



Multi-path streaming: Optimization of load distribution

A.L.H. Chow^{a,*}, L. Golubchik^b, J.C.S. Lui^c, W.-J. Lee^d

^a *Computer Science Department, University of Southern California, Los Angeles, CA 90089, United States*

^b *CS&EE-Systems Departments, IMSC, ISI, University of Southern California, Los Angeles, CA 90089, United States*

^c *Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, PR China*

^d *Department of Computer Science, University of Maryland, College Park, MD 20742, United States*

Available online 10 August 2005

Abstract

Under the current Internet infrastructure, quality of service (QoS) in the delivery of continuous media (CM) is still relatively poor and inconsistent. In this paper we consider providing QoS through the exploitation of multiple paths existing in the network. Previous work has illustrated the advantages of this approach. Here we extend this work by considering a more expressive model for characterizing the network path losses. In particular, we propose a variation on the Gilbert model wherein the loss characteristics of a path depend on an application's transmission bandwidth. Using this model, we show the benefits of multi-path streaming over best single-path streaming, under optimal load distribution among the multiple paths. We use extensive simulation and measurements from a system prototype to quantify the performance benefits of our techniques.

© 2005 Elsevier B.V. All rights reserved.

Keywords: Multi-path streaming; Path diversity; Optimization; Path loss characterization; Modeling; Internet-based streaming

1. Introduction

Under the current Internet infrastructure, quality of service (QoS) in delivery of continuous media (CM) is still relatively poor and inconsistent. Degradation in quality of CM applications, involving delivery of video and audio, is partly due to variations in delays and losses experienced by packets sent through

* Corresponding author.

E-mail address: lhchow@cs.usc.edu (A.L.H. Chow).

wide-area networks. Although many such applications can tolerate some degree of missing information, significant losses degrade an application's QoS. One approach to providing QoS for CM 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, in this work we consider a more deployable approach to providing QoS guarantees through the exploitation of *multiple paths* existing in the network between a set of senders and a receiver, i.e., the CM data is fragmented into packets and the different packets take alternate routes to the receiver.¹ Previous work (discussed below) illustrated advantages of this approach in pushing the complexity of QoS provision to the network edge and hence improving the scalability and deployment characteristics while at the same time providing a certain QoS level.

As stated in [9] a number of issues need to be considered in designing a multi-path streaming system, but we limit the scope of this paper by focusing on (a) delivery of *pre-stored* video, e.g., as in video-on-demand applications (in contrast to delivery of “live” data² as in video-conferencing applications), (b) application-level schemes (which are deployable today over the current Internet), (c) accomplishment of multiple paths to the same receiver by *distributing servers* across wide-area networks and streaming data from multiple senders simultaneously, and (d) network related streaming issues *only* (rather than, e.g., considering server-related problems such as server load balancing), i.e., 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 this paper, we consider a system where server i sends fraction α_i of the data expected by the receiver, where $0 \leq \alpha_i \leq 1$ and $\sum_i \alpha_i = 1$. The receiver assembles the data from multiple senders and plays it in the appropriate order. Note that the multi-path approach we consider here injects the *same* total amount of data as the single path approach, when no error erasure codes are used. For example, if an application's bandwidth requirement is 1.5 Mbps, then either approach will inject traffic into the network at the total rate of 1.5 Mbps—the difference is that a single path approach injects traffic at that rate from a single server over a single path and a multi-path approach does it from multiple servers over multiple paths, where each path carries a fraction of that traffic. When error erasure codes are used, the total amount of data that the multi-path approach will inject into the network may be less than or equal to that of the single path approach and yet achieve the same or higher viewing quality, as illustrated in the remainder of the paper.

Multi-path streaming and exploitation of path diversity has attracted much attention recently, and [3] provides a broad overview of the general area. Due to lack of space, here we give only a brief survey of existing work on this topic focusing on those works which consider loss characteristics, or can be deployed over best-effort networks, or do not rely on specific coding techniques, as these are considerations in our work as well.

The work in [15] proposes the use of multi-servers for streaming on the Internet; this is later extended to include the use of erasure codes (EC) [14]. The focus is on designing a receiver-driven transport protocol which includes a rate allocation algorithm and a packet partitioning algorithm. The case of a last mile bottleneck is examined in [13]. In these proposals, a (what we call below) conventional Gilbert

¹ Such paths do not have to be completely disjoint; it is sufficient for them to have disjoint points of congestion, which can be detected, e.g., as in [17].

² For “live” data, it may be necessary to use a collection of relay hosts or proxies to “force” paths different from those provided by the network.

model is adopted (refer to Section 2) while a more expressive model is used in our work. In addition to focusing on the loss rate, we also propose an optimization approach using other loss characteristics, e.g., lag-1 autocorrelation. The use path diversity (as well erasure codes) on an overlay network as a way of emulating multiple sources is examined in [16]. Our approach can also be implemented on an overlay structure as well as with the aid of relay nodes, e.g., a P2P prototype implementation of our approach is described in [2].

