Multi-path Continuous Media Streaming: What are the Benefits?

L. Golubchik [a,1] J. C.S. Lui [b] T.F. Tung [b] A. L.H. Chow [b]
W.-J. Lee [c] G. Franceschinis [d] C. Anglano [d]

[a] *Computer Science Department, University of Southern California*

[b] *Department of Computer Science & Eng., Chinese University of Hong Kong*

[c] *Department of Computer Science, University of Maryland at College Park*

[d] *Dipartimento di Scienze e Tecnologie Avanzate, Università del Piemonte Orientale.*

Abstract

Quality of service (QoS) in delivery of continuous media over the Internet is still relatively poor and inconsistent. Although many such applications can tolerate some degree of missing information, significant losses degrade an application's QoS. In this paper we investigate the potential benefits of mitigating this problem through the exploitation of multiple paths existing in the network between a set of senders and a receiver of continuous media. Our focus in this work is on providing a fundamental understanding of the benefits of using multiple paths to deliver continuous media over best-effort wide-area networks. Specifically, we consider pre-recorded continuous media applications and use the following metrics in evaluating the performance of multi-path streaming as compared to single-path streaming: (a) data loss rate, (b) conditional error burst length distribution, and (c) lag1-autocorrelation. The results of this work can be used in guiding the design of multi-path continuous media systems streaming data over best-effort wide-area networks.

1 Introduction

Quality of service (QoS) in streaming of continuous media over the Internet is still poor and inconsistent. The degradation in quality of continuous media applications is partly due to variations in delays as well as losses experienced

[1] This work was partly done while the author was with the Department of Computer Science and UMIACS at the University of Maryland.

by packets sent through wide-area networks. Although many such applications can tolerate some degree of missing information, significant losses degrade an application's quality of service. One approach to providing QoS for continuous media applications over the Internet is to use the IntServ model for signaling (e.g., RSVP) and resource reservation in all routers along the streaming path. However, this approach suffers from scalability and deployment problems. In contrast, we investigate the potential benefits of providing QoS guarantees in continuous media delivery through the *exploitation of multiple paths* existing in the network between a set of senders and a receiver. One advantage of this approach is that the complexity of QoS provision can be pushed to the network edge and hence improve the scalability and deployment characteristics while at the same time provide a certain level of QoS guarantees. Our focus in this work is on providing a fundamental understanding of the benefits of using *multiple paths* to deliver continuous media data destined for a particular receiver, i.e., this data is fragmented into packets and the different packets take alternate routes to the receiver.

There are a number of approaches to accomplishing a multi-path data delivery, and we describe the specific approach considered in our system below. We first note that such paths do not have to be completely disjoint, i.e., it is sufficient for them to have disjoint points of congestion or bottlenecks. Existence of multiple paths with *disjoint bottlenecks* includes the following potential benefits.

- *Reduction in correlation between consecutive packet losses.* Although a continuous media (CM) application can tolerate some missing information, a large number of consecutive packet losses not only contributes to significant degradation in CM quality but also diminishes ability to correct such losses through error correction techniques. As we will show in this paper, sending data through multiple paths can potentially reduce burst lengths and correlations between consecutive losses and thus improve the quality of delivered data.
- *Increased throughput.* In delivery of continuous media one can reduce the amount of bandwidth needed to deliver the data at the cost of its quality. Sending data through multiple paths potentially increases the amount of (aggregate) bandwidth available to the application and hence increases the quality of delivered data.
- *Ability to adjust to variations in congestion patterns on different parts of the network.* CM applications are often long lasting. Hence, it is reasonable to expect that network conditions will change throughout the delivery of data to a CM application. Since not all paths, in general, would experience the same traffic patterns and congestion, sending data through multiple paths potentially improves the ability to adapt to changes in network conditions.

In general, the use of multiple paths in designing of distributed (over best-effort wide-area networks) CM applications requires consideration of the following issues.

- *Determining bottlenecks, joint points of congestion, and network characteristics in general.* To gain the benefits of multi-path streaming described above, one must first determine the paths to be used in delivery of the data. Since it is reasonable to characterize a path using its bottleneck link [1], what we need to be able to do is determine whether a number of paths share points of congestion [18]. Although this is not necessary in our approach, other approaches to multi-path streaming might require fairly accurate estimation of various network characteristics (refer to Section 4). These are non-trivial problems which are outside the scope of this paper. However, we note that currently we use [18] in our system for detecting shared points of congestion.
- *Effects of redundancy and error erasure schemes.* Some amount of lost data can be reconstructed in CM applications through the use of redundant information. Hence, in constructing multi-path streaming techniques one should take into consideration the effect of redundant information on the final quality of the data and how the erasure codes interact with multi-path delivery.
- *Adaptation schemes under changes in network conditions.* When network conditions change, one can improve the quality of CM by adapting how the data is streamed on multiple paths.
- *Data placement.* Proper placement of data on the servers is an issue in the context of CM applications delivering *pre-stored* data. Inappropriate data placement can adversely affect servers' performance. For instance, this can occur due to load imbalance problems arising from the fact that only specific parts of the data are being delivered from a particular server as well as the fact that specific data required might change over the course of the application, as the system adapts to congestion patterns in the network. This in turn reduces the quality of service experienced by the CM application (in this case due to server rather than network performance). We note that these problems can be more severe when adaptation schemes are used.
- *Data dispersion.* Given that one cannot necessarily rely on the network layer to provide multi-path routing, another consideration is how to accomplish the dispersion of data over multiple paths existing in the network between a sender and a receiver of data. This may be an especially important consideration for applications where data is generated live, e.g., a video conferencing application, in contrast to applications where data is pre-recorded.
- *Need for protocol/network support.* Lastly, some mechanisms for streaming application data over multiple paths might require support from lower layers, such as the network layer. In this case, ease of deployment is an issue.

Although all these issues are of importance, in this paper we narrow the scope by focusing on:

- delivery of *pre-stored* video, e.g., as in video-on-demand applications (in contrast to delivery of "live" data as in video-conferencing applications);
- *application-level* schemes (which are deployable today over the current Internet) — that is, we assume the use of best-effort IP-based networks, where a specific path is used between any pair of hosts (sender and receiver) on the network **and** this path is determined by a network-level routing algorithm; furthermore, our system does not require specific knowledge of the paths, only the ability to determine whether two paths share a point of congestion, e.g., using [18];
- accomplishment of multiple paths to the same receiver by *distributing servers* across wide-area networks and streaming data from multiple servers simultaneously;
- streaming over the network issues *only* (rather than, e.g., considering server-related problems mentioned above); that is, for the purposes of this paper we assume that the data is fully replicated at all servers and hence any server can deliver any fraction of the CM data.

In our system, server i sends fraction α_i of the data expected by the receiver, where $0 \leq \alpha_i \leq 1$ and $\sum_i \alpha_i = 1$. In general, we assume that the setting and possible adaptation of these fractions (as the delivery of data progresses) is done by the receiver (based on its perceived quality of data and determination of joint points of congestion). The receiver assembles the data from multiple senders and plays it in the appropriate order.

In the remainder of the paper, our focus is on providing fundamental understanding and on *characterizing the benefits* of the multi-path approach to streaming of pre-stored continuous media data over wide-area networks, under the setup described above. Specifically, we focus on loss characteristics as they are an indication of the resulting quality of the delivered data stream. We believe that the understanding of loss characteristics under a multi-path approach is non-trivial and deserves further attention. We also believe that the work presented here is a step in the right direction. The contributions of this paper are as follows. We give an analytical characterization of when a multi-path approach is beneficial, as compared to a single path approach, using the following metrics (a) packet loss rate, (b) lag-1 autocorrelation of packet losses, and (c) burst length distribution. We also extend this analysis to information loss rate, i.e., we consider the resulting losses after an application of an erasure code. Secondly, we extend the evaluation of the multi-path approach benefits using simulations of the analytical model. These are also performed with and without the use of an erasure code. Our results indicate that: (1) in general, multi-path streaming exhibits better loss characteristics than single-path streaming, (2) use of an erasure code may not necessarily improve data loss characteristics in the case of single-path streaming, while multi-path streaming (with or without use of an erasure code) can improve data loss characteristics, and (3) lag1-autocorrelation of multi-path streaming is usually closer to zero than that of single path streaming, and we believe that this will also result in a higher viewing quality of the received CM.

