To appear in *ICIP 2002, Sept 22-25, 2002.*

PERFORMANCE OF A MULTIPLE DESCRIPTION STREAMING MEDIA CONTENT DELIVERY NETWORK

John G. Apostolopoulos, Wai-tian Tan, Susie J. Wee

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

ABSTRACT

Content delivery networks (CDNs) have been widely used to provide reduced delay and packet loss, fault tolerance, and improved scalability for web content delivery. Additional benefits are provided for video streaming when one designs a Streaming Media CDN (SM-CDN) for either conventional single description (SD) or multiple description (MD) coding. Specifically, when precise network conditions and topology are known, simulations show that an MD-SM-CDN can provide 20 to 40% reduction in distortion over a conventional SD-SM-CDN, even when the underlying CDN is not designed with MD streaming in mind [1].

This paper examines the performance of an MD-SM-CDN as a function of different network topologies and loss conditions, and compares it with a conventional SD-SM-CDN. This examination provides insight into an MD-SM-CDN's performance when knowledge of network topology and conditions is imprecise or uncertain. Our simulations indicate that an MD-SM-CDN can provide improved performance over a conventional SD-SM-CDN over a wide range of network topologies and loss conditions.

1. INTRODUCTION

Content delivery networks (CDN) were developed to overcome performance problems, such as network congestion and server overload, that arise when many users access popular content. CDNs improve end-user performance by caching popular content on edge servers located closer to users [2]. This provides a number of advantages. First, it helps prevent server overload, since replicated content can be delivered to users from edge servers. Furthermore, since content is delivered from the closest edge server and not from the origin server, the content is sent over a shorter network path, thus reducing the request response time, the probability of packet loss, and the total network resource usage. While CDNs were originally intended for static web content, the same benefits have been exploited by state-of-the-art streaming video systems.

Video communication quality over lossy packet networks can be improved by employing multiple description coding (MDC) and path diversity [3]. The design of a multiple description streaming media CDN (MD-SM-CDN) can create an effective platform for providing and exploiting path diversity. Specifically, we use MDC to code a media stream into multiple descriptions, and distribute copies of these descriptions across servers in the CDN. When a client requests a media stream, it is directed to multiple nearby servers which host complementary descriptions of the stream. The client receives the complementary descriptions through different paths and from different servers. In this way, midstream disruption occurs only in the less likely case when losses afflict both descriptions simultaneously.

Realizing an effective MD coding and path diversity scheme in a CDN context requires solving a number of fundamental CDN design problems adapted to the context of an MD-SM-CDN: (1) where to deploy the servers (server placement), (2) how to distribute the descriptions across the servers (content distribution), and (3) how to select for each client the best servers hosting complimentary descriptions (server selection). Assuming precise knowledge of network condition and topology, the models developed in [4] describe the expected distortion for MD and path diversity and can be used to determine high-quality, and in some cases optimal, solutions for these MD-SM-CDN design problems. However, various attributes of the network topology and loss conditions may be unknown or time-varying. This uncertainty can have a detrimental effect on the solution of these MD-SM-CDN problems and may influence whether an MD-SM-CDN or a conventional SD-SM-CDN provides better performance. Thus, the goal of this paper is to (1) examine how the client performance in an MD-SM-CDN varies as a function of different network topologies and loss conditions, (2) determine under what conditions an MD-SM-CDN or a conventional SD-SM-CDN provides better performance, and (3) use this knowledge to gain an understanding of the effect of network uncertainty on MD-SM-CDN performance.

This paper continues with a brief overview of MD and path diversity in Sec. 2 and the design of an MD-SM-CDN in Sec. 3. Models for estimating MD and SD distortion are discussed in Sec. 4, and the performance as seen by the client is examined in Sec. 5.

2. MULTIPLE DESCRIPTION VIDEO CODING AND PATH DIVERSITY

Video communication over lossy packet networks such as the Internet is hampered by limited bandwidth and packet loss. MD coding is one form of compression that may help to overcome this problem [5, 6, 7]. In MDC a signal is coded into a number of separate bitstreams, each containing a complementary description of the signal. MD coding enables a useful reproduction of the signal when *any* description is received; for this reason it is beneficial to increase the probability that at least one description is received correctly at any point in time. MDC provides improved reconstructed video quality as more descriptions are received.