Other works try to achieve multi-path streaming with the assistance from the lower layers, e.g., [8] utilizes bandwidth on multiple paths and re-distributes workload according to congestion detected by a receiver. But, it requires network layer knowledge and a centralized routing server. Given information about the underlying network graph, [6] proposes multi-path routing heuristics for unicast and multicast scenarios and a data scheduling algorithm at the server. Similarly, given network link information, [4] discusses a heuristic for finding a set of paths that minimizes the streaming distortion for a Multiple Description coded stream. These works assume significant knowledge (i.e., link bandwidth and delay) and support (i.e., ability to control routing paths) from the underlying network. In contrast, our approach only deals with end-hosts and hence allows easy deployment on the Internet. Our only requirement is that chosen paths do not share points of congestion, which can be detected at the end-hosts using schemes such as [17]. Also, we focus on packet loss characteristics rather than bandwidth and delay.

This paper extends the work presented in [9] where we illustrate the potential benefits of using multi-path streaming to improve the quality of the delivered CM as compared to single path streaming. This is done by illustrating lower loss burst lengths and lower correlations in consecutive packet losses. In particular, the loss characteristics on a network path are characterized by the Gilbert model (refer to Section 2)—here loss rates experienced at the receiver are independent of the sending rate on a path. We later [1] study the load distribution problem in multi-path streaming and show that both the packet loss rate and the loss correlations are important when choosing an optimization objective. However, these works are performed under the assumption that the application's sending rate does not affect the loss rate on a path. In contrast, in this work, we illustrate the utility of considering an application's sending rate and the resulting effects on the loss characteristics of the streaming application. Specifically, we consider a more general loss model than in [9] which takes into consideration effects of the data streaming rate on the loss characteristics of a path. We refer to this model as the *functional* Gilbert model (FGM); we refer to the original Gilbert model used in [9] as a conventional Gilbert model, which is a special case of the FGM. The motivation for considering the FGM is that it is potentially a more realistic representation of a network path's loss characteristics (as motivated further in Section 2).

In this paper, we use the FGM to address the following main questions: (a) what is an optimal splitting of traffic among the multiple paths, and (b) what resulting benefits can be expected. We also consider what gains can be expected when we employ erasure codes to further improve QoS. In considering these questions, throughout the paper (as in [9]) we use the following performance metrics: mean loss rate, lag-1 autocorrelation of consecutive losses, and burst length distribution of lost packets. (The importance of the latter two metrics in the resulting visual quality of CM is, e.g., described in [10].) We show that the conventional Gilbert model may not be sufficient to represent the loss characteristics of long duration streaming applications. We also present a methodology for optimizing load distribution (i.e., traffic splitting) among paths in the context of multi-path streaming. Our results illustrate the benefits of multi-path streaming over single-path streaming, under optimal splitting of traffic among the multiple paths.

2. Background and introduction of the functional Gilbert model

Previous works, e.g., [5], use a stationary continuous time Gilbert model (GM) to characterize potential correlations between consecutive packet losses on a network path. For a GM, 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 the transmitted packet is received correctly by the receiver; the transmitted packet is considered lost if $X_k(t) = 1$. Let $\mu_0(k)$ be the transition rate from state 0 to 1 and $\mu_1(k)$ be the transition rate from state 1 to 0. The stationary distribution of this GM is $\pi(k) = [\pi_0(k), \pi_1(k)]$ 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., $p_{i,j}^{(k)}(\tau) = P(X_k(t + \tau) = j | X_k(t) = i)$. Then, $p_{i,j}^{(k)}(\tau)$ can be computed for all $\tau > 0$ from previous work (refer, e.g., to [9]).

Using a GM to characterize the loss process of a path, [9] studies and compares the performance of streaming pre-stored CM data under single path and multi-path settings. In particular, the performance metrics considered are as follows. Firstly, *loss rate*, P_n , is the fraction of lost packets as seen by the receiver when one uses $n \geq 1$ paths for CM streaming. Secondly, *lag-1 autocorrelation function*, $R[\mathcal{X}_t, \mathcal{X}_{t+\delta}]$, measures the degree of dependency of consecutive packet losses as seen by the receiver, where \mathcal{X}_t is a random variable indicating whether the packet sent at time t is lost or received properly (depending on the state of the GM) and $1/\delta$ is the bandwidth requirement (in units of packets/second) of the streaming application.³ Thirdly, *burst length of lost packets* is the probability mass function of consecutively lost packets as seen by the receiver. Note that if the lost packets burst length is large, it can (a) significantly affect the viewing quality of the CM object and (b) reduce the effectiveness of an error correction scheme, if some form of an erasure code is deployed. Moreover, [9] illustrates that improvements in these metrics can be obtained by employing a multi-path streaming technique. However, these results were shown under the Gilbert model defined above and referred to as a “conventional Gilbert model” in the remainder of the paper. One major limitation of using a conventional Gilbert model is that the loss process of a path is *independent* of the bandwidth requirements of the streaming application.