2 Analytical Evaluation

In this section, we present our analysis of the single-path and the multi-path streaming approaches. We first consider these approaches without the use of erasure codes, so as to understand the basic differences between single and multi-path streaming. We then also consider the changes in loss characteristics when redundant information is added, as this is another approach to dealing with packet losses. Specifically, we consider a variation of such codes, which we refer to as FEC, as defined below. As in [1], we use a two-state Markov chain model, known as the Gilbert model; as in [1] we characterize the path by its bottleneck link. This model allows for dependence in consecutive packet losses and should be a more accurate representation of the loss process in the network than an independent loss model.

We use the following performance measures to quantify the merits of the different streaming approaches: (1) mean data packets loss rate (with and without FEC), (2) conditional burst length distribution, conditioned on there being at least one error (with and without FEC), (3) lag-1 auto-correlation (with and without FEC). The first performance measure is an obvious approach to comparing single and multi-path streaming (when losses, rather than throughput, are of importance). The other two performance measures are less obvious; however, we believe that they can significantly affect the quality of the viewed continuous media.

To illustrate this point, we first give a brief motivation for considering above given performance metrics, and specifically, for considering burst lengths and correlations between losses. We discuss this in the context of video data. Ideally, one would like to have a measure of the perceptual quality of the viewed video, as a function of loss characteristics. To the best of our knowledge, there is no such widely accepted measure, and often the quality of a video is evaluated using human observers. However, some metrics have been used in the past, for instance, signal to noise ratio of the resulting video [7]. Hence, we performed an experiment to illustrate the effects of bursty losses on the quality of the resulting video (and specifically on the signal to noise ratio). In this experiment, we introduced losses in a variety of "patterns", e.g., they can be evenly spaced throughout video, or they can be more bursty. (Due to lack of space, we do not give the details of the experiment and refer the interested reader to [6].) From this experiment, we observed that given the same amount of information loss, the signal to noise metric can be significantly lower for the more bursty loss patterns, and hence, the perceptual quality of a video may potentially be degraded. Thus, we believe that burst length distribution and correlations between losses are appropriate metrics for evaluating the goodness of a streaming approach as they can reflect the quality of the received video.

2.1 Model

Let us now state the path model used in this paper. As in [1], we use a stationary continuous time Gilbert model to characterize the potential correlations between consecutive losses on a path. Under a stationary continuous time Gilbert model, the packet loss process along path k is described by a two state continuous time Markov chain $\{X_k(t)\}$ where $X_k(t) \in \{0, 1\}$. If a packet is transmitted at time t when the state of path k is $X_k(t) = 0$, then no packet loss occurs. On the other hand, the transmitted packet is considered lost if $X_k(t) = 1$. The infinitesimal generator for this Gilbert model of path k is:

$$\boldsymbol{Q}_k = \begin{bmatrix} -\mu_0(k) & \mu_0(k) \\ \mu_1(k) & -\mu_1(k) \end{bmatrix}.$$

The stationary distribution of this Gilbert model is $\boldsymbol{\pi}(k) = [\pi_0(k), \pi_1(k)]$ where $\pi_0(k) = \mu_1(k)/(\mu_0(k) + \mu_1(k))$ and $\pi_1(k) = \mu_0(k)/(\mu_0(k) + \mu_1(k))$. Let $p_{i,j}^{(k)}(\tau)$ be the probability that path k is in state j at time $t + \tau$, given that it was in state i at time t, i.e., $p_{i,j}^{(k)}(\tau) = P(X_k(t + \tau) = j | X_k(t) = i)$. From [14], we have that

$$p_{i,j}^{(k)}(\tau) = \begin{cases} \frac{\mu_1(k)}{\mu_0(k)+\mu_1(k)} \left(1 - e^{-[\mu_0(k)+\mu_1(k)]\tau}\right) & i = 1, j = 0, \\ \frac{\mu_0(k)}{\mu_0(k)+\mu_1(k)} \left(1 - e^{-[\mu_0(k)+\mu_1(k)]\tau}\right) & i = 0, j = 1, \\ \frac{\mu_0(k)+\mu_1(k)e^{-(\mu_0(k)+\mu_1(k))\tau}}{\mu_0(k)+\mu_1(k)} & i = 1, j = 1, \\ \frac{\mu_1(k)+\mu_0(k)e^{-(\mu_0(k)+\mu_1(k))\tau}}{\mu_0(k)+\mu_1(k)} & i = 0, j = 0 \end{cases} \tag{1}$$

for all $\tau > 0$.

Throughout the paper we refer to single path streaming as SP streaming and multi-path streaming with N paths as MP streaming. Without loss of generality, when paths are homogeneous, we assume that SP streaming always transmits data along path 1. In the evaluation of MP streaming, we assume that the multiple paths have disjoint bottlenecks (or points of congestion) and hence the Gilbert models representing them are independent. Note that, since we represent a path by its bottleneck link, multiple paths with joint points of congestion could just be represented by a single Gilbert model. Lastly, note that our focus is on a streaming application which generates packets at a constant rate; hence our derivations below are done under this assumption.

2.2 Performance Analysis of SP vs. MP Streaming (without FEC)

Let us first derive the average packet loss rate. Unless stated otherwise, below we consider a special case of MP streaming, namely dual path, round robin (DPRR) streaming. There are a number of different approaches to distributing data along the multiple paths; unless otherwise stated, we consider a simple case, i.e., DPRR, wherein each path carries half the application's traffic and the packet transmission is carried out in a round robin manner. That is, odd numbered packets are transmitted along path 1 while even numbered packets are transmitted along path 2. We use this simple scheme for dual path streaming to illustrate the basic performance differences between SP and MP streaming, so as to gain some basic understanding.

If we assume that the streaming rate does not affect the channel loss characteristics (i.e., the parameters of the Gilbert model), then for the SP case, the average packet loss rate is simply

$$P_{sp}[\text{loss packet}] = \pi_1(1) \;=\; \frac{\mu_0(1)}{\mu_0(1)+\mu_1(1)}. \tag{2}$$

For the MP case, assume that we have $N \geq 1$ paths and let α_i be the fraction of the application's workload that is sent along path i where $\sum_{i=1}^{N} \alpha_i = 1$. Then the average packet loss rate for the MP case is

$$P_{mp}[\text{loss packet}] = \sum_{i=1}^{N} \alpha_i \pi_1(i) \;=\; \sum_{i=1}^{N} \alpha_i \left(\frac{\mu_0(i)}{\mu_0(i)+\mu_1(i)} \right).$$

If the N paths are homogeneous, then we can simplify as follows

$$P_{mp}[\text{loss packet}] = \frac{\mu_0(1)}{\mu_0(1)+\mu_1(1)}. \tag{3}$$

Remark: the implication of Equations (2) and (3) is that if the application's sending rate does not affect the loss characteristics of a path then splitting data between multiple homogeneous paths does *not* reduce the average packet loss rate, as compared to a single path with the same loss characteristics.

On the other hand, if the application's sending rate can affect the loss characteristics of the path (e.g., sending data with a higher bandwidth may increase the losses), then the average loss rate of the MP approach can be different from that of the SP approach. To illustrate this effect, let λ be the application's mean sending rate, and let $\mu_0(i) = \mathcal{F}(\lambda)$ and $\mu_1(i) = \mathcal{B}(\lambda)$, where $\mathcal{F}$ ($\mathcal{B}$) is a continuous non-decreasing (non-increasing) function of λ. Then, we have the following result.