In [3] it was shown that combining MD video coding with a path diversity transmission system, where different descriptions are explicitly transmitted through different network paths (as opposed to the default scenarios where they would proceed along a single path), can improve the effectiveness of MD coding over a packet network by increasing the likelihood that the loss probabilities for each description are independent. A path diversity communication system enables the end-to-end application to effectively see a virtual channel with improved loss characteristics [8, 3]. For example, the application effectively sees an average path behavior, which generally provides better performance, e.g. reduced variability, as compared to a single path. Furthermore, the probability that all of the multiple paths are simultaneously congested is much less than the probability that a single path is congested. An important property of our selected MD video coder [7, 3], as opposed to other proposed MD video coders, is that it enables repair of corrupted frames in a description using uncorrupted frames in other descriptions so that usable quality can be maintained even when there are losses in all descriptions, as long as the losses do not afflict all descriptions simultaneously. This MD video coder can be designed as a standard-compatible enhancement of MPEG-4 V2, H.263 V2, and H.26L.

3. MULTIPLE DESCRIPTION STREAMING MEDIA CONTENT DELIVERY NETWORK

The previous section described how the combination of MD coding and path diversity can achieve improved error-resilient streaming over lossy packet networks. In [1], we show how the benefits of MD coding and path diversity can be achieved in a CDN infrastructure. Specifically, media is coded into multiple descriptions which are distributed across the servers of a CDN. Then, when a client requests a stream, it is directed to multiple nearby servers hosting complementary descriptions. These descriptions are sent to the client from different servers over different paths. The paths between the distributed servers and the client may have different degrees of jointness or disjointness; greater degrees of disjointness imply greater path diversity, which when coupled with MD coding can lead to improved error resilience for the MD-SM-CDN.

3.1. MD-SM-CDN Architecture Design

Traditional CDN design principles can be used to design SD-SM-CDN's. However, the use of MD coding in an MD-SM-CDN presents a number of important new research problems that are fundamentally different from their SD counterparts.

The **Server Placement Problem** addresses how to determine the best locations to deploy servers to optimize client performance. Conventional SD-SM-CDN design would place servers to minimize the distance between clients and the closest single server. However, in MD-SM-CDN design it is important to minimize the distance from the client to multiple servers which host complementary descriptions of the content. Furthermore, it is important to maximize path disjointness between the client and these multiple servers to exploit the path diversity gain achievable with MD coding. Thus, the MD server placement problem requires optimization for both server distance and path diversity.

Given an existing CDN infrastructure, the **Content Distribution Across Servers Problem** addresses how to optimally distribute the MD streams across an existing set of servers. In a conventional SD-SM-CDN, the content is stored at each single chosen server. In an MD-SM-CDN the content is coded into complementary descriptions which are distributed across multiple servers. In this system, each server may have zero, one, or both descriptions. The MD content distribution problem involves optimally distributing the descriptions across the servers such that both descriptions are close to every client while maximizing path diversity.

For each client request, the **Server Selection Problem** addresses how to select the optimal pair of servers that have complementary descriptions to maximize the quality at the client. In a conventional SD-SM-CDN, each client request is typically served by the single closest server. However, for an MD-SM-CDN the goal is to select multiple servers with complementary descriptions, and both server distance and path diversity must be considered to optimize performance.

In certain cases it may be beneficial for an MD-CDN to leverage an existing CDN infrastructure, instead of deploying servers from scratch; since the server placement is predetermined, only the content distribution and server selection problems must be solved.

4. MODELING MD CODING AND PATH DIVERSITY

In order to better understand, characterize, and compare MD coding and path diversity versus conventional approaches, analytical models that accurately and quickly predict the performance of the system for different path diversity topologies and network conditions were developed [4]. Specifically, [4] presented models that express (1) MD video distortion behavior as a function of different loss events and (2) the loss process for a two-path path diversity transmission system. In addition, models were presented for conventional single description video coders (e.g. MPEG-4) and the loss process for a single path. These models enable one to compare the performance of MD video coding and path diversity against conventional SD video over a single path for different network topologies and channels exhibiting isolated and bursty packet loss. It was shown that MD video and path diversity provides significant benefits for a variety of settings that exhibit path diversity.

To summarize the MD and path diversity models in [4], the MD and SD distortion behavior for different loss events can be captured for a given video sequence by 5 and 7 parameters, respectively. Each path is modeled as the concatenation of multiple independent links, where the loss characteristics of each link is modeled with a Gilbert model parameterized by $\{p, q\}$ which can reflect different packet loss rates and burst lengths, and the effects of joint and disjoint links are captured. This provides (1) a model for the loss process of a two-path path diversity system, and (2) a distortion model that maps the loss model to MD distortion values. Specifically, assuming identical Gilbert parameters $\{p_0, q_0\}$ for each link, SD distortion for a given video source over a path of N_{Single} links is given by $D_{SD}(N_{Single}, p_0, q_0)$, and that for MD is given by $D_{MD}(N_{Disjoint\text{-}1}, N_{Joint}, N_{Disjoint\text{-}2}, p_0, q_0)$, where $N_{Disjoint\text{-}1}, N_{Joint}$, and $N_{Disjoint\text{-}2}$ denote the number of joint and disjoint links along the two paths.