To illustrate the potential dependence of the loss rate on an application’s bandwidth requirements, we carried out the following Internet experiment. We transmitted 1400 byte UDP packets from Hong Kong to the West Coast of the USA, using a number of rates from 120 pkts/s (around 1.34 Mbps) to 1200 pkts/s (around 12.8 Mbps), with a step size of 120 pkts/s interval. For each sending rate, the streaming experiment is carried out for 6 min, while measuring the corresponding achieved loss rate at the receiver. The experiment is carried out during daytime in Hong Kong, which corresponds to nighttime on the West Coast of the USA. Fig. 1 illustrates the achieved loss rate for each experimental setting, which is the *fraction* of lost packets as measured at the receiver.⁴ Fig. 1 supports our hypothesis that the conventional Gilbert model may not be sufficient for characterizing the loss process of a path. Since this evidence suggests that an application’s sending rate can significantly affect the achieved loss rate, we propose to use a *functional Gilbert model* (FGM) as a general approach to characterizing the bursty loss nature of a path as well as its dependency on an application’s bandwidth requirements.

³ A high positive value of $R[\mathcal{X}_t, \mathcal{X}_{t+\delta}]$ implies that a lost packet is very likely to be followed by another lost packet. A high negative value of $R[\mathcal{X}_t, \mathcal{X}_{t+\delta}]$ implies that a lost packet is likely to be followed by a successful packet arrival. If the statistics of the consecutive packet losses are not correlated, then $R[\mathcal{X}_t, \mathcal{X}_{t+\delta}] = 0$.

⁴ Similar experiments using NS2 [11] gave qualitatively similar results.

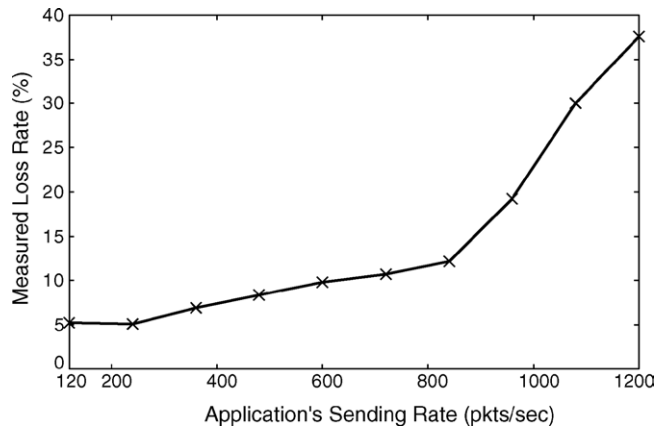


Fig. 1. Loss rate vs. sending rate.




Fig. 2. Functional Gilbert model.

Let λ denote an application’s average sending rate, in units of packets/second. For a stationary continuous time FGM, the packet loss process along path k is described by a two state continuous time Markov chain $\{X_k(t)\}$ where $X_k(t) \in \{0, 1\}$. Fig. 2 depicts the state transition diagram of this model. Similarly to the conventional Gilbert model’s definition, if a packet is transmitted at time t when the state of path k is $X_k(t) = 0$, then no packet loss occurs; the transmitted packet is considered lost if $X_k(t) = 1$. The transition rate⁵ from state 0 to 1 takes a *functional* form of $\mathcal{F}_k(\lambda)$. The transition rate from state 1 to 0 also takes a *functional* form of $\mathcal{B}_k(\lambda)$. In this paper, we assume that $\mathcal{F}_k(\lambda)$ and $\mathcal{B}_k(\lambda)$ are continuous and furthermore that $\mathcal{F}_k(\lambda)$ is a *non-decreasing* function of λ and $\mathcal{B}_k(\lambda)$ is a *non-increasing* function of λ . We note that intuitively these assumptions make sense, and hence, in practice, they should not be restrictive. When multiple paths have the same functions in the FGMs (i.e., $\mathcal{F}_i(\lambda) = \mathcal{F}_j(\lambda)$ and $\mathcal{B}_i(\lambda) = \mathcal{B}_j(\lambda)$, for all $i, j \leq M$ where M is the number of paths in the system), we say that they are *homogeneous* paths; otherwise they are *heterogeneous* paths. When one uses an FGM to characterize the loss process of a path, we have the following result.

