



## Path Diversity for Enhanced Media Streaming

John G. Apostolopoulos, Mitchell D. Trott  
Mobile and Media Systems Laboratory  
HP Laboratories Palo Alto  
HPL-2004-123  
July 14, 2004\*

video streaming,  
path diversity,  
multipath  
streaming, multiple  
description coding

Media streaming over best-effort packet networks such as the Internet is quite challenging because of the dynamic and unpredictable available bandwidth, loss rate, and delay. Recently, streaming over multiple paths to provide path diversity has emerged as an approach to help overcome these problems. This paper provides an overview of the benefits and use of path diversity for media streaming. The different approaches for media coding and streaming over multiple paths are examined, together with architectures for achieving path diversity between single or multiple senders and a single receiver. Important examples include using the distributed servers in a content delivery network to provide path diversity to a requesting client, using multiple 802.11 wireless access points to provide path diversity to a mobile client, and using relays to provide low-latency media communication. The design, analysis, and operation of media streaming systems that use path diversity are considered, with emphasis on the accurate performance models needed to select the best paths or best servers.

\* Internal Accession Date Only

Approved for External Publication

© Copyright IEEE. To be published in IEEE Communications Magazine, special issue on Proxy Support for Streaming on the Internet, August 2004.

## PATH DIVERSITY FOR ENHANCED MEDIA STREAMING

John G. Apostolopoulos and Mitchell D. Trott

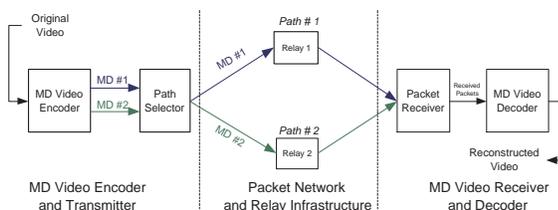
Streaming Media Systems Group  
Hewlett-Packard Labs, Palo Alto, CA 94304

### ABSTRACT

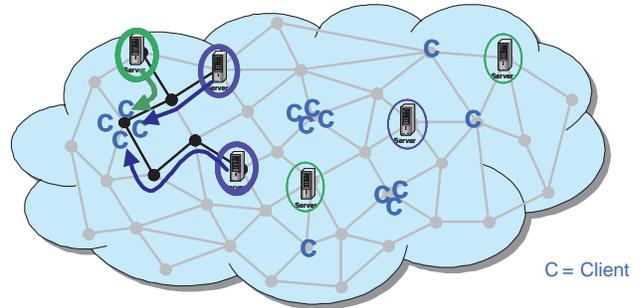
Media streaming over best-effort packet networks such as the Internet is quite challenging because of the dynamic and unpredictable available bandwidth, loss rate, and delay. Recently, streaming over multiple paths to provide path diversity has emerged as an approach to help overcome these problems. This paper provides an overview of the benefits and use of path diversity for media streaming. The different approaches for media coding and streaming over multiple paths are examined, together with architectures for achieving path diversity between single or multiple senders and a single receiver. Important examples include using the distributed servers in a content delivery network to provide path diversity to a requesting client, using multiple 802.11 wireless access points to provide path diversity to a mobile client, and using relays to provide low-latency media communication. The design, analysis, and operation of media streaming systems that use path diversity are considered, with emphasis on the accurate performance models needed to select the best paths or best servers.

### 1. INTRODUCTION

*Path diversity* is a transmission technique that sends data simultaneously through two or more paths in a packet-based network. The paths may originate from single [1, 2, 3, 4, 5, 6] or multiple sources [7, 8, 9]. The single-source case in a typical video streaming application is shown in Figure 1 [1]. A more complex multiple-source scenario is illustrated in Figure 2 [8].



**Fig. 1.** Path diversity: media packets are sent over multiple paths in a packet network. In this case different media descriptions are explicitly sent over different paths via relays.



**Fig. 2.** Path diversity in a content delivery network: multiple servers stream content to each client. A basic question: one green and one blue server should be selected to stream to the clients—which blue server should be selected? The closer one which shares a link with the green server, or the further blue server?

The use of multiple paths through the transport network for streaming has been proposed to help overcome the loss [1] and delay [2] problems that afflict streaming media and low-latency communication. In addition, it has long been known that multiple paths can improve fault tolerance and link recovery for data delivery, as well as provide larger aggregate bandwidth, load balancing, and faster bulk data downloads.

