC enhancement layer, and shows on the right the remaining $N_2' = 30$ PET packets containing source data beyond source byte $R_{N_1}$ (optimized for $\epsilon_2 = 10\%$ packet loss).

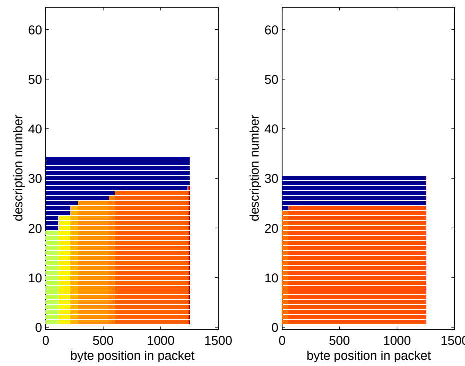


**Fig. 7**. Typical MDC base and enhancement layer packetizations using unequal erasure protection.

### 3.3. Layered MDC by Overlapping Layers

The next construction, denoted *layered MDC by overlapping layers*, devotes all $N_2$ packets in the MDC enhancement layer to PET packetization of the source data beyond source byte $R_\ell$, where $\ell$ determines the range of source bytes, $R_\ell$ through $R_{N_1}$, that are contained in both base and enhancement descriptions. The idea here is to repeat some of the less protected bytes from the base layer — viz., bytes $R_\ell$ through $R_{N_1}$ — in the enhancement layer, since the inability to recover these bytes would render the new (and possibly more heavily protected) bytes contained in the enhancement layer useless. By adjusting the breakpoints $R_0'$, $R_1', \ldots, R'_{N_2}$, it is possible to minimize the expected distortion seen by the high-bandwidth client. Let $k \leq \ell$ be the minimum number of the first $N_1$ packets necessary to recover up through source byte $R_\ell$. Then the expected distortion can be expressed

$$\begin{aligned} D(\mathbf{R}') &= \sum_{m=0}^{N_1} p_{1,m} \sum_{n=0}^{N_2} p_{2,n} d_m(R_n') \\ &= \sum_{n=0}^{N_2} p_{2,n} D_{mod}(R_n'), \end{aligned} \tag{1}$$

where $d_m(R') = D(R_m)$ (a constant in $R'$) if $m < k$, $d_m(R') = \min\{D(R_m), D(R')\}$ otherwise, and $D_{mod}(R')$ $= \sum_{m=0}^{N_1} p_{1,m} d_m(R')$ is a modified distortion measure. Hence using the modified distortion measure, the enhancement layer breakpoints $\mathbf{R}'$ can be optimized by minimizing (1) subject to $R_\ell = R_0' \leq \cdots \leq R'_{N_2}$ in the usual way.

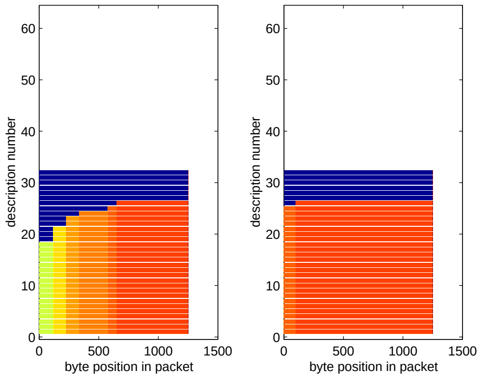


**Fig. 8**. Typical MDC base and enhancement layer packetizations using overlapping layers.

Figure 8 shows on the left the $N_1 = 32$ MDC base layer packets optimized for $\epsilon_1 = 10\%$ packet loss and on the right the $N_2 = 32$ MDC enhancement layer packets optimized for $\epsilon_2 = 10\%$ packet loss.

We note here that it is trivial to combine the last two constructions by substituting $N_1'$, $N_2'$, $p_{1,m}'$, and $p_{2,n}'$ for $N_1$, $N_2$, $p_{1,m}$, and $p_{2,n}$ in the above equations.