Theorem 1. *Let there be $M \geq 1$ homogeneous paths available for CM streaming. Define $\alpha = [\alpha_1, \alpha_2, \dots, \alpha_M]$ as the vector which determines how the traffic is split among these M paths, where*

⁵ In what follows, we drop the path designation from the notation whenever it is clear from context, e.g., when paths are homogeneous.

$\alpha_i \geq 0$ and $\sum_i^M \alpha_i = 1$. The achieved loss rate via the multi-path streaming approach (P_M) is less than or equal to the achieved loss rate via the single path streaming approach (P_1), for all possible valid traffic splitting vectors α .

Proof. Let P_M be the achieved loss rate for the multi-path streaming approach, with all paths being *homogeneous*, i.e., $P_M = \sum_{j=1}^M \alpha_j \mathcal{F}(\alpha_j \lambda) / (\mathcal{F}(\alpha_j \lambda) + \mathcal{B}(\alpha_j \lambda))$. Since paths are homogeneous, the achieved loss rate under single path streaming is $P_1 = \mathcal{F}(\lambda) / (\mathcal{F}(\lambda) + \mathcal{B}(\lambda))$. Let path k^* be the path that has the largest loss rate, that is, $\mathcal{F}(\alpha_{k^*} \lambda) / (\mathcal{F}(\alpha_{k^*} \lambda) + \mathcal{B}(\alpha_{k^*} \lambda)) \geq \mathcal{F}(\alpha_j \lambda) / (\mathcal{F}(\alpha_j \lambda) + \mathcal{B}(\alpha_j \lambda))$ for all $j \in \{1, \dots, M\}$. Then we have $P_M \leq \sum_{j=1}^M \alpha_j \mathcal{F}(\alpha_{k^*} \lambda) / (\mathcal{F}(\alpha_{k^*} \lambda) + \mathcal{B}(\alpha_{k^*} \lambda)) = \mathcal{F}(\alpha_{k^*} \lambda) / (\mathcal{F}(\alpha_{k^*} \lambda) + \mathcal{B}(\alpha_{k^*} \lambda)) = P_M^*$. Note that $P_M^* - P_1 \leq 0$ implies that $P_M - P_1 \leq 0$. Since $P_M^* - P_1 = \mathcal{F}(\alpha_{k^*} \lambda) / (\mathcal{F}(\alpha_{k^*} \lambda) + \mathcal{B}(\alpha_{k^*} \lambda)) - \mathcal{F}(\lambda) / (\mathcal{F}(\lambda) + \mathcal{B}(\lambda))$, we need to show that $\mathcal{F}(\alpha_{k^*} \lambda) (\mathcal{F}(\lambda) + \mathcal{B}(\lambda)) - \mathcal{F}(\lambda) (\mathcal{F}(\alpha_{k^*} \lambda) + \mathcal{B}(\alpha_{k^*} \lambda)) \leq 0$, to prove that $P_M^* - P_1 \leq 0$ and hence $P_M - P_1 \leq 0$. Expanding the terms gives

$$\begin{aligned} & \mathcal{F}(\alpha_{k^*} \lambda) \mathcal{F}(\lambda) + \mathcal{F}(\alpha_{k^*} \lambda) \mathcal{B}(\lambda) - \mathcal{F}(\lambda) \mathcal{F}(\alpha_{k^*} \lambda) - \mathcal{F}(\lambda) \mathcal{B}(\alpha_{k^*} \lambda) \\ &= \mathcal{F}(\alpha_{k^*} \lambda) \mathcal{B}(\lambda) - \mathcal{F}(\lambda) \mathcal{B}(\alpha_{k^*} \lambda) \leq \mathcal{F}(\alpha_{k^*} \lambda) \mathcal{B}(\alpha_{k^*} \lambda) - \mathcal{F}(\lambda) \mathcal{B}(\alpha_{k^*} \lambda) \\ &= \mathcal{B}(\alpha_{k^*} \lambda) [\mathcal{F}(\alpha_{k^*} \lambda) - \mathcal{F}(\lambda)] \leq 0. \quad \square \end{aligned}$$

Remark. This theorem implies that under the homogeneous path assumption, *any* valid traffic splitting in a multi-path streaming approach will do no worse than single path streaming in terms of the packet loss rate metric. An interesting question to ask is what is the *right metric* to optimize in determining the traffic split among the M paths. We consider this and the resulting performance effects below.

3. Optimal traffic splitting

In this section, we present a framework for determining appropriate traffic splitting between the multiple paths used for CM streaming. We use the achieved loss rate and lag-1 autocorrelation (as described in Section 2) as our objective functions. Although, we consider a single performance metric at a time in this optimization process,⁶ in Section 4 we illustrate the effects of this optimization process on the other performance metrics.