Theorem 1 *If the parameters of the Gilbert model are specified by functions $\mathcal{F}$ and $\mathcal{B}$, then the average packet loss rate under the single path streaming approach will be greater than or equal to the average packet loss rate under the multi-path streaming approach wherein these paths have the same Gilbert model parameters.*

Proof: Omitted due to lack of space; please refer to [6]. ∎

Let us now consider the conditional burst length distribution, of both SP and MP cases, conditioned on there being a loss. Let λ_1 be the mean streaming rate (in units of packets per second) along path 1 and $\delta_1 = 1/\lambda_1$ be the time between two consecutively transmitted packets. Then, in the SP case (as also derived in [1] for a voice-over-IP type application), the probability of having a packet error burst of length $m \geq 1$ is:

$$\mathrm{P}_{sp}[\text{length of error burst} = m] = \begin{cases} \pi_0(1)p_{0,1}^{(1)}(\delta_1)p_{1,0}^{(1)}(\delta_1) & \text{for } m = 1, \\ \pi_0(1)p_{0,1}^{(1)}(\delta_1)\left[p_{1,1}^{(1)}(\delta_1)\right]^{m-1} p_{1,0}^{(1)}(\delta_1) & \text{for } m \geq 2. \end{cases} \quad (4)$$

The probability of having a packet error burst of any length is therefore

$$\mathrm{P}_{sp}[\text{error burst}] = \sum_{m=1}^{\infty} \mathrm{P}_{sp}[\text{length of error burst} = m] = \pi_0(1)p_{0,1}^{(1)}(\delta_1).$$

Moreover, the conditional probability of having a packet error burst of length $m \geq 1$, conditioned on there being a loss, is equal to

$$\mathrm{P}_{sp}[\text{error burst of length } m|\text{ error burst}] = \frac{\mathrm{P}_{sp}[\text{length of error burst} = m]}{\mathrm{P}_{sp}[\text{error burst}]} = \left[p_{1,1}^{(1)}(\delta_1)\right]^{m-1} p_{1,0}^{(1)}(\delta_1) \text{ for } m \geq 1. \quad (5)$$

In the MP case, let us consider the special case of DPRR streaming, i.e., $N = 2$. Let λ_2 be the streaming rate (in units of packets per second) along path 1 or path 2. Note that under DPRR, $\lambda_2 = \lambda_1/2$. Then, the time between two consecutively transmitted packets along the same path is $\delta_2 = 1/\lambda_2 = 2\delta_1$.

To understand the basic tradeoff between SP and MP streaming, we also assume that both paths are homogeneous such that they are characterized by a stationary continuous time Gilbert model of the same parameters (i.e., $\mu_0(1) = \mu_0(2)$ and $\mu_1(1) = \mu_1(2)$). Given this simplification, the stationary distributions for both paths are the same (i.e., $\pi_0(1) = \pi_0(2)$; $\pi_1(1) = \pi_1(2)$) and we can express all performance measures using the parameters of path 1. Under these assumptions, the probability of having a packet error burst of length $m \geq 1$ is:

$$\begin{aligned} &\mathrm{P}_{dp}[\text{length of error burst} = m] \\ &\quad = \begin{cases} \pi_0(1)\pi_1(1)p_{0,0}^{(1)}(2\delta_1) & \text{for } m = 1, \\ \pi_0(1)\pi_1(1)\left[p_{1,1}^{(1)}(2\delta_1)\right]^{m-2} p_{0,1}^{(1)}(2\delta_1)p_{1,0}^{(1)}(2\delta_1) & \text{for } m \geq 2. \end{cases} \end{aligned} \tag{6}$$

and the probability of having a packet error burst of any length is therefore:

$$\begin{aligned} \mathrm{P}_{dp}[\text{error burst}] &= \sum_{m=1}^{\infty} \mathrm{P}_{dp}[\text{length of error burst} = m] \\ &= \pi_0(1)\pi_1(1)p_{0,0}^{(1)}(2\delta_1) + \sum_{m=2}^{\infty} \pi_0(1)\pi_1(1)[p_{1,1}^{(1)}(2\delta_1)]^{m-2}p_{0,1}^{(1)}(2\delta_1)p_{1,0}^{(1)}(2\delta_1) \\ &= \pi_0(1)\pi_1(1)\left[p_{0,0}^{(1)}(2\delta_1) + \frac{p_{0,1}^{(1)}(2\delta_1)p_{1,0}^{(1)}(2\delta_1)}{1 - p_{1,1}^{(1)}(2\delta_1)}\right] \\ &= \pi_0(1)\pi_1(1)\left[p_{0,0}^{(1)}(2\delta_1) + p_{0,1}^{(1)}(2\delta_1)\right] = \pi_0(1)\pi_1(1). \end{aligned}$$

Then, the conditional probability of having a packet error burst of length $m \geq 1$, conditioned on there being a packet error, is equal to:

$$\begin{aligned} \mathrm{P}_{dp}[\text{error burst of length } m | \text{ error burst}] &= \frac{\mathrm{P}_{dp}[\text{length of error burst} = m]}{\mathrm{P}_{dp}[\text{error burst}]} \\ &= \begin{cases} p_{0,0}^{(1)}(2\delta_1) & \text{for } m = 1, \\ \left[p_{1,1}^{(1)}(2\delta_1)\right]^{m-2} p_{0,1}^{(1)}(2\delta_1)p_{1,0}^{(1)}(2\delta_1) & \text{for } m \geq 2. \end{cases} \end{aligned} \tag{7}$$

We can now state the conditions under which the DPRR approach will have a smaller conditional burst error length than the SP approach. Before we present this result, let us present the definition and a basic lemma of stochastic comparison [17].

Definition 1 *We say that the random variable X is stochastically larger than the random variable Y, written $X \geq_{st} Y$, if $P[X \geq z] \geq P[Y \geq z]$ for all z.*

Lemma 1 *We say that $X \geq_{st} Y$ iff $E[f(X)] \geq E[f(Y)]$ for all increasing functions f.*

Now, let $\mathcal{B}_{sp}$ and $\mathcal{B}_{dp}$ be the random variables representing the conditional packet error burst length, given that there is at least one packet error, under the SP and the homogeneous DPRR approaches, respectively. Then, we have the following result.

Theorem 2 *If $p_{0,1}(2\delta_1)p_{1,0}(2\delta_1) \leq p_{1,1}(\delta_1)p_{1,0}(\delta_1)$, then $\mathcal{B}_{sp} \geq_{st} \mathcal{B}_{dp}$.*

Proof: First, note that $p_{1,1}(t)$ is an non-increasing function of t. If $p_{0,1}(2\delta_1)p_{1,0}(2\delta_1)$ $\leq p_{1,1}(\delta_1)p_{1,0}(\delta_1)$, then from Equations (5) and (7), we can deduce that for $m \geq 2$, we have

$$\mathrm{P}_{dp}[\text{error burst of length } m| \text{ error burst}] \leq \mathrm{P}_{sp}[\text{error burst of length } m| \text{ error burst}]$$

Since

$$\begin{aligned} \sum_{m=1}^{\infty} P_{sp}[\mathcal{B}_{sp} = m] &= \sum_{m=1}^{\infty} P_{dp}[\mathcal{B}_{dp} = m] = 1 && \text{and} \\ \sum_{m=j}^{\infty} P_{sp}[\mathcal{B}_{sp} = m] &\geq \sum_{m=j}^{\infty} P_{dp}[\mathcal{B}_{dp} = m] && \text{for } j \geq 2, \end{aligned}$$

we can conclude that $\mathcal{B}_{sp} \geq_{st} \mathcal{B}_{dp}$. ∎

Remark: Note that $\mathcal{B}_{sp} \geq_{st} \mathcal{B}_{dp}$ implies (based on Lemma 1) that $E[f(\mathcal{B}_{sp})] \geq E[f(\mathcal{B}_{dp})]$ for all increasing functions f. Therefore, we can conclude that for all moments of $\mathcal{B}_{sp}$ and $\mathcal{B}_{dp}$, we have $E[\mathcal{B}_{sp}^k] \geq E[\mathcal{B}_{dp}^k]$ for $k \geq 1$, where $E[\mathcal{B}_{sp}^k]$ and $E[\mathcal{B}_{dp}^k]$ refer to the k^{th} moments of $\mathcal{B}_{sp}$ and $\mathcal{B}_{dp}$, respectively. One implication of the above theorem is that the homogeneous DPRR approach will have a lower mean conditional burst length than the SP approach, given that the theorem's condition is satisfied.