Diversity techniques have been studied for many years for wireless communication, e.g., frequency, time, and spatial diversity. However, path diversity over packet networks has received limited attention until recently. The early work in this area, e.g. [10, 11], focuses on reliable data delivery through error correction coding or retransmission, and is generally performed in the context of multiple virtual circuits on connection-oriented systems such as ATM.

A number of thorough experimental studies have demonstrated the great variability in the end-to-end performance observed over the Internet. This variability is analogous to the variability of wireless links, and motivates the application of wireless diversity techniques to wired and wireless Internet communication. Further motivation for the use of multiple paths is that the default (single) path between two nodes on a network is often not the best path. For example, the measurement study [12] comparing the paths between

two hosts on the Internet found that “in 30-80% of the cases, there is an alternate path with significantly superior quality” to the default path, where quality is measured in terms of metrics such as round-trip-time, loss rate, and bandwidth.

This paper outlines the benefits, implementation questions, and open problems associated with using path diversity for streaming media delivery. We begin in Section 2 by describing two basic components of the path diversity systems we consider: media streaming and media coding. In Section 3 we survey the benefits of path diversity in a streaming-media context. Many of these benefits accrue only for media streaming applications and have no direct analogy, for example, to benefits associated with classical wireless diversity. In Section 4 we describe various architectures that support path diversity, including source routing, overlay networks, and IEEE 802.11 wireless infrastructure. Lastly, Section 5 considers a variety of research problems in the design, analysis, and operation of path diversity systems.

## 2. BUILDING BLOCKS FOR MEDIA STREAMING

This section describes the general problem of media streaming, focusing in particular on how it differs from transporting generic data. We then provide an overview of conventional single description, scalable, and multiple-description compression techniques, which are important ingredients in streaming methods that exploit multiple transport paths.

### 2.1. Media Streaming

The problem of streaming media such as voice and video over best-effort packet networks is complicated by a number of factors. Unlike static content delivery or file download, streaming involves the simultaneous delivery and playback of media and is characterized by delay constraints on each transmitted packet. Packets that arrive after their decoding or display deadline are useless. These inevitable packet losses are not necessarily fatal; most media streams have some tolerance to packet loss, albeit limited by the use of temporally predictive compression. Delay constraints and tolerance to loss are primary factors that distinguish media streaming from ordinary data transport.

If the delay constraints are loose relative to typical transport and decoding delays, for example more than a few seconds, then media can be delivered using mechanisms that resemble file download. Download is convenient for delivering pre-encoded objects of small size such as images or short audio clips.

Larger objects, such as video clips, require long download times and significant storage at the receiver. A better approach for delivering these objects is streaming, which involves splitting the object into parts, transmitting the parts

in succession, and decoding and playing back the parts as they are received.

Streaming is the only option for live (meaning not pre-recorded) media objects. Conversational or interactive applications have particularly tight delay constraints, typically 100 milliseconds or less. Meeting these constraints in the presence of packet losses, queuing delays, and network outages is quite challenging, and provides motivation for some of the path diversity techniques which follow.

Note that conventional approaches for overcoming packet loss and network outage in data delivery, such as reliable transport via TCP, generally are not possible for streaming because the persistent retransmissions lead to excessive delay. Moreover, many applications (e.g., multicast or broadcast) lack a back channel or other means for retransmissions.

### 2.2. Media Coding: Single Description, Scalable, and Multiple Description

Conventional single-description (SD) video coding algorithms, such as MPEG-1/2/4 and H.261/3/4, compress video into a single bitstream. In scalable (also called layered) coding, video is coded into multiple bitstreams, beginning with a base layer that provides low but usable quality, and one or more enhancement layers that improve quality. The scalable coding bitstreams are partially ordered, and different subsets can represent the video at different spatial, temporal, or amplitude resolutions and fidelities. Applications and networks that support prioritization can exploit scalable coding by assigning higher delivery priorities to coding layers that are more important. In best-effort networks like the Internet, however, all packets are equally likely to be lost or delayed. This fundamental mismatch—prioritized data on a non-prioritized network—makes scalable coding by itself difficult to exploit using the Internet.