3.1. Optimization based on achieved loss rate

We first consider the minimization of the loss rate, P_M , achieved at the receiver as our objective. For path j , let $\mathcal{F}_j(b)$ denote the functional transition rate from state 0 to 1 when the streaming traffic on path j is b pkts/s. Similarly, $\mathcal{B}_j(b)$ denotes the functional transition rate from state 1 to 0 when the streaming traffic on path j is b pkts/s. Let us first consider a simple case wherein there are two paths available for CM streaming. We define $F_j(\alpha_j \lambda) = \mathcal{F}_j(\alpha_j \lambda) / (\mathcal{F}_j(\alpha_j \lambda) + \mathcal{B}_j(\alpha_j \lambda))$, for $j = 1, 2$ and express the achieved loss rate as $P_2 = \alpha_1 F_1(\alpha_1 \lambda) + (1 - \alpha_1) F_2((1 - \alpha_1) \lambda)$. We can then state the following theorem for a 2-paths homogeneous system.

⁶ If multiple (equivalent) optimum solutions exist, secondary performance metrics can be used to “break the tie”, e.g., as described in Section 3.2.

Theorem 2. Under a 2-paths system, if the two paths are homogeneous, and $F_j(\alpha_j\lambda)$ is a convex non-decreasing function ($j = 1, 2$), then the traffic splitting vector $\alpha_{\text{even}} = [\frac{1}{2}, \frac{1}{2}]$ is optimal, i.e., it achieves the lowest loss rate.

Proof. Assume there exists an $\alpha' = [\alpha_1, \alpha_2]$, which results in the lowest loss rate and $\alpha_1 \neq \alpha_2$. Let $\alpha_1 < \frac{1}{2}$ and $\alpha_2 > \frac{1}{2}$. As paths are homogeneous, the achieved loss rate is $P_2(\alpha') = \alpha_1 F(\alpha_1\lambda) + \alpha_2 F(\alpha_2\lambda)$. Given that $F(\alpha_j\lambda)$ is a non-decreasing function, as well as $\alpha_1 < \frac{1}{2}$ and $\alpha_2 > \frac{1}{2}$, we have $P_2(\alpha') \geq 1/2 F(\alpha_1\lambda) + 1/2 F(\alpha_2\lambda)$. Due to the property of convex functions and $\alpha_1 + \alpha_2 = 1$, $P_2(\alpha') \geq F(\frac{1}{2}(\alpha_1\lambda + \alpha_2\lambda)) = F(\frac{1}{2}\lambda)$. It is easy to show that $F(\frac{1}{2}\lambda) = P_2(\alpha_{\text{even}})$, which is the achieved loss rate corresponding to α_{even} . Thus, we have $P_2(\alpha') \geq P_2(\alpha_{\text{even}})$. This implies that in a 2-paths homogeneous streaming system where $F(\alpha_j\lambda)$ is a convex non-decreasing function, the traffic splitting vector $\alpha = [\frac{1}{2}, \frac{1}{2}]$ is optimal, i.e., it achieves the lowest loss rate. \square

To illustrate the performance gains due to multi-path streaming, we consider a family of functional transition rates. In particular, we consider the FGM for path j , for $j = 1, \dots, M$, to have the form of

$$\mathcal{F}_j(b) = \beta_j b^{\theta_j} + S_j, \quad \mathcal{B}_j(b) = \frac{\kappa_j}{b^{\chi_j}} + \varphi_j \quad (1)$$

where $\beta_j, \kappa_j, \chi_j \geq 0$, and $\theta_j, S_j, \varphi_j > 0$.⁷ In other words, \mathcal{F}_j is a non-decreasing function while \mathcal{B}_j is a non-increasing function of the traffic bandwidth b . These forms can represent a large family of Gilbert models. For example, for a conventional Gilbert model, a constant transition rate from state 0 to 1 can be represented by setting $S_j > 0$, $\beta_j = 0$, and $\theta_j = 0$, and a constant transition rate from state 1 to 0 can be represented by setting $\kappa_j = 0$ and $\varphi_j > 0$. A linearly increasing function of $\mathcal{F}_j(b)$ can be represented by setting $\beta_j > 0$ and $\theta_j = 1$. A function $\mathcal{B}_j(b)$ which decreases as an inverse of the packet rate can be represented by setting $\kappa_j > 0$ and $\chi_j = 1$.