## 4. RESULTS

As in the figures in the previous sections, for the results in this section we have chosen packet loss rates $\epsilon_1 = \epsilon_2 = 10\%$, number of descriptions $N_1 = N_2 = 32$, target rate $R^* = 1250$ bytes per packet, and an operational distortion-rate function $D(R)$ obtained by encoding the first second of the standard MPEG test sequence *foreman* using a fine-grain scalable video codec. For these parameters, the bit rate to the low-bandwidth client is 320 Kbps and the bit rate to the high-bandwidth client is 640 Kbps for a GOF duration equal to 1 second. This might be typical for a scenario where the low-bandwidth clients are DSL subscribers with at most 384 Kbps downlink, and the high-bandwidth clients are DSL subscribers with at most 768 Kbps downlink.

Figure 9 shows the trade-off between the distortion seen by the low-bandwidth client (on the horizontal axis) and the distortion seen by the high-bandwidth client (on the vertical axis), for the various constructions. The star ($*$) in the lower left corner of the graph indicates the lowest possible distortion achievable at the low- and high-bandwidth clients, using separate multiple description codes as shown in Figures 3 and 4. The $\times$-marked solid curve indicates the performance of a single multiple description code optimized for the mixture of the low- and high-bandwidth channels, $\mathbf{p} = (1-\beta)\mathbf{p}_1 + \beta\mathbf{p}_2$, as $\beta$ runs from $0$ to $1$. When $\beta = 0$, the MDC is optimized for the low-bandwidth client and the high-bandwidth client suffers high distortion (upper left corner), while when $\beta = 1$, the MDC is optimized

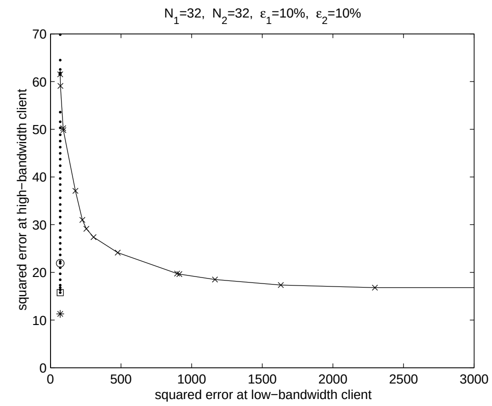


**Fig. 9**. Expected distortion of low-bandwidth client vs. expected distortion of high-bandwidth client for different layered MDC schemes, with 10% packet loss. $N_1 = N_2 = 32$.

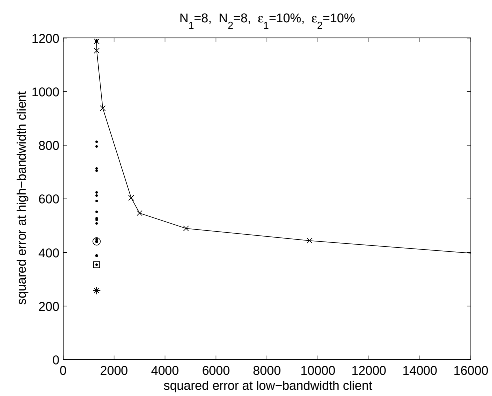


**Fig. 10**. Expected distortion of low-bandwidth client vs. expected distortion of high-bandwidth client for different layered MDC schemes, with 10% packet loss. $N_1 = N_2 = 8$.

for the high-bandwidth client and the low-bandwidth client suffers high distortion (lower right corner). The dots represent the performances of *layered* multiple description codes for all possible combinations of the number of packets $q = 0, 1, \ldots, N_2 - 1$ in the MDC enhancement layer used to protect base source data (bytes $R_0$ through $R_{N_1}$), and the amount $\ell = 0, 1, \ldots, N_1$ of overlap (bytes $R_\ell$ through $R_{N_1}$) covered in both the MDC enhancement and base layers. For every combination, the distortion seen by the low-bandwidth client is at its optimum, by construction. The circled dot represents the best performance when the amount of extra protection is minimal ($q = 0$), and the boxed dot represents the best performance when the amount of overlap is minimal ($\ell = N_1$). The best overall performance is apparently achieved when the amount of overlap is minimal. The best overall performance is still 1.4 dB away from the minimum possible distortion for the high-bandwidth client. Figure 10 shows similar results for the case $N_1 = N_2 = 8$.

| Schemes | Redundancy |
|---|---|
| High BW Optimal | 1.26 |
| Low BW Optimal | 2.66 |
| Mixed BW Optimal | 1.58 |
| Unequal Erasure Code | 1.30 |
| Layer Overlapping | 1.28 |

**Table 1**. Redundancy in various layered MDCs, given the same transmission rate.

Another measure of performance is redundancy, which is the ratio of the total number of source plus parity bytes ($(N_1 + N_2) \times R^*$) to the number of source bytes ($R'_{N_2}$) sent to the high-bandwidth client, per GOF. Table 1 shows the redundancies for each of the constructions for the case $N_1 = N_2 = 32$.