5. EXAMINING PERFORMANCE

Given knowledge of the network topology, including knowledge of joint and disjoint links, and loss characteristics for each link, the analytical models presented in the prior section enable us to accurately estimate expected distortion and thereby enable the reliable solution of the problems presented in Section 3. However, typically one does not have complete knowledge of the network topology and loss conditions, and therefore it is important to characterize how robust the solutions are to partial or noisy estimates of this information. This is particularly important for an MD-SM-CDN, since it requires additional information as compared to a conventional SD-SM-CDN (e.g. knowledge of the number of joint and disjoint links) and therefore may be more vulnerable to uncertainties in this information. This section examines the variation in the expected distortion at the client for different network topologies and loss conditions.

The MD and path diversity models depend on accurate knowledge of path lengths, number of joint and disjoint links, and loss characteristics of each link. In many cases it may be possible to estimate the total number of links between a client and each server of a pair of servers. For example, the *ping* utility may provide this information in certain contexts. However, as illustrated in Figure 1, it may be difficult to determine the number of links that are joint and the number that are disjoint. Therefore an important consideration is the effect of inaccuracies in estimating the number of joint links, given that the total number of links for each path is known. These inaccuracies can be better understood by examining the performance under different loss conditions (packet loss rate and burst length). In the following, we assume two MD streams are sent over two paths, and the path diversity is expressed by the triplet $\{N_{\text{Disjoint-1}}, N_{\text{Joint}}, N_{\text{Disjoint-2}}\}$. The SD path is chosen as the shorter of the two MD paths, $N_{\text{Single}} = N_{\text{Joint}} + min(N_{\text{Disjoint-1}}, N_{\text{Disjoint-2}})$. We also assume each link is IID with identical Gilbert parameters $\{p_0, q_0\}$. The results for the *Foreman* and *Bus* sequences are plotted on the left and right, respectively. Details about the MD video coder are available in [4].

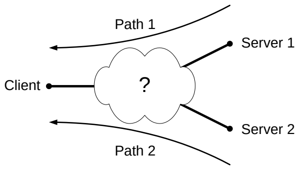

Fig. 1. A client communicates with two servers over two paths, however neither the client nor the MD-SM-CDN may know the path topology (joint and disjoint links) and loss characteristics.

5.1. Experimental Results

Performance as a Function of Joint versus Disjoint Losses Figure 2 examines MD and SD performance for 5% end-to-end average packet loss rate, as we vary the fraction of losses that are joint versus disjoint. Another interpretation is that the total length of each path is known to be 10 links, however it is not known how many links are joint: the number of joint links can vary from 0 (completely disjoint) to 10 (completely joint) and the number of disjoint links therefore varies from 10 to 0. We assume $\{p_0, q_0\} = \{.0043, .835\}$ for each link, which corresponds to 5% end-to-end average packet loss rate for 10 links, and average burst length of 1.25 packets which is the longest average burst length (for 30 msec sampling) that we are aware of in the literature [9, 10].

In these plots, the SD performance is constant because SD is sent along a single path that is independent of jointness. For maximally disjoint paths, the MD distortion is minimized and it significantly outperforms SD. As the jointness of the paths increases, the MD distortion increases approximately linearly to the case of completely joint paths at which point its performance is similar to that of SD. Specifically, in the Foreman sequence, MD outperforms SD across all jointness ratios; while in the Bus sequence, MD outperforms SD for jointness ratios between 0% and 90%, with a crossover point at around 94%. In general, we observe that the MD distortion is sensitive to the ratio of joint to disjoint links, contributing to a distortion variation of a factor of two in our test conditions. Furthermore, we see that MD comfortably outperforms SD for most jointness ratios. Thus, we conclude that when coupled with MD coding, even small amounts of path diversity or disjointness can improve the performance of an SM-CDN.

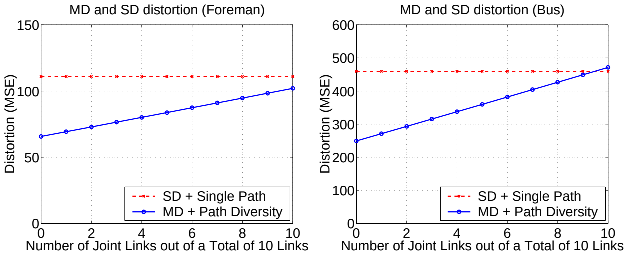

Fig. 2. MD and SD distortion as a function of joint versus disjoint losses. When 8 out of the 10 links in each path are joint, then 80% of the end-to-end losses are joint losses: 4% and 1% end-to-end packet loss on joint and disjoint links, respectively.