Multiple description (MD) coding is an alternative to scalable coding that produces two or more bitstreams, or *descriptions*, each of roughly equal importance. MD coding has two important properties. First, each description can be independently decoded to give a usable reproduction of the original signal. Second, the multiple descriptions contain complementary information so that the quality of the decoded signal improves with the number of descriptions received. Note that the first property contrasts with scalable coding, which has a base layer that is critically important and if lost renders the other bitstream(s) useless. MD coding enables a useful reproduction of the signal when *any* description is received. For example, a simple MD video coder can be achieved by splitting a video into even and odd frames and coding them separately.

*Repairable* MD coding (also known as *state recovery* in MD coding) is an extension of conventional MD coding beneficial for temporally predictively coded media such as

video [1]. Predictive coding causes errors in one frame, due for example to lost packets, to propagate to potentially many subsequent frames. In repairable MD coding the decoder stops error propagation in one description by repairing it using undamaged frames from the other. Such error-resilient decoding is effective so long as packets are not simultaneously lost for both descriptions — thus coupling nicely with path diversity. This comparatively recent decoder-side optimization is complementary to the ongoing efforts to improve the rate vs. quality tradeoff of MD encoders, which typically do not consider partial losses of both descriptions.

The relative compression efficiencies of SD, scalable, and MD coding depend on the media type. For image and audio sources scalable coding can be as efficient as SD coding, and many new coding standards, e.g., JPEG-2000, are scalable. In addition, scalable image or audio coding may be combined with unequal error protection (UEP) to provide an efficient MD-like system: more redundancy is allocated to the high-priority data relative to the low-priority data, improving the probability that the high-priority data is correctly received. For video sources both scalable coding and MD coding are currently less efficient than SD coding, however they provide valuable properties for streaming.

### 3. BENEFITS OF PATH DIVERSITY

Streaming over multiple paths can provide several types of benefits depending on how the paths are used. In this section we briefly describe some of these benefits, as itemized in Table 1. For simplicity we assume ideal path diversity where each path is independent with identical attributes. The case of realistic paths is examined in Section 5.

Benefits of Path Diversity	
Improved (aggregate) BW	Merging
Load Balancing	Splitting
Reduced variability (losses)	Average
Reduced burst losses	Subsampling
Improved error resilience	OR
Reduced outages	AND
Reduced latency (queue diversity)	Min

**Table 1.** Benefits of path diversity and the associated mathematical operations that intuitively expresses why the benefit occurs.

A straightforward benefit of multiple paths is increased bandwidth available by using all paths at once. A complementary benefit is load balancing, i.e., decreasing per-path bandwidth by splitting a stream across multiple paths.

Another benefit of path diversity is reduced variability of packet losses, e.g., reduced excursions between periods of no loss and high loss that are common on the Internet.

The end-to-end application sees the average network behavior across the paths, which generally has reduced variability. This reduced variability enables methods for combating losses, such as forward error correction (FEC), to become more effective [7, 13].

FEC adds specialized inter-packet redundancy that enables data recovery up to a loss threshold. The amount of redundancy can be adjusted dynamically to compensate for changes in the loss rate, but adaptation is problematic when the network changes quickly: the FEC is inevitably over-designed (and therefore inefficient) or under-designed (and therefore ineffective). Combining FEC with long interleaving also helps combat loss variability, but the added delay makes interleaving unsatisfactory for media streaming. Path diversity may be viewed as a form of interleaving across paths rather than across time. Combined with FEC, path diversity provides benefits similar to time interleaving with smaller associated delay.

Path diversity also reduces the frequency and length of burst losses, i.e., losses of consecutive packets. Distributing packets across multiple paths increases the inter-packet spacing on each path, and therefore for a network congestion event of a given duration fewer packets are lost. Reducing burst losses provides a number of benefits for media streaming. For example, for video it is easier to recover from multiple isolated losses than from an equal number of consecutive losses [1]. For two paths with equal average packet loss rates, sending even packets on one path and odd packets on the other has no effect on the end-to-end loss rate but does reduce burst losses. This provides 0.5–1.0 dB gain in video quality in certain SD-coded situations [5].