## 5. CONCLUSION

We have presented and evaluated constructions for two-layer multiple description codes. We have found that the best performing such codes can offer the low-bandwidth clients a distortion 0 dB worse than their minimum possible distortion while offering the high-bandwidth clients a distortion 1.4 dB worse than their minimum possible distortion. In the future it would be interesting to investigate whether there might be a trade-off such that, for example, layered multiple description codes pay no more than 1 dB penalty for each of the low- and high-bandwidth clients relative to separate multiple description codes.

The gap between the "optimal" distortion (represented by the star ($*$) in the lower left corner of Figures 9 and 10) and that attained by our layered multiple description codes might at first glance suggest that we would be better off with non-layered multiple description codes optimized separately for low-bandwidth and high-bandwidth clients. However, this is not necessarily so. A layered approach enables clients to share the base layer (and possibly other layers), thereby reducing bandwidth requirements on shared bottleneck links. Furthermore, adding or dropping layers is a less disruptive and more seamless way of adapting to changes in network congestion than switching between separate (non-layered) low-bandwidth and high-bandwidth channels. This is especially so when we are in a position to control the relative loss rates experienced by the different layers. In [12], we consider one such scenario — peer-to-peer multicast — and outline an approach for congestion control based on layered MDC. In summary, when the larger issues pertaining to the network and to content distribution are considered, layering offers many benefits.

## 6. REFERENCES


[1] V. K. Goyal. Multiple description coding: Compression meets the network. *IEEE Signal Processing Magazine*, pages 74–93, September 2001.

[2] P. A. Chou and K. Ramchandran. Clustering source/channel rate allocations for receiver-driven multicast with error control under a limited number of streams. In *Proc. Int'l Conf. Multimedia and Exhibition*, New York, NY, July 2000. IEEE.

[3] V. N. Padmanabhan, H. J. Wang, P. A. Chou, and K. Sripanidkulchai. Distributing streaming media content using cooperative networking. In *Proc. Network and Operating System Support for Digital Audio and Video (NOSSDAV)*, Miami, FL, May 2002.

[4] S. R. McCanne. *Scalable Compression and Transmission of Internet Multicast Video*. PhD thesis, The University of California, Berkeley, CA, December 1996.

[5] P. A. Chou, A. E. Mohr, A. Wang, and S. Mehrotra. Error control for receiver-driven layered multicast of audio and video. *IEEE Trans. Multimedia*, 3(1):108–122, March 2001.

[6] A. Albanese, J. Blömer, J. Edmonds, M. Luby, and M. Sudan. Priority encoding transmission. *IEEE Trans. Information Theory*, 42:1737–1744, November 1996.

[7] G. Davis and J. Danskin. Joint source and channel coding for image transmission over lossy packet networks. In *Conf. Wavelet Applications to Digital Image Processing*, Denver, CO, August 1996. SPIE.

[8] A. E. Mohr, E. A. Riskin, and R. E. Ladner. Unequal loss protection: Graceful degradation of image quality over packet erasure channels through forward error correction. *IEEE J. Selected Areas in Communications*, 18(6):819–829, June 2000.

[9] R. Puri and K. Ramchandran. Multiple description source coding through forward error correction codes. In *Proc. Asilomar Conference on Signals, Systems, and Computers*, Asilomar, CA, October 1999. IEEE.

[10] A. E. Mohr, R. E. Ladner, and E. A. Riskin. Approximately optimal assignment for unequal loss protection. In *Proc. Int'l Conf. Image Processing*, Vancouver, BC, September 2000. IEEE.

[11] V. Stanković, R. Hamzaoui, and Z. Xiong. Packet loss protection of embedded data with fast local search. In *Proc. Int'l Conf. Image Processing*, Rochester, NY, September 2002. IEEE.

[12] P. A. Chou, V. N. Padmanabhan, and H. J. Wang. Resilient peer-to-peer streaming. Technical Report MSR-TR-2003-11, Microsoft Research, Redmond, WA, March 2003.