Performance as a Function of Packet Loss Rate, Expected Burst Length, and Topology The packet loss rate and expected burst length are generally unknown in advance and usually are time-varying. Therefore, it is important to characterize the effect of different packet loss rates and burst lengths on MD and SD performance. Figure 3 examines the performance for different packet loss rates for three different single-link expected burst lengths of {1.11, 1.25, 1.43} packets corresponding to q_0 = {.9, .8, .7}, and three different path topologies. As anticipated, the expected distortion for SD and MD increases with the packet loss rate, and for MD the slope increases with the path jointness [3]. An important observation is that the MD distortion is sensitive to losses in joint links, but is largely insensitive to losses in disjoint links.

Furthermore, the plots show that MD outperforms SD for all packet loss rates and burst lengths for the completely disjoint and half disjoint topologies. Meanwhile, for the completely joint topology SD only begins to outperforms MD for longer burst lengths. This is shown more clearly in the plots in Figure 4. These plots examine the effect of expected single-link burst lengths ranging from 1 to 1.5 packets corresponding to q_0 varying from 1 to .667. p_0 is chosen as .0235 which corresponds to 2.3 to 3.4 % packet loss for a single link (depending on q_0) and 4.5 to 6.7 % end-to-end packet loss over two links. Note that the end-to-end loss characteristics depend on whether the path length is 1 or 2 links. Once again, the characteristics vary as a function of the type of path diversity that exists, and the combination of joint links and large expected burst length corresponds to an increased probability of simultaneous loss of both descriptions for MD, corresponding to significantly higher distortion. It is worth stressing that the longest expected end-to-end burst length that has been experimentally measured and cited in the literature [9, 10] that we are aware of is 1.25 packets, and in the region of expected burst length less than 1.25 packets, the expected distortion of MD is less than that of SD.

SD distortion depends primarily on the loss rate and to a lesser extent on the expected burst length (relatively small spread in Fig 3 and small slope in Fig 4). In contrast, MD is relatively insensitive to loss rate and loss burstiness on disjoint links (small spread in Fig 3 and small slope in Fig 4 for topology {1,0,1}), but is sensitive to loss rate and loss burstiness on the joint link (larger spread in Fig 3 and larger slope in Fig 4 for topology {0,1,0}). This is because

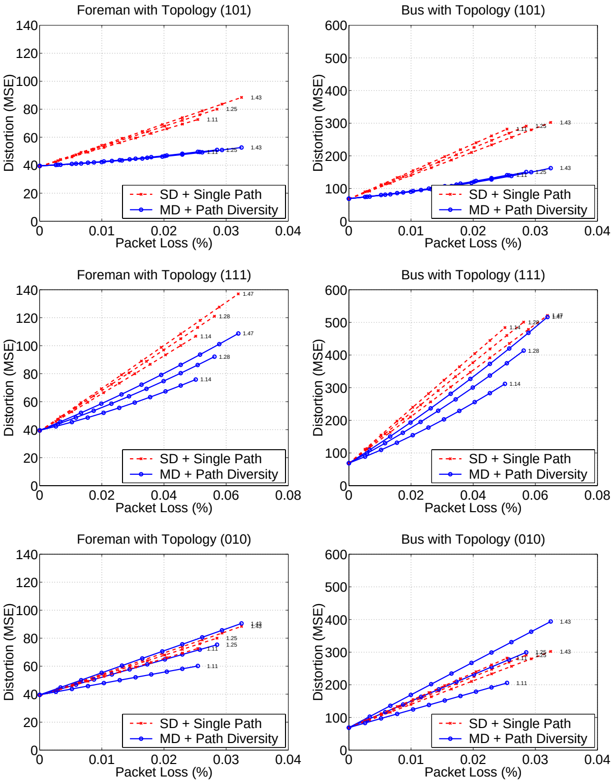

Fig. 3. MD vs SD distortion as packet loss rate is varied for three different expected burst lengths and three different degrees of path diversity: (top) completely disjoint {1,0,1}, (middle) half disjoint and half joint {1,1,1}, and (bottom) completely joint {0,1,0}.

the dominant term for MD distortion results from simultaneously losing both descriptions, which is more likely when losses in the joint link are heavy or bursty. Also note that the total MD distortion for a topology given by the concatenation of disjoint links and joint links is approximately additive, specifically the distortion for topology {1,1,1} is approximately given by the sum of the baseline distortion for the case of no-loss and the incremental distortions for {1,0,1} and {0,1,0}.