Improved error resilience is a powerful, synergistic benefit that accrues from combining path diversity with repairable MD coding. To better understand this combination, observe that repairable MD coding is useful without path diversity because it reduces error propagation. Conversely, path diversity is useful with non-repairable MD coding because the distortion remains manageable even if one description is lost. But both of these “partial” approaches have weaknesses: in the former case burst losses on the single path simultaneously damage frames in both descriptions and defeat the error-resilient decoder; in the latter, even isolated errors on the two paths ruin both descriptions. The combination of repairable MD coding and path diversity allows the system to survive multiple losses—even burst losses—on the two paths, so long as the losses are not simultaneous.

Path diversity can dramatically decrease the probability of outage, where an *outage* is an event where communication along a network path is lost for a sizable length of time. A path diversity system experiences complete outage only when all paths are undergoing outages. A system employing MD video coding benefits from such outage protection because the media stream remains useful even if the traffic

on one path is lost (e.g., 50% packet loss).

Packets sent over independent paths are unlikely to suffer simultaneously from large packet delays. In heuristic terms, the probability of a bad delay event improves from  $p$  to  $p^2$  when the minimum of two path delays governs whether information is late. We term this effect *queue diversity* because network delays are often due to backlogged queues, and to emphasize the benefit of multiple parallel queues. Queue diversity offers little benefit for loss-intolerant data because the receiver must wait for all packets from all paths. In contrast, loss-tolerant delay-intolerant voice or video stream remain useful when only some packets arrive promptly. The affected stream may realize the minimum of two or more path delays (effectively chopping off the long tails in the end-to-end delay distribution), albeit at some loss in quality. Examples of such delay reduction include MD voice over IP (VoIP) [2] and SD video over 802.11 wireless networks [5].

## 4. ARCHITECTURES FOR PATH DIVERSITY

This section examines different architectures for media streaming with path diversity over packet networks. Specifically, we discuss: the different media coding approaches and how they can be coupled with path diversity; two approaches for path diversity from a single sender to a single receiver; the use of the distributed infrastructure of a CDN to stream from multiple senders to a single receiver; the use of multiple 802.11 access points for streaming to a wireless client; and the use of multiple paths in an ad-hoc wireless network.

### 4.1. SD versus Scalable versus MD coding?

Effective media streaming with path diversity requires a careful co-design of the compression and the path diversity components. The coding approaches reviewed in Section 2.2 provide a number of tradeoffs, and there are many open research questions as to when each is preferable. In general, the answer depends crucially on the context, such as the specific media type, source coder, the required playback delay, the use of retransmissions or FEC, transmission over best-effort or QoS-enabled network, and whether the transmission is unicast or multicast.

For example, while (prioritized) scalable coding is not a good match to the (unprioritized) best-effort Internet, the ability to perform prioritized retransmission or unequal error protection at the end hosts can provide a virtual end-to-end channel that can better exploit the scalable coding properties. As a further example, the combination of MD coding and path diversity was first proposed for feedback-free video applications over lossy packet networks where either the delay of retransmission or FEC is not acceptable (as in real-time video communications) or where retransmissions are not possible (as in broadcast or multicast) [1].

### 4.2. Single-Sender Path Diversity

Path diversity can be achieved over a packet network by explicitly sending packets through different network paths. In certain circumstances it is possible to specify the set of nodes or “source route” for each packet to traverse. Path diversity can then be achieved by explicitly specifying different source routes for different subsets of packets. While IP source routing is theoretically straightforward, there are a number of problems that limit its use in the current Internet, though it may be useful within a private network. It has been argued that IPv6 may allow source routing to be adopted more widely.

Another approach to explicitly send different streams over different paths is to send each stream to a different relay placed at a strategic node in the network, where each relay performs a simple forwarding operation [1, 13]. The relay infrastructure appears promising for today’s Internet, and corresponds to an application-specific overlay network on top of the conventional Internet. The proposed system routes traffic through semi-intelligent nodes at strategic locations, thereby providing a service of improved reliability while leveraging the infrastructure of the Internet. These paths should be selected to avoid joint bottleneck links, as discussed in Section 5.