Consider now an application with an average bandwidth requirement of 120 pkts/s (e.g., a representative 1.28 Mbps MPEG 1 video stream with 1400 byte packets). We first consider a system with homogeneous paths having the following simple functions for the FGM: $\mathcal{F}(b) = 0.2333 \times b$ and $\mathcal{B}(b) = 24750/b$, as an illustration. One motivation for choosing simple forms for illustration is that they will be easier to extract from measurements in a real system. Fig. 3 depicts the loss rates for a system under two homogeneous paths as a function of α_1 (the corresponding α_2 is $1 - \alpha_1$). Fig. 4 illustrates the contour map of loss rates for three homogeneous paths (with the same parameters as the 2-paths system) at various values of α . Under this homogeneous paths example, the loss rate using single path streaming is 11.953%. Under the 2-paths streaming approach, the achieved loss rate is reduced to 3.283%, with $\alpha^* = [0.5, 0.5]$, and the corresponding lag-1 autocorrelation is 0. The 3-paths system reduces the loss rate further to 1.486%, with $\alpha^* = [1/3, 1/3, 1/3]$.

Now let us consider an example with the same bandwidth requirement of 120 pkts/s but under a heterogeneous paths setting. We first consider a system with two heterogeneous paths wherein $\mathcal{F}_1(b) = 0.4 \times b$, $\mathcal{B}_1(b) = 21000/b$, $\mathcal{F}_2(b) = 0.0667 \times b$, and $\mathcal{B}_2(b) = 28500/b$, as an illustration. Note that given the same packet rate on a path, path 2 has better loss characteristics than path 1. Fig. 5 depicts the loss rate as a function of α_1 (with $\alpha_2 = 1 - \alpha_1$). Fig. 6 depicts the contour map of loss rates at various values of α when path 3 is added with $\mathcal{F}_3(b) = 0.2333 \times b$ and $\mathcal{B}_3(b) = 24750/b$ (i.e., path 3 has loss characteristics

⁷ We also tried other forms, e.g., $\mathcal{F}_j(b) = \beta_j(b/\sigma_j)^{\theta_j} + S_j$, $\mathcal{B}_j(b) = \kappa_j e^{-\chi_j b} + \varphi_j$. The results are similar to those presented in this paper.




Fig. 3. Loss rate with two homogeneous paths.

in between paths 1 and 2). Under the *best* single path streaming approach (i.e., path 2 in this case), one can achieve a loss rate of 3.259%. Using 2-paths streaming, we reduce the loss rate to 1.814%, with $\alpha^* = [0.260, 0.740]$, and the corresponding lag-1 autocorrelation is 0.012. We reduce the loss rate further via 3-paths streaming to 0.976%, with $\alpha^* = [0.190, 0.541, 0.269]$. This example illustrates that the traffic splitting flexibility of the multi-path approach provides us an opportunity to reduce the mean loss rate of a streaming application to a point which would not be possible with a single *best-path* type approach. Note that the additional path which was not present in the 2-paths example is *not* the best of the three, yet

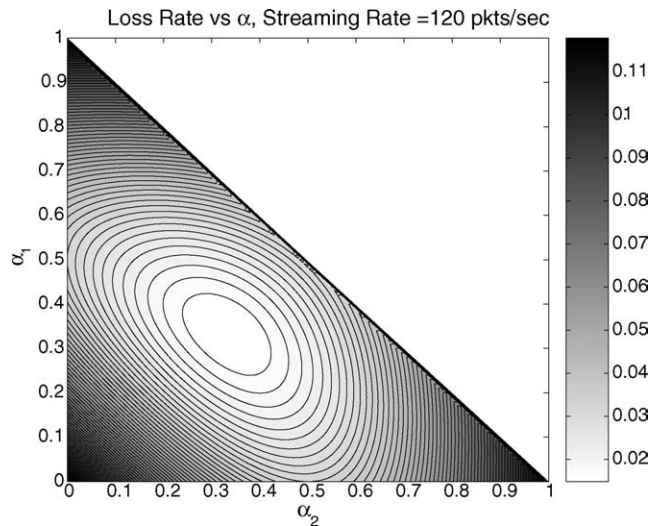


Fig. 4. Contour map of loss rate with three homogeneous paths.




Fig. 5. Loss rate with two heterogeneous paths.

it allows us to reduce the loss rate further. The above given examples illustrate the benefits of multi-path streaming when optimizing the achieved loss rate. We explore this in more detail in Section 4.

3.2. Optimization based on lag-1 autocorrelation

As explained in Section 2, metrics other than the loss rate can have a significant effect on the viewing quality of CM, e.g., the lag-1 autocorrelation. Hence, we now consider the problem of finding an appropriate traffic splitting when one wants to optimize the achieved lag-1 autocorrelation function, R_M ,