Let us now consider the lag-1 autocorrelation of packet errors metric. We begin with the SP approach. The lag-1 autocorrelation function $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}]$ measures the degree of dependency of consecutive packet errors, where $\mathcal{X}_t$ is a random variable indicating whether the packet sent at time t is lost (indicated by state 1 in Figure 1) or received properly (indicated by state 0 in Figure 1). For example, a high positive value of $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}]$ implies that a lost packet is very likely to be followed by another lost packet. On the other hand, a high negative value of $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}]$ implies that a lost packet is likely to be followed by a successful packet arrival. Also, if the statistics of the consecutive packet losses are not correlated [2], then $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}] = 0$.

The lag-1 autocorrelation for the SP approach is

$$R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}] = \frac{E[(\mathcal{X}_t - \overline{\mathcal{X}})(\mathcal{X}_{t+\delta_1} - \overline{\mathcal{X}})]}{E[(\mathcal{X}_t - \overline{\mathcal{X}})^2]} = \frac{E[\mathcal{X}_t\mathcal{X}_{t+\delta_1} - \overline{\mathcal{X}}^2]}{E[\mathcal{X}_t^2 - \overline{\mathcal{X}}^2]}.$$

Since $\overline{\mathcal{X}} = \pi_1(1) = \mu_0(1)/[\mu_0(1) + \mu_1(1)]$, $E[\mathcal{X}_t\mathcal{X}_{t+\delta_1}] = \pi_1(1)p_{1,1}^{(1)}(\delta_1)$ and $E[\mathcal{X}_t^2] = \pi_1(1) = \mu_0(1)/[\mu_0(1)+\mu_1(2)]$, substituting these expressions into the above equation, gives us

$$R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}] = \frac{\pi_1(1)p_{1,1}^{(1)}(\delta_1) - \pi_1^2(1)}{\pi_1(1)[1 - \pi_1(1)]} = \frac{[\mu_0(1) + \mu_1(1)]p_{1,1}^{(1)}(\delta_1) - \mu_0(1)}{\mu_1(1)}. \quad (8)$$

Lemma 2 *For a high (low) bandwidth streaming application, the lag-1 autocorrelation of the SP streaming approach is positively correlated (tends to zero).*

Proof: Note that when $\delta_1 \to 0$ and $p_{1,1}^{(1)}(\delta_1) \to 1$, the lag-1 autocorrelation $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}]$ approaches 1. In other words, if the streaming application has a high bandwidth requirement such that the inter-packet spacing tends to zero, then the consecutive packet losses are *"positively"* correlated. On the other hand, when $\delta_1 \to \infty$ and $p_{1,1}^{(1)}(\delta_1) \to \mu_0(1)/[\mu_0(1)+\mu_1(1)]$, the lag-1 autocorrelation $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}] \to 0$. This implies that for low bandwidth streaming applications, wherein the inter-packet spacing is very large, the lag1-autocorrelation tends to zero. ∎

Let us also derive the lag-1 autocorrelation of the homogeneous DPRR approach. The lag-1 autocorrelation in this case is:

$$R[\mathcal{X}_t^{(1)}\mathcal{X}_{t+\delta_1}^{(2)}] = \frac{E[(\mathcal{X}_t^{(1)} - \overline{\mathcal{X}^{(1)}})(\mathcal{X}_{t+\delta_1}^{(2)} - \overline{\mathcal{X}^{(2)}})]}{\sqrt{E[(\mathcal{X}_t^{(1)} - \overline{\mathcal{X}^{(1)}})^2]E[(\mathcal{X}_t^{(2)} - \overline{\mathcal{X}^{(2)}})^2]}}. \quad (9)$$

where $\mathcal{X}_t^{(i)}$ is a random variable indicating whether the packet sent at time t on path i $(i = 1, 2)$ is lost or received properly. Because both paths are homogeneous, we can simplify as follows:

[2] Note that if the lag-1 autocorrelation, $R[\mathcal{X}_t\mathcal{X}_{t+\delta_1}]$, is equal to 0, it does not necessarily imply that consecutive packet losses are not correlated.

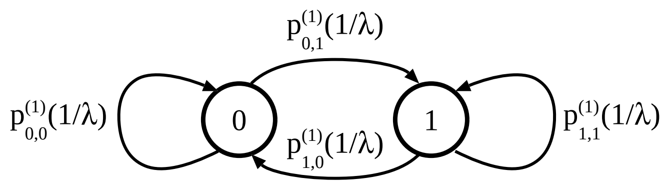

Fig. 1. An Embedded Markov Chain which describes whether a transmitted packet is lost or not.

$$R[\mathcal{X}_t^{(1)}\mathcal{X}_{t+\delta_1}^{(2)}] = \frac{E[\mathcal{X}_t^{(1)}\mathcal{X}_{t+\delta_1}^{(2)} - \overline{\mathcal{X}^{(1)}}^2]}{E[(\mathcal{X}_t^{(1)} - \overline{\mathcal{X}^{(1)}})^2]} = \frac{E[\mathcal{X}_t^{(1)}\mathcal{X}_{t+\delta_1}^{(2)}] - E[\overline{\mathcal{X}^{(1)}}^2]}{E[\left(\mathcal{X}_t^{(1)}\right)^2] - E[\left(\overline{\mathcal{X}^{(1)}}\right)^2]}$$
$$= \frac{\left(\frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}\right)\left(\frac{\mu_1(2)}{\mu_0(2)+\mu_1(2)}\right) - \left(\frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}\right)^2}{\mu_0(1)\mu_1(1)/\left(\mu_0(1)+\mu_1(1)\right)^2}$$
$$= \frac{\left(\frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}\right)^2 - \left(\frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}\right)^2}{\mu_0(1)\mu_1(1)/\left(\mu_0(1)+\mu_1(1)\right)^2} = 0 \tag{10}$$

In fact, we can see that the consecutive packet losses under the homogeneous DPRR approach are *"uncorrelated"* since we have assumed independence of the two paths.

2.3 Performance Analysis of SP vs. MP Streaming (with FEC)

We have shown that loss characteristics can be improved with multi-path streaming as compared to single path streaming, under conditions and metrics specified above. However, an interesting question that remains is whether there are still benefits to be gained once some form of redundancy is added to the stream. Specifically, we consider the use of an erasure code (as defined below), to which we will refer as FEC in the remainder of the paper. FEC-based techniques are widely used, e.g., as in [1,5]; hence, in this section we focus on the basic understanding of the performance of SP vs. MP streaming when FEC is added to the stream.

Since numerous coding schemes exist, we first give the details of the simple FEC scheme considered here. We divide a video file into groups of data packets such that each group consists of k data packets. Given each group of k data packets, we generate $n > k$ packets. We refer to these n packets as a FEC group. The encoding scheme is such that, if the number of lost packets within a FEC group is less than or equal to $(n - k)$, then we can reconstruct the original k data packets within that FEC group.