### 4.3. Multiple-Sender Path Diversity

Content delivery networks (CDNs) have been widely used to provide low latency, scalability, fault tolerance, and load balancing for the delivery of web content and more recently streaming media. CDNs improve end-user performance by caching popular content on edge servers located closer to users. This helps prevent server overload and improves scalability. Furthermore, delivering content from a nearby edge server reduces the request response time, the probability of packet loss, and total network resource usage.

Another advantage of CDNs is their distributed infrastructure, which can be used to explicitly achieve path diversity between each client and multiple nearby edge servers, as in Figure 2 [8]. Furthermore, appropriately coupling MD coding with a CDN can provide improved reliability to packet losses, link outages, and server failures. This system is referred to as a Multiple Description Streaming Media Content Delivery Network or an MD-CDN for short.

An MD-CDN operates in the following manner [8]. The media is coded into multiple complementary descriptions which are distributed across different servers in the CDN. When a client requests a stream, complementary descriptions are simultaneously sent from different nearby servers through different network paths. This architecture simultaneously reaps the benefits of path diversity and of CDNs.

The multiple servers in a CDN can also provide path diversity with SD coding and FEC [7]. In this case, the

content is replicated on multiple mirror servers and multiple servers stream disjoint sets of SD and FEC packets to the single client. As discussed earlier, the reduced variability of losses afforded by path diversity makes FEC more efficient.

#### 4.4. Path Diversity over Wireless Networks

Wireless networks are quite challenging for streaming since they generally offer time-varying and unpredictable behavior caused by multiple users, interference, propagation effects, and mobility. At the network level, these effects appear as variable delays, losses, and bandwidth. Wireless links, rather than the wired infrastructure, are often the bottlenecks. As mentioned in the introduction, various forms of diversity have been used for many years to overcome these problems, particularly in the cellular environment.

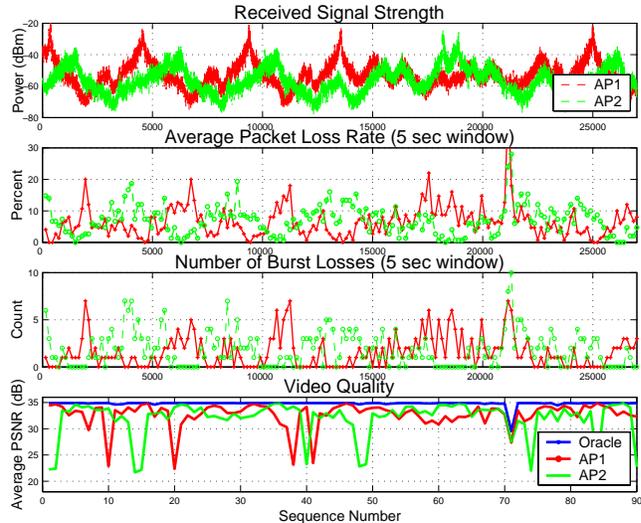
Wireless LANs such as IEEE 802.11 are becoming widespread as they provide simple connectivity and data delivery. One approach to reliable streaming over 802.11 connections, in particular for low-latency or interactive communication, is to exploit the potential path diversity between each mobile client and multiple 802.11 access points (APs) in the infrastructure [5]. A mobile 802.11 wireless client is often in range of multiple APs, each offering a different relationship to the client with respect to distance, obstructions, multipath, signal strength, contention, available bandwidth, neighboring interferers, and potential hidden nodes. Figure 3 illustrates the variability seen by a mobile client in an 802.11 network with two access points.

By using multiple paths simultaneously or by switching between multiple paths (site selection) as a function of channel characteristics, our results indicate that significant benefits may be achieved by using a wireless path diversity system compared to a conventional single access point system. These benefits may be achieved using a single client radio on a single channel [5]; there is no need to perform physical layer combining or multichannel operation.

Path diversity is also beneficial in ad-hoc wireless networks, where the topology may change or nodes may come or go. Many ad-hoc routing protocols identify multiple paths between the sender and receiver, and the use of multiple paths can improve the reliability of the connection between the two wireless hosts. The effectiveness of path diversity over an ad-hoc wireless network for MD image communication is examined in [3].