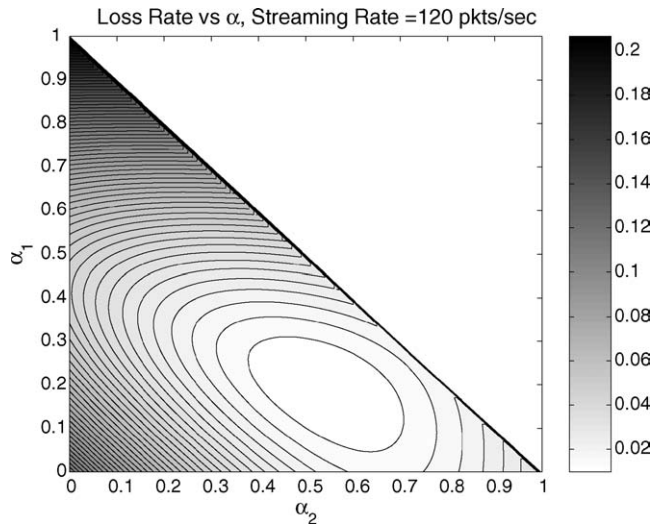


Fig. 6. Contour map of loss rate with three heterogeneous paths.

rather than the loss rate, P_M . As stated earlier, the lag-1 autocorrelation function, $R[\mathcal{X}_t \mathcal{X}_{t+\delta}]$, measures the degree of dependency of consecutive packet losses as seen by the receiver, where \mathcal{X}_t is a random variable indicating whether the packet sent at time t is lost or received properly (depending on the state of the FGM). It can be derived as

$$\begin{aligned} R[\mathcal{X}_t \mathcal{X}_{t+\delta}] &= \frac{E[(\mathcal{X}_t - \bar{\lambda})(\mathcal{X}_{t+\delta} - \bar{\lambda})]}{E[(\mathcal{X}_t - \bar{\lambda})^2]} = \frac{E[\mathcal{X}_t \mathcal{X}_{t+\delta}] - \bar{\lambda}^2}{E[\mathcal{X}_t^2] - \bar{\lambda}^2} = \frac{E[\mathcal{X}_t \mathcal{X}_{t+\delta}] - \bar{\lambda}^2}{\bar{\lambda} - \bar{\lambda}^2} \\ &= 1 - \frac{\bar{\lambda} - E[\mathcal{X}_t \mathcal{X}_{t+\delta}]}{\bar{\lambda} - \bar{\lambda}^2} \approx 1 - \frac{\text{Pr}[\text{burst}]}{\bar{\lambda}(1 - \bar{\lambda})} \end{aligned} \quad (2)$$

where $\text{Pr}[\text{burst}]$ is the probability of there being an error burst of any length.⁸ For instance, assuming equally spaced inter-arrival times on a path, we have the following for $M = 2$:

$$\begin{aligned} \text{Pr}[\text{burst}] &= \begin{cases} [\alpha_2 - \alpha_1]\pi_0(2)p_{0,1}^{(2)}(\delta_2) + \alpha_1\pi_0(1)\pi_1(2) + \alpha_1\pi_0(2)\pi_1(1) & \alpha_1 < \alpha_2 \\ \frac{1}{2}\pi_0(1)\pi_1(2) + \frac{1}{2}\pi_0(2)\pi_1(1) & \alpha_1 = \alpha_2 = \frac{1}{2} \\ [\alpha_1 - \alpha_2]\pi_0(1)p_{0,1}^{(1)}(\delta_1) + \alpha_2\pi_0(1)\pi_1(2) + \alpha_2\pi_0(2)\pi_1(1) & \alpha_1 > \alpha_2 \end{cases} \quad \text{and} \\ R[\mathcal{X}_t \mathcal{X}_{t+\delta}] &= \begin{cases} 1 - \frac{[\alpha_2 - \alpha_1]\pi_0(2)p_{0,1}^{(2)}(\delta_2) + \alpha_1\pi_0(1)\pi_1(2) + \alpha_1\pi_0(2)\pi_1(1)}{[\alpha_1\pi_0(1) + \alpha_2\pi_0(2)][\alpha_1\pi_1(1) + \alpha_2\pi_1(2)]} & \alpha_1 < \alpha_2 \\ 1 - \frac{\pi_0(1)\pi_1(2) + \pi_0(2)\pi_1(1)}{\frac{1}{2}[\pi_1(1) + \pi_1(2)][\pi_0(1) + \pi_0(2)]} & \alpha_1 = \alpha_2 = \frac{1}{2} \\ 1 - \frac{[\alpha_1 - \alpha_2]\pi_0(1)p_{0,1}^{(1)}(\delta_1) + \alpha_2\pi_0(1)\pi_1(2) + \alpha_2\pi_0(2)\pi_1(1)}{[\alpha_1\pi_0(1) + \alpha_2\pi_0(2)][\alpha_1\pi_1(1) + \alpha_2\pi_1(2)]} & \alpha_1 > \alpha_2 \end{cases} \end{aligned} \quad (3)$$

where δ_k denotes the time interval between two consecutively transmitted packets on path k (i.e., $\delta_k = \frac{1}{\alpha_k \lambda}$). Note that $\pi_0(k)$, $\pi_1(k)$ and $p_{i,j}^{(k)}$ are obtained from the FGM. (A detailed derivation is given in [7]; we omit it here due to lack of space.)