Let us first derive the average packet loss rate under the SP approach. As before, assume that we use path 1 which is characterized by a Gilbert model, as defined above, with parameters $\mu_0(1)$ and $\mu_1(1)$. The streaming application generates packets at a rate of λ (in units of packet/sec)[3]. Whenever a packet is transmitted along this path, it may be lost (if the state of the path is "1") or it may arrive successfully at the receiver (if the state of the path is "0"). Figure 1 depicts an embedded Markov chain of this path wherein the two consecutive embedded points are $1/\lambda$ units apart. The derivation of transition probabilities of this DTMC is based on Equation (1); hence they are a function of the Gilbert model's parameters $\mu_0(1)$ and $\mu_1(1)$ as well as the packet transmission rate λ. The steady state probabilities of this embedded Markov chain are $\pi_0(1) = \frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}$ and $\pi_1(1) = \frac{\mu_0(1)}{\mu_0(1)+\mu_1(1)}$.

We are now interested in deriving $P^{(1)}(j,n)$, which is the probability of losing j packet in an n packet transmission. We define

$$P_i^{(1)}(j,n) = \text{Prob}(j,n|\text{initial state of the path is } i) \qquad i \in \{0,1\}$$

as the probability of j lost packets in an n packet transmission, given that the *first* packet was transmitted when the path was in state i (where $i \in \{0,1\}$). We then have:

$$P^{(1)}(j,n) = P_0^{(1)}(j,n)\pi_0(1) + P_1^{(1)}(j,n)\pi_1(1) \qquad j = 0,1,\ldots,n. \tag{11}$$

Let $L_i^{(1)}(j,n)$ ($H_i^{(1)}(j,n)$) be the probability that we have j lost packets in an n packet transmission and that the *last* packet was transmitted when the path was in state 0 (state 1), given that the *first* packet was transmitted when the path was in state i, where $i \in \{0,1\}$. Then we have:

$$P_i^{(1)}(j,n) = L_i^{(1)}(j,n) + H_i^{(1)}(j,n) \qquad i \in \{0,1\} \text{ and } j = 0,1,\ldots,n. \tag{12}$$

We can also express $L_i^{(1)}(j,n)$ and $H_i^{(1)}(j,n)$, for $j < n$, using the following recursive forms:

$$L_i^{(1)}(j,n) = L_i^{(1)}(j,n-1)(1-p_{0,1}^{(1)}(1/\lambda)) + H_i^{(1)}(j,n-1)p_{1,0}^{(1)}(1/\lambda), \tag{13}$$
$$H_i^{(1)}(j,n) = L_i^{(1)}(j-1,n-1)p_{0,1}^{(1)}(1/\lambda) + H_i^{(1)}(j-1,n-1)(1-p_{1,0}^{(1)}(1/\lambda)) \tag{14}$$

where we also have the following boundary conditions:

$$L_i^{(1)}(j,m) = 0 \qquad i \in \{0,1\};\ j = 0,1,\ldots,n \text{ and } m \le j \tag{15}$$

[3] Note that here, "packets" includes both data packets and packets carrying redundant information.

$$
\begin{aligned}
L_0^{(1)}(0,m) &= (1-p_{0,1}^{(1)}(1/\lambda))^{m-1} && \text{for } m=1,2,\ldots,n && (16)\\
L_1^{(1)}(0,m) &= 0 && \text{for } m=1,2,\ldots,n && (17)\\
H_i^{(1)}(j,m) &= 0 && i\in\{0,1\}; j=1,2,\ldots,n \text{ and } m<j && (18)\\
H_i^{(1)}(0,m) &= 0 && \text{for } i\in\{0,1\} \text{ and } m=0,1,\ldots,n && (19)\\
H_0^{(1)}(m,m) &= 0 && \text{for } m=1,2,\ldots,n && (20)\\
H_1^{(1)}(m,m) &= (1-p_{1,0}^{(1)}(1/\lambda))^{m-1} && \text{for } m=1,2,\ldots,n. && (21)
\end{aligned}
$$

Remark: To compute the value of $P^{(1)}(j,n)$ in Equation (11), we need to compute the values of the four square matrices $\boldsymbol{L}_0^{(1)}, \boldsymbol{L}_1^{(1)}, \boldsymbol{H}_0^{(1)}$, and $\boldsymbol{H}_1^{(1)}$, whose entries can be computed using Equations (13) through (21). Each of these matrices is of size $(n+1)\times(n+1)$. In other words, computing the values of $P^{(1)}(j,n)$ (for all j) has a computational complexity of $\Theta(4(n+1)^2)$.

Let P_{sp} be the probability of an irrecoverable error within a FEC group. It is equal to

$$
\begin{aligned}
P_{sp} = \sum_{j=n-k+1}^{n} P^{(1)}(j,n) &= \sum_{j=n-k+1}^{n} \left[P_0^{(1)}(j,n)\pi_0(1) + P_1^{(1)}(j,n)\pi_1(1)\right]\\
&= \sum_{j=n-k+1}^{n} \left[\left(L_0^{(1)}(j,n) + H_0^{(1)}(j,n)\right)\left(\frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}\right) + \right.\\
&\qquad \left. + \left(L_1^{(1)}(j,n) + H_1^{(1)}(j,n)\right)\left(\frac{\mu_0(1)}{\mu_0(1)+\mu_1(1)}\right)\right].
\end{aligned}
$$

To derive the average data packet loss rate (with the use of FEC) for the SP approach, denoted by $\mathcal{L}_{sp}$, we consider the following two cases, based on the number of lost packets, $j\in\{0,1,\ldots,n\}$, within a FEC group.

Case 1: $j \le n-k$
If j, the number of lost packets within a FEC group, is less than or equal to $n-k$, then all k data packets can be reconstructed at the receiver. Hence, this case does not contribute to *information* loss and $\mathcal{L}_{sp}=0$.

Case 2: $j > n-k$
In this case, the lost data packets cannot be fully reconstructed and some information will be lost. However, given that there j lost packets within a FEC group, there are a number of different ways to distribute these losses among the n packets of the FEC group. To understand this effect, let us illustrate it using an example. Assume that $n=5$ and $k=4$. If $j=2$, then there are two possible ways to distribute these two lost packets among the packets of the FEC group: (1) the two lost packets are the data packets within the FEC group, or (2) one lost packet is a data packet and the other lost packet

corresponds to redundant information in the FEC code. In the first case, we lost 2 data packets out of a 4 data packet transmission. In the second case, we lost 1 data packet out of a 4 data packet transmission. Using the same argument, if $j = 5$, then there is only one way to distribute these five lost packets among packets of the FEC group. That is, all data packets are lost. Therefore, given that there are j lost packets, the number of ways to distribute the j lost packets among the packets of a FEC group is $\mathcal{W} = \mathcal{M} - j + (n-k) + 1$ where $\mathcal{M} = \min\{j, k\}$. Let $\mathcal{L}(j)$ be the average data packet loss rate given that there are j lost packets in a FEC group. We approximate $\mathcal{L}(j)$ by

$$\begin{aligned}\mathcal{L}(j) &= \frac{1}{\mathcal{W}} \sum_{i=j-(n-k)}^{\mathcal{M}} \frac{i}{k} \\ &= \left[\frac{1}{\mathcal{M} - j + (n-k) + 1}\right] \times \left[\frac{1}{k}\right] \times \\ &\qquad \left[\frac{\mathcal{M}(\mathcal{M}+1)}{2} - \frac{(j-(n-k))(j-(n-k)-1)}{2}\right]. \end{aligned} \tag{22}$$

The approximation comes from the assumption that the probability of distributing j lost packets among the packets within a FEC group is equally likely. It is now easy to derive $\mathcal{L}_{sp}$, the average data packet loss rate (with the use of FEC) for the SP approach as follows:

$$\begin{aligned}\mathcal{L}_{sp} &= \sum_{j=n-k+1}^{n} P^{(1)}(j,n)\mathcal{L}(j) \\ &= \sum_{j=n-k+1}^{n} \left[P_0^{(1)}(j,n)\pi_0(1) + P_1^{(1)}(j,n)\pi_1(1)\right] \mathcal{L}(j) \\ &= \sum_{j=n-k+1}^{n} \left[\left(L_0(j,n) + H_0(j,n)\right)\left(\frac{\mu_1(1)}{\mu_0(1)+\mu_1(1)}\right)\mathcal{L}(j) + \right. \\ &\qquad \left. \left(L_1^{(1)}(j,n) + H_1^{(1)}(j,n)\right)\left(\frac{\mu_0(1)}{\mu_0(1)+\mu_1(1)}\right)\mathcal{L}(j)\right]. \end{aligned} \tag{23}$$