### 5. DESIGN, ANALYSIS, AND OPERATION OF MULTIPATH STREAMING SYSTEMS

The prior sections discuss the potential benefits of path diversity for media streaming and highlight scenarios where path diversity arises. This section discusses issues that arise



**Fig. 3.** Performance statistics for path diversity between a mobile 802.11 wireless client and two APs in the distributed infrastructure [5]: received packet signal strength variation (top), average packet loss rate (second), and number of loss events of burst length  $\geq 2$  (third) as a function of packet sequence number for 15-minute packet trace. Bottom plot illustrates video quality when using only AP1 (red) or only AP2 (green), and the upper bound on achievable performance by selecting the best AP for every packet (blue).

in effectively designing and using path diversity in streaming media systems.

#### 5.1. Joint and Disjoint Paths

Multiple paths are not guaranteed in practice to be independent. A path may be expressed abstractly as a sequence of links connected through nodes. Two paths often share links and nodes. The benefits of path diversity do not depend on whether paths are completely disjoint, but rather on whether bottlenecks occur on shared or disjoint portions. Shared bottlenecks reduce the impact of path diversity, while disjoint bottlenecks do not. Identifying bottlenecks and avoiding them if possible is an important element in effective use of path diversity. Joint bottleneck detection is currently an active area of research, e.g., [14].

#### 5.2. How Many Paths to Use?

How do diversity benefits scale with the number of paths and when is it worthwhile to use more than one path? This complicated question depends on the specifics of the application, the path diversity benefits that one is trying to exploit, and the characteristics of the available paths. For example, for repairable MD video coding, the improved er-

ror recovery arises with two paths, and while increasing the number of paths helps there are no further jumps in performance. End-to-end network characteristics improve in different ways with the number of paths. Total aggregated bandwidth (in principle) increases additively with the number of paths. Probability of outage decreases exponentially with the number of independent paths. Delay variability, measured in terms of standard deviation, decreases as  $1/\sqrt{N}$  where  $N$  is the number of independent paths. For media streaming in a real world scenario, we would expect that typically a small number of paths would provide a good balance between complexity and performance.

### 5.3. Streaming and Packet Scheduling across Asymmetric Paths

Different paths may offer different bandwidths, loss rates, and delay characteristics. A path diversity system should compensate for and exploit these asymmetries. The first step is to estimate the characteristics of each path. Path measurement is an active area of research, and many techniques developed for single paths may be applied to the multiple path case, hence we do not cover this topic in further depth.

When paths have different, potentially time-varying, available bandwidths, it is important to adapt the rate across paths. In the point-to-point case rate adaptation can be performed in a centralized manner. However in the multipoint-to-point case a distributed algorithm is generally required. The distributed rate allocation should not only perform the appropriate adaptation, but also for SD coding should ensure that the different senders send disjoint sets of packets (no duplication) to the receiver [7].

Paths with different packet loss rates or delays provide the opportunity to perform packet scheduling across the paths [9]. For example, more important packets may be sent over the path with the lower packet loss rate. Packets with tight delay constraints may be sent over the path with shorter delay. Furthermore, real-time video encoding can be adapted to react to the time-varying path characteristics and losses [4].

### 5.4. Selecting the Best Paths or Best Servers

As discussed above, the effectiveness of a pair of paths depends less on their lengths than on their shared vs. disjoint topology. Of still greater importance is whether bottlenecks occur on shared portions. Thus, it may be preferable to have longer paths or longer shared portions of paths if the resulting bottlenecks are not shared. This observation hints at the complexity of the important basic question: How to select the best paths to use?

Determining the best paths is important for point-to-point multipath streaming, however its impact becomes even more apparent when streaming from multiple servers, as in the case of a CDN. The fundamental problems that arise

when designing and operating a CDN are changed in important ways when using multipath streaming [8]. Three key CDN problems are:

1. Where to deploy the servers?
2. How to distribute the content?
3. How to select the best server for each client?

In a conventional CDN, where each client receives a stream from a single server over a single path, the above problems are solved by minimizing some notion of distance between a client and a server. However, streaming from multiple servers to a single client, as in an SD- or MD-CDN, requires fundamentally different metrics [8]: minimizing distance from each client to multiple servers while maximizing path diversity. Note that minimizing distance to multiple servers and maximizing diversity usually conflict. Accurately evaluating these metrics is necessary for designing and operating a CDN that uses path diversity. In addition, the server selection problem must be solved for every client request.