6. CONCLUDING REMARKS

This paper provided an overview of design issues for multiple description streaming media CDNs. MD-SM-CDN performance was examined in terms of expected distortion at the client as a function of different network topologies (different degrees of path diversity) and different network loss conditions (packet loss rate and expected burst length). Our studies show that for topologies with even modest amounts of path diversity, MD-SM-CDNs outperform SD-SM-CDNs over a wide range of typical operating conditions, and that the improvement becomes quite significant for larger degrees of path diversity. SD-SM-CDNs perform better for the case of completely joint paths with long expected burst lengths, since in this extreme case the MD coder is highly likely to be afflicted with simultaneous losses in both descriptions. However, in most typical operating ranges MD-SM-CDNs provide better performance.

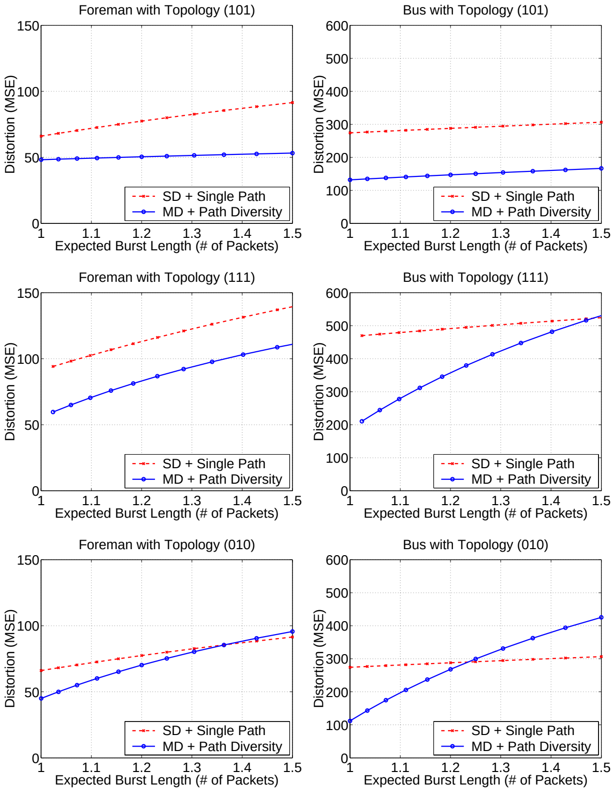

Fig. 4. MD vs SD distortion as the expected single-link burst length varies from 1 to 1.5 packets, for three different degrees of path diversity: (top) completely disjoint {1,0,1}, (middle) half disjoint and half joint {1,1,1}, and (bottom) completely joint {0,1,0}.

7. REFERENCES

[1] J.G. Apostolopoulos, T. Wong, W. Tan, and S.J. Wee, "Multiple description streaming media content delivery networks," *INFOCOM*, June 2002.

[2] L. Qiu, V. Padmanabhan, and G. Voelker, "On the placement of web server replicas," *INFOCOM*, 2001.

[3] J.G. Apostolopoulos, "Reliable video communication over lossy packet networks using multiple state encoding and path diversity," *VCIP*, January 2001.

[4] J.G. Apostolopoulos, W. Tan, S.J. Wee, and G.W. Wornell, "Modeling path diversity for multiple description video communication," *ICASSP*, May 2002.

[5] V. Vaishampayan and S. John, "Interframe balanced-multiple-description video compression," *ICIP*, Oct. 1999.

[6] A.R. Reibman, H. Jafarkhani, Y. Wang, M.T. Orchard, and R. Puri, "Multiple description coding for video using motion compensated prediction," *IEEE Inter. Conf. Image Processing*, October 1999.

[7] J.G. Apostolopoulos, "Error-resilient video compression via multiple state streams," *Proc. International Workshop on Very Low Bitrate Video Coding (VLBV'99)*, pp. 168–171, October 1999.

[8] N.F. Maxemchuk, *Dispersity Routing in Store and Forward Networks*, Ph.D. thesis, University of Pennsylvania, May 1975.

[9] M. Yajnik, S. Moon, J. Kurose, and D. Towsley, "Measurement and modelling of temporal dependence in packet loss," *INFOCOM'99*.

[10] J. Wenyu and H. Schulzrinne, "Modeling of packet loss and delay and their effects on real-time multimedia service quality," *NOSSDAV'00*.