To derive the average data packet loss rate (with the use of FEC) for the MP approach, let us first consider a simple case of dual-path streaming. Assume that there are two servers S_1 and S_2 that use two different, possibly heterogeneous, paths. We use the same FEC scheme as described above to generate a stream of data divided into n packet FEC groups. To transmit the packets within a FEC group, server S_1 transmits n_1 packets while server S_2 transmits n_2 packets such that $n_1 + n_2 = n$. Based on a similar argument we made above in the SP case, we have

$$P^{(1)}(j,n_1) = P_0^{(1)}(j,n_1)\pi_0(1) + P_1^{(1)}(j,n_1)\pi_1(1) \qquad j = 0,1,\ldots,n_1 \tag{24}$$

$$P^{(2)}(j,n_2) = P_0^{(2)}(j,n_2)\pi_0(2) + P_1^{(2)}(j,n_2)\pi_1(2) \qquad j = 0,1,\ldots,n_2. \tag{25}$$

The computation of $P_i^{(h)}(j,n_h)$ where $i \in \{0,1\}$ and $h \in \{1,2\}$ is similar to the approach mentioned above, that is, it is done by evaluating the entries of the corresponding four matrices. The computational complexity would then be $\Theta(4(n_1+1)^2 + 4(n_2+1)^2)$.

Let P_{2p} be the probability of an irrecoverable error within a FEC group. It is equal to

$$P_{2p} = \sum_{j=n-k+1}^{n} \sum_{h=0}^{j} P^{(1)}(h,n_1)P^{(2)}(j-h,n_2), \tag{26}$$

which involves a convolution between the two probability mass functions, $P^{(1)}(j,n_1)$ and $P^{(2)}(j,n_2)$. Let $\mathcal{L}_{2p}$ be the average data packet loss rate (with the use of FEC) for the dual path approach. Then, we have

$$\mathcal{L}_{2p} = \sum_{j=n-k+1}^{n} \sum_{h=0}^{j} P^{(1)}(h,n_1)P^{(2)}(j-h,n_2)\mathcal{L}(j). \tag{27}$$

In general, if we employ N servers $S_1, S_2, \ldots, S_N$, then the probability of an irrecoverable error within a FEC group is

$$P_{Np} = \sum_{j=n-k+1}^{n} \left(\sum_{i_1+\ldots+i_N=j} P^{(1)}(i_1,n_1)P^{(2)}(i_2,n_2)\cdots P^{(N)}(i_N,n_1) \right). \tag{28}$$

The average data packet loss rate with FEC under MP streaming with N paths is

$$\mathcal{L}_{Np} = \sum_{j=n-k+1}^{n} \left(\sum_{i_1+\ldots+i_N=j} P^{(1)}(i_1,n_1)P^{(2)}(i_2,n_2)\cdots P^{(N)}(i_N,n_1) \right) \mathcal{L}(j). \tag{29}$$

In the case of the other two performance measures, namely, the conditional burst length distribution and the lag-1 autocorrelation, we resort to the use of simulation, as described in the following section.

3 Analytical Model Based Evaluation

In this section, we further evaluate the loss characteristics of the SP vs. MP methods using simulations of the Gilbert model described in Section 2. The simulations allow us to consider the loss characteristics under more sophisticated scenarios than in Section 2. (Due to lack of space, we only present a small set of experiments; more detailed results can be found in [6].) Specifically, we assume an MPEG-1 video streaming application which generates packets at a rate of 120 packets per second with each packet containing 1400 bytes. We consider at most three senders (S_1, S_2, S_3) and one receiver C. Sender S_i uses path i to transmit its fraction of the data; unless otherwise stated, these paths are assumed to be independent. Moreover, in the figures given below (unless otherwise stated), the curves corresponding to SP streaming use path 1, the curves corresponding to MP streaming with 2 senders use paths 1 and 2, and the curves corresponding to MP streaming with 3 senders use all three paths. Unless stated otherwise, the packet assignment is carried out in a round-robin manner, e.g., if we use all three senders, then sender S_i transmits data packets at a rate of 40 packets per second. The loss process of path i is modeled by a continuous stationary Gilbert model (as defined in Section 2). Unless stated otherwise, we use $\mu_0(i) = 20$ and $\mu_1(i) = 70$, for $i = 1, 2, 3$.

Experiment 1 (Data Loss Rate): In this experiment, we study the data packet loss rate of the SP and MP approaches, using only two paths, 1 and 2. The path parameters are as described above except that we vary the $\mu_0(2)$ parameter from 5 to 50. Table 1 illustrates the data loss rate for the single path(s) and the dual-path approaches (in each case, with and without the use of FEC, where the parameters for the FEC scheme are $n = 5$ and $k = 4$). We can observe that in this experiment:

- Without the use of FEC, the data packet loss rate of the dual path is approximately the mean of the data packet loss rates of paths 1 and 2. These results are consistent with the derivation of Section 2.
- With the use of FEC, (in this case $n = 5$ and $k = 4$), the achieved data packet loss rate can be *less* than the average of the data packet loss rates of the two corresponding single paths. This may occur due to the fact that error burst lengths in dual-path streaming tend to be shorter than in single-path streaming (refer Theorem 2 in Section 2), and hence a chance of recovery of lost data (using FEC) should also be higher.

This experiment also illustrates the potential advantages of multi-path streaming over "best path" streaming, even when losses (rather than throughput) are the important consideration. That is, when multiple paths are available (but throughput is not the issue), another approach might be to stream the data over the "best" available path (and as congestion conditions change keep switching the streaming of the data to the best available path at the time). Our experiment shows that MP streaming could provide better loss characteristics (e.g., when FEC is used) than the "best" available path (refer to [6] for further comparison to a best-path type approach).

Loss rate: ($\mu_0(2)$)	single path: path 1 w/o FEC	single path: path 2 w/o FEC	dual-path without FEC	single path: path 1 with FEC	single path: path 2 with FEC	dual-path with FEC
5	0.221743	0.066767	0.144351	0.189053	0.053048	0.101264
15	0.221743	0.176153	0.199395	0.189053	0.147171	0.141632
20	0.221743	0.221743	0.222255	0.189053	0.189053	0.158861
35	0.221743	0.332848	0.278178	0.189053	0.297647	0.201947
50	0.221743	0.416609	0.319230	0.189053	0.385602	0.235681

Table 1
Data Loss rate with Heterogeneous Paths.

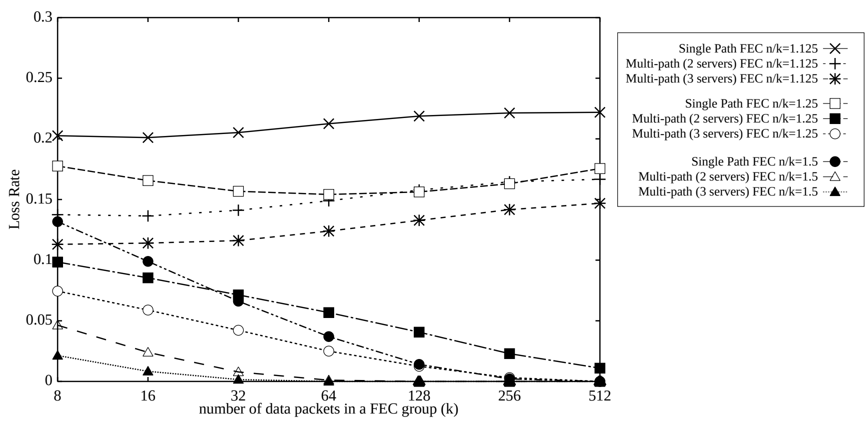

Fig. 2. Loss rate as a function of n/k and k

Experiment 2 (Data Loss Rate as a function of FEC parameters): In this experiment, we study the effects of FEC parameters on the data loss rate. In general, there are two ways to vary the FEC parameters. We can:

(1) Increase the degree of redundancy (e.g., for a given value of k, increase the value of n). Note that by increasing the degree of redundancy, we also increase the amount of traffic on the network.
(2) Increase the values of n and k but keep the same ratio of n/k. This implies that we increase the FEC group size, and hence the application needs to maintain a larger receiving buffer (for reconstruction purposes in case of loss) as well as experience potentially higher latency (since a larger amount of information must be received prior to reconstruction of missing information).