#### 5.4.1. Modeling Path Diversity Performance

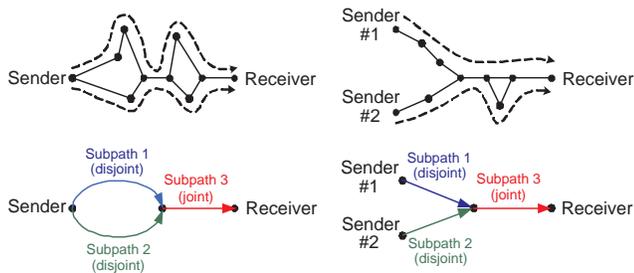
The above discussion highlights the importance of accurately modeling and predicting the performance of a path diversity system. An accurate model is needed for selecting the best subset of paths from a set of possible paths, selecting the best subset of servers from a set of possible servers, comparing path diversity scenarios, or simply evaluating the merits of a path diversity system relative to a conventional single path system.

Determining an appropriate performance model depends on the benefits (reviewed in Section 3) one intends to capture and on the particular characteristics of the media and the network. In the following we examine the problem of accurately predicting SD and MD video quality over a lossy packet network as a function of path diversity and loss characteristics. Analytical models were proposed in [15, 8] and are briefly highlighted here; some attributes of these models may be useful in other scenarios.

Our performance metric is application-level video quality, measured as mean-square error (MSE) distortion at the receiver. Different types of packet losses affect conventional SD and MD video differently. The model must therefore distinguish between isolated packet losses, burst losses of various lengths, and for MD whether the loss afflicts one or both descriptions at the same time. To account for losses on joint vs. disjoint links, each path is modeled as the concatenation of bursty single links (e.g., 2-state Gilbert model), where each link is separately identified as shared or disjoint.

Modeling media streaming using path diversity is potentially quite complex, as the sender(s) and receiver may be connected through a wide range of different topologies.

Sample topologies with one and two senders are illustrated in the top left and top right of Figure 4, respectively. It can be shown [15, 8], however, that the important end-to-end properties of a path diversity network are captured using the dramatically simplified three-subpath topologies shown in the bottom of Figure 4, where subpaths 1 and 2 are formed by the disjoint links along the first and second paths, respectively, and subpath 3 is formed by the joint links along both paths. To summarize, the loss process for two-path path diversity can be accurately modeled by three subpaths, where the associated burst loss behavior of each subpath is derived from the corresponding portion of the original path.

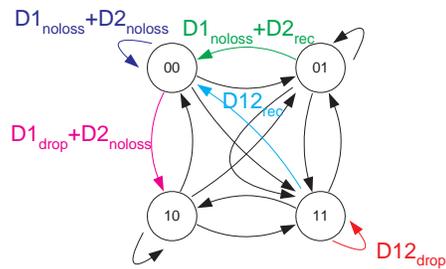


**Fig. 4.** Modeling path diversity: Examples of various path diversity topologies (top), and simplified path diversity models which preserve the end-to-end characteristics important for accurately predicting video distortion (bottom).

To form a complete application-level performance model, the packet loss model defined by the simplified topologies in Figure 4 must be coupled with a distortion model for SD and MD media streams. The result in the MD case is the finite-state model shown in Figure 5. The four states of the model represent at each point in time whether both descriptions are correctly received (state 00), one is correctly received and one is afflicted by loss (states 01 and 10), or both are simultaneously afflicted by loss (state 11). The state transition probabilities are derived from the path diversity loss model, while the distortions accumulated on each transition are derived from an MD distortion model (see [15, 8]).

The combined model in Figure 5 may be used to compare SD and MD methods in a variety of scenarios. A basic and important comparison is examining conventional SD over a single path versus MD with two-path diversity. Detailed analysis and results are given in [15, 8], but to summarize for most video sequences, as we vary the fraction of losses on the joint and disjoint links (with, e.g., 5% end-to-end average packet loss rate and average burst length of 1.25 packets) when the losses are on disjoint links then MD with path diversity is beneficial, however if the losses are on joint links then the path diversity gain vanishes and SD sometimes perform better. See [15, 8] for more details.

Additional recent work includes [9] which investigates



**Fig. 5.** Combined MD performance model: the four states identify loss conditions, the 16 transition arcs correspond to different distortions (only 5 are labeled), and the transition probabilities are determined by the network loss model.

rate-distortion optimized packet transmission schedules across multiple paths, [6] where fast heuristics are presented for quickly performing path selection when a large number of candidate path pairs are available, while accounting for the loss and on-time arrival probability of each path, and [16] which provides a recent overview of MD video coding and its transport over multiple paths.

## 6. SUMMARY

Media streaming with path diversity has gained significant interest in the last few years, as it provides valuable benefits for overcoming some of the challenges that afflict best-effort packet networks, including dynamic and unpredictable available bandwidth, delay, and loss rate. Studies have shown promise for both the single sender to single receiver case, and the multiple sender to single receiver case as in a streaming media CDN, as well as for both wired and wireless networks. Path diversity helps us take a step closer to feedback-free video streaming, which could simplify a range of applications from low-latency communication to multicast or broadcast streaming. The combination of media streaming and path diversity provides significant promise, and we look forward to its continuing evolution and adoption.

## 7. REFERENCES

- [1] J.G. Apostolopoulos, “Reliable video communication over lossy packet networks using multiple state encoding and path diversity,” *VCIP*, pp. 392–409, January 2001.
- [2] Y.J. Liang, E.G. Steinbach, and B. Girod, “Real-time voice communication over the internet using packet path diversity,” *Proc. ACM Multimedia*, pp. 431–440, Sept/Oct 2001.

- [3] N. Gogate, D. Chung, S.S. Panwar, and Y. Wang, "Supporting image/video applications in a mobile multihop radio environment using route diversity and multiple description coding," *IEEE Trans. CSVT*, pp. 777–792, Sept 2002.
- [4] Y. Liang, E. Setton, and B. Girod, "Channel-adaptive video streaming using packet path diversity and rate-distortion optimized reference picture selection," *IEEE Fifth Workshop on Multimedia Signal Processing*, pp. 420–423, Dec 2002.
- [5] A. Miu, J. Apostolopoulos, W. Tan, and M. Trott, "Low-latency wireless video over 802.11 networks using path diversity," *IEEE ICME*, pp. 441–444, July 2003.
- [6] A. Begen, Y. Altunbasak, and O. Ergun, "Fast heuristics for multi-path selection for multiple description encoded video streaming," *IEEE ICME*, pp. 517–520, July 2003.
- [7] T. Nguyen and A. Zakhor, "Distributed video streaming with forward error correction," *Packet Video Workshop*, April 2002.
- [8] J.G. Apostolopoulos, T. Wong, W. Tan, and S.J. Wee, "On multiple description streaming with content delivery networks," *IEEE INFOCOM*, pp. 1736–45, June 2002.
- [9] J. Chakareski and B. Girod, "Rate-distortion optimized packet scheduling and routing for media streaming with path diversity," *IEEE DCC*, pp. 203–212, April 2003.
- [10] N.F. Maxemchuk, *Dispersity Routing in Store and Forward Networks*, Ph.D. thesis, University of Pennsylvania, May 1975.
- [11] A. Banerjea, "Simulation study of the capacity effects of dispersity routing for fault-tolerant real-time channels," *Computer Communications Review (ACM SIGCOMM'96)*, vol. 26, no. 4, pp. 194–205, October 1996.
- [12] S. Savage, A. Collins, E. Hoffman, J. Snell, and T. Anderson, "The end-to-end effects of internet path selection," *ACM SIGCOMM*, pp. 289–299, October 1999.
- [13] T. Nguyen and A. Zakhor, "Path diversity with forward error correction (pdf) system for packet switched networks," *IEEE INFOCOM*, pp. 663–672, April 2003.
- [14] D. Rubenstein, J. Kurose, and D. Towsley, "Detecting shared congestion of flows via end-to-end measurement," *ACM SIGMETRICS*, pp. 145–155, 2000.
- [15] J.G. Apostolopoulos, W. Tan, S.J. Wee, and G.W. Wornell, "Modeling path diversity for multiple description video communication," *ICASSP*, pp. 2161–2164, May 2002.
- [16] Y. Wang, A. Reibman, and S. Lin, "Multiple description coding for video communications," *Submitted to Proceedings of the IEEE*, 2004.