Figure 2 illustrates the effects of FEC parameters on the data loss rate, and specifically, it depicts data loss rates for SP and MP streaming with $n/k = 1.125, 1.25$ and 1.5 as well as with different FEC group sizes (where we vary the number of data packets in a FEC group (k) from 8 to 512 packets). In this case the path parameters are $\mu_0(1) = 20$, $\mu_1(1) = 70$, $\mu_0(2) = \mu_0(3) = 10$, and $\mu_1(2) = \mu_1(3) = 80$. We observe that:

- Increasing the amount of redundancy (e.g., from $n/k = 1.125$ to 1.5) in SP or MP streaming can reduce the data loss rate. However, one can achieve a lower data packet loss rate with MP streaming with a smaller n/k ratio (as compared to SP streaming). In other words, without introducing additional network traffic, we can obtain better performance with MP streaming.
- Increasing the number of data packets in a FEC group (while keeping the same ratio of n/k) may not necessary reduce the data loss rate. For example, consider SP streaming; as we increase k, the data loss rate actually increases in some cases. Similar observations have been made in other works, e.g., in [1] for an Internet-telephony application. One might consider addressing this problem by adapting the characteristics of the FEC scheme, e.g., as in [1], as well as by considering more efficient FEC schemes, e.g., as in [5].

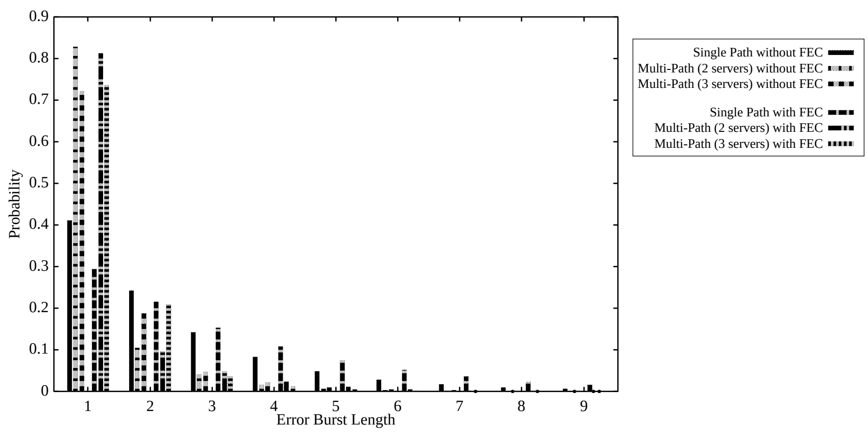

Fig. 3. Conditional probability mass functions of error burst length.

Experiment 3 (Conditional Error Burst Length): In this experiment, we compare the conditional burst length distribution, conditioned on there being at least one error. Figure 3 illustrates the conditional probability mass functions of error burst length (as defined in Section 2). In this experiment, we observe that the packet error burst length is indeed stochastically less than the error burst length of the single path streaming. We also note, that the condition of Theorem 2 in Section 2 holds in this experiment. This relationship also holds when we employ FEC.

Experiment 4 (Lag-1 Autocorrelation): In this experiment, we study the lag-1 autocorrelation of packet losses for both SP and MP streaming (as defined in Section 2). Figure 4 illustrates the lag-1 autocorrelation where $\mu_1(1) = \mu_1(2) = \mu_1(3) = 70$ and $\mu_0(i)$ is varied (identically) for all three paths.

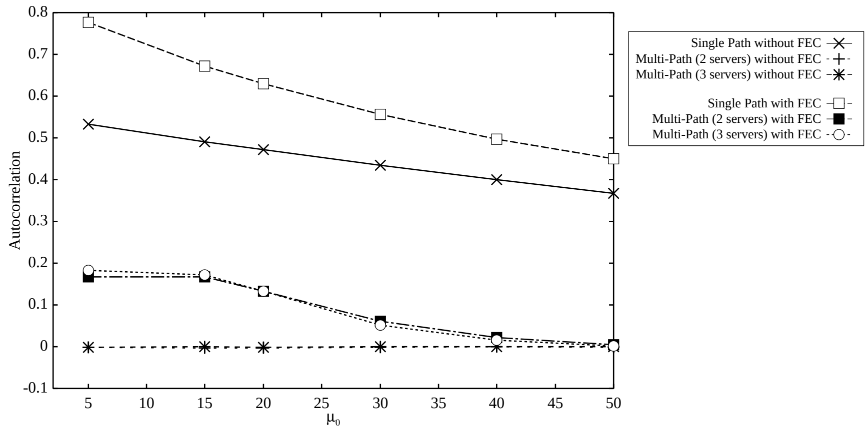

Fig. 4. Lag-1 autocorrelation.

We make the following observations.

- When we use MP streaming without FEC, the lag-1 autocorrelation is nearly zero while the lag-1 autocorrelation of SP path streaming (with or without FEC) can be highly correlated.
- The use of FEC may increase the lag-1 autocorrelation (for both approaches). This may be explained as follows. The irrecoverable losses (after the error correction process) are likely to end up "closer" in the resulting data stream than in the original data stream (one without the use of erasure codes), and hence the lag-1 autocorrelation in this new stream behaves similarly to lag-h autocorrelation of the original stream, where $h > 1$. However, we still observe that the lag-1 autocorrelation of MP streaming is significantly closer to zero as compared to SP streaming, even with the use of FEC.

4 Other Considerations and Related Work

In this section we first briefly discuss some of the issues that should be explored when considering the use of MP streaming. We then survey related work.

Considerations in Use of MP Streaming.
We note that one should also consider the potential costs or detrimental effects of multi-path streaming. For instance, MP streaming might have an adverse effect on the resulting delay characteristics observed at the receiver. As a result, it might also require a large amount of receiver buffer space. We also note that one might consider interleaving of packets during the transmission, as another approach to achieving a similar effect of a reduction in loss correlations. However, in this case, one should also consider the potential adverse effects of interleaving on delay characteristics observed at the receiver.

In addition, the overheads associated with sending data over multiple paths and then assembling it into a single stream at the receiver should also be considered. Moreover, the overheads and complexity due to measurements needed to achieve better performance with MP streaming should also be considered. For instance, in our case, we employ detection of shared points of congestion [18] to improve the performance of our MP streaming system. Other approaches to MP streaming might require even more detailed information about the network (see below) which is likely to result in a need for more "intrusive" and complex measurements. Furthermore, scalability of such measurement schemes is an issue as well.

We also note that if the bottleneck points on each path are at the client, then an MP streaming approach, of course, will not improve the situation as all servers will have a joint point of congestion. However, in wide-area networks, bottlenecks can also occur: (a) closer to the servers (e.g., due to high demand for service) and (b) at the peering points which are often cited as points of congestion in wide-area data transfers [8] and, in fact, as an important motivation for services like Akamai [8]. Moreover, even if Akamai-type servers are widely deployed but are not in a particular client's last hop ISP, there are still potential benefits to an MP streaming approach.

Lastly, when multiple paths are available (but throughput is not the issue), another approach might be to stream the data over the "best" available path (e.g., by using application-level re-routing techniques) and as congestion conditions change, potentially keep switching the streaming of the data to the best available path at the time. A comparison between simultaneously exploiting multiple paths and application-level re-routing, although in the context of large scale data transfers rather than streaming, is given in [3]. Furthermore, our experiments (refer to [6] for details) also show that MP streaming could provide better loss characteristics than the "best" available path.

Related Work.
We now give a brief survey of existing work on this topic, and specifically, we focus on those that either consider loss characteristics or can be deployed over best-effort networks (as these are considerations in our work as well). Earlier efforts on dealing with losses through the use of multiple independent paths (although at lower layers of the network) include dispersity routing, as proposed by Maxemchuk [11–13]. The focus in this work was on reducing delay, which includes reducing the number of retries needed to deliver a message without error, by sending the pieces of the data over multiple independent paths. Addition of redundant information is also possible under such a scheme. An important difference in our work is that we focus on streaming applications where the data transmission rate is determined by the application's needs rather than on delivering the data to its destination as fast as possible. Hence, in our case the data is sent through the network at a specific rate and that has an effect on loss characteristics. Also, we do not consider retransmissions

due to the real-time nature of the CM applications.

The use of multiple paths in routing data has of course been considered at the network layer, although not generally done in the current Internet. Hence, higher layer mechanisms should be considered. Another set of works on the topic considers higher level mechanisms, but requires some assistance from the lower layers and/or assumes significant knowledge of network topology and/or link capacities and delays (on all links used for data delivery). Given such knowledge, algorithms are proposed for selecting paths which can avoid congested routes. For instance, in [4], the authors focus on adaptation of delivery rate along the different paths, based on losses observed at the receiver. And, [2] considers proper scheduling of the initial portion of the video so as to reduce the start-up delay. In contrast, our approach does not rely on specific knowledge of topologies, capacities, delays, etc., and only considers whether a set of paths do or do not share joint points of congestion, as can be detected at the end-hosts. Moreover, our focus in this paper is on characterizing the benefits, with respect to loss characteristics, of a multi-path approach as compared to a single path approach. Hence, our interest is in the more basic understanding of this problem.

Recent literature on this topic also includes works on voice-over-IP type applications. For instance, [10,9] proposes a scheme for real-time audio transmission using multiple independent paths between a single sender and a single receiver, where multiple description coding (MDC) is used in multi-path delivery and a FEC approach is used in single-path delivery. These approaches are evaluated through simulation and experiments. In contrast, we believe that it is important to understand the effects of multi-path delivery on loss characteristics, even without the use of coding techniques. Hence, a great deal of our paper focuses on that. We also note that “live” applications (such as voice-over-IP) have different characteristics than pre-recorded applications (as we are considering here). For instance, one such difference is the need to disperse data in real-time, whereas in our case, we can distribute it to the multiple senders ahead of time; this makes our application-level implementation simpler and possibly more efficient. Another difference might be the ability to address the potentially adverse effects of MP streaming on delay characteristics (as mentioned above).

The recent work in [16] is closest to ours in that it also considers delivery of pre-recorded video from multiple senders distributed across the network. However, this work focuses on a transport protocol as well as on optimization algorithms for rate and packet distribution among the paths, with the objective of minimizing the loss rate at the receiver. In [15] FEC techniques are added (as compared to [16]), where distribution algorithms are considered but with the objective of minimizing the probability of an irrecoverable error. In contrast, due to the nature of the application, we believe that it is important to consider loss characteristics even when the losses cannot be fully recovered. That is, since we are considering delivery of video (which can be displayed even under some losses), it is important to consider other metrics. Hence, in our evaluation of potential benefits of multi-path streaming, we consider data loss rate, burst length distribution, as well as lag-1 autocorrelation (all with and without the use of FEC).

5 Conclusions

In this paper we investigated the potential benefits of an application-layer multi-path streaming approach to providing QoS over best-effort wide-area networks. As already mentioned, an advantage of this approach (as compared to approaches that require support of lower layers) is that the complexity of QoS provision can be pushed to the network edge and hence improve the scalability and deployment characteristics while at the same time provide a certain level of QoS guarantees. Our focus in this paper was on providing a fundamental understanding of the benefits of using multiple paths to deliver pre-recorded continuous media over best-effort wide-area networks, with loss characteristics being the main concern.

Our results indicate that in general, multi-path streaming exhibits better loss characteristics than single-path streaming (with or without use of an erasure code), which should result in a higher viewing quality of the received continuous media. These results can be used in guiding the design of multi-path continuous media systems. Overall, we believe that these results are quite encouraging and warrant further study of multi-path streaming over wide-area networks.

Acknowledgements. The authors are grateful to Bill Cheng, Don Towsley, and Edmundo de Souza e Silva for fruitful discussions, as well as to the anonymous referees for their helpful and insightful comments. This research was supported in part by the NSF ANI-0070016 grant and by the RGC and the CUHK Mainline Research Grant.

References

[1] Jean-Chrysostome Bolot, Sacha Fosse-Parisis, and Don Towsley. Adaptive FEC-Based Error Control for Internet Telephony. In *INFOCOM*, 1999.

[2] J.-C. Chen and S.-H. Chan. Multipath routing for video unicast over bandwidth-limited networks. In *the IEEE Globecom*, San Antonio, Texas, November 2001.

[3] C.-F. Chou, Y.-C. Wan, W.C. Cheng, L. Golubchik, and S. Khuller. A performance study of a large-scale data collection problem. In *Proceedings of the International Workshop on Web Content Caching and Distribution (WCW)*, August 2002.

[4] H. Chu and K. Nahrstedt. Dynamic multi-path communication for video traffic. In *the Hawiian Intl. Conference on System Science*, Hawaii, January, 1997.

[5] Daniel Ratton Figueiredo and Edmundo de Souza e Silva. Efficient mechanisms for recovering voice packets in the internet. In *IEEE Globecom*, 1999.

[6] L. Golubchik, John C.S. Lui, T.F. Tung, Alix L.H. Chow, W.J. Lee, G. Franceschinis, and C. Anglano. Multi-path continuous media streaming: What are the benefits? Technical Report CS-TR-2002-01, Computer Science and Engineering Department, CUHK, 2002.

[7] Rafael C. Gonzalez and Richard E. Woods. *Digital Image Processing*. Addison Wesley, 1993.

[8] T. Leighton. *The Challenges of Delivering Content on the Internet*. Keynote address at the ACM SIGMETRICS 2001 Conference, Cambridge, Massachusetts, June 2001.

[9] Y. J. Liang, E. G. Steinbach, and B. Girod. Multi-stream voice over ip using packet path diversity. In *the IEEE Fourth Workshop on Multimedia Signal Processing*, Cannes, France, October 2001.

[10] Y. J. Liang, E. G. Steinbach, and B. Girod. Real-time voice communication over the internet using packet path diversity. In *the ACM Multimedia Conference*, Ottawa, Canada, September/October 2001.

[11] N. F. Maxemchuk. Dispersity routing. In *the IEEE International Conference on Communications*, San Francisco, California, June 1975.

[12] N. F. Maxemchuk. Dispersity routing in high-speed networks. In *the Workshop on Very High Speed Networks*, Greenbelt, Maryland, March 1991.

[13] N. F. Maxemchuk. Dispersity routing in an atm network. In *the IEEE INFOCOM*, San Francisco, California, March/April 1993.

[14] P.M. Morse. *Queues, Inventories, and Maintenance*. John Wiley, 1958.

[15] T. Nguyen and A. Zakhor. Distributed video streaming with forward error correction. In *the International Packetvideo Workshop*, Pittsburg, Pennsylvania, April 2002.

[16] T. Nguyen and A. Zakhor. Distributed video streaming over internet. In *the SPIE Conference on Multimedia Computing and Networking*, San Jose, California, January 2002.

[17] S.M. Ross. *Stochastic Processes*. John Wiley, 1996.

[18] D. Rubenstein, J. Kurose, and D. Towsley. Detecting shared congestion of flows via end-to-end measurement. In *the ACM Sigmetrics Conference*, Santa Clara, California, June, 2000.