We now define an optimal traffic split as a split which minimizes the absolute value of the lag-1 autocorrelation function. The intuition here is that we aim to reduce correlations between packet losses, as explained in Section 2. Therefore, to find an optimal traffic split, $\alpha^* = [\alpha_1^*, 1 - \alpha_1^*]$, we equate the numerator in Eq. (3) to 0. This equation can be solved using standard numerical methods for finding roots. If more than one solution for α_1 satisfies this equation, then we can use a secondary objective. For instance, we can check the corresponding loss rates and choose the one with the lowest loss rate. Note that, when the two paths are homogeneous (i.e., $F_1(\cdot) = F_2(\cdot)$ in this case), then equal splitting of traffic along these two paths (i.e., $\alpha_1 = \alpha_2 = 1/2$) is a critical point in this optimization problem. This claim can be easily verified as follows:

$$1 - \frac{\pi_0(1)\pi_1(2) + \pi_0(2)\pi_1(1)}{\frac{1}{2}[\pi_1(1) + \pi_1(2)][\pi_0(1) + \pi_0(2)]} = \frac{-[\pi_1(1) - \pi_1(2)]^2}{[\pi_0(1) + \pi_0(2)][\pi_1(1) + \pi_1(2)]}$$

⁸ Essentially, $\text{Pr}[\text{burst}]$ is the probability of encountering the beginning of a burst (i.e., a no loss followed by a loss) when inspecting a stream at a random point.

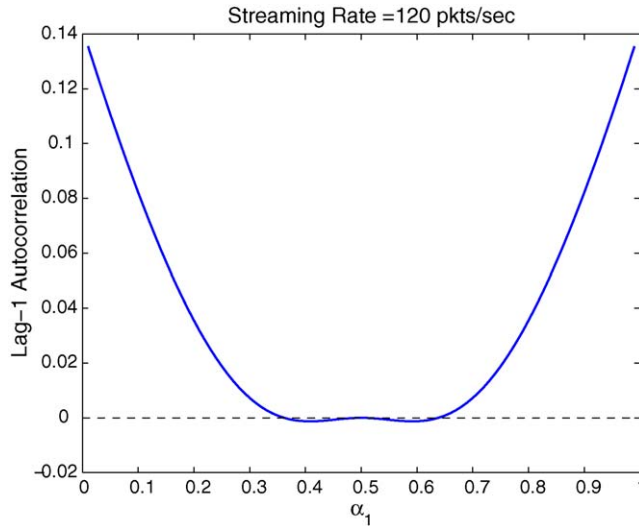


Fig. 7. Lag-1 autocorrelation with two homogeneous paths.

which equals to 0 when $\pi_1(1) = \pi_1(2)$, i.e., when $F_1(\alpha_1\lambda) = F_2(\alpha_2\lambda)$. As paths are homogeneous, this equation is satisfied when $\alpha_1 = \alpha_2 = 1/2$.

To illustrate the optimization based on lag-1 autocorrelation, we again consider paths characterized by the FGM with functional forms given in Eq. 1 and an application with a bandwidth requirement of 120 pkts/s. We first consider a system with homogeneous paths having the following simple functions for the FGM: $\mathcal{F}(b) = 0.2333 \times b$ and $\mathcal{B}(b) = 24750/b$, as an illustration. Fig. 7 depicts the corresponding lag-1 autocorrelation as a function of α_1 (the corresponding α_2 is $1 - \alpha_1$). In this example, the optimized lag-1 autocorrelation is 0, with $\alpha^* = [0.5, 0.5]$ and a corresponding loss rate of 3.283%.

Now let us consider an example with the same bandwidth requirement of 120 pkts/s but under a heterogeneous paths setting. Specifically, we consider a system with two heterogeneous paths wherein $\mathcal{F}_1(b) = 0.4 \times b$, $\mathcal{B}_1(b) = 21000/b$, $\mathcal{F}_2(b) = 0.0667 \times b$, and $\mathcal{B}_2(b) = 28500/b$. Fig. 8 illustrates the lag-1 autocorrelation as a function of α_1 (with $\alpha_2 = 1 - \alpha_1$). In this example, we can determine two α_1 's (one being smaller than 0.5 and the other being larger than 0.5) which both achieve lag-1 autocorrelation of zero. By comparing the corresponding loss rates, we choose $\alpha^* = [0.347, 0.653]$, corresponding to the loss rate of 2.035%. (The second optimal lag-1 autocorrelation point corresponds to $\alpha^* = [0.690, 0.310]$ and a loss rate of 8.06%.) We further explore the potential of optimizing traffic splitting based on lag-1 autocorrelation in Section 4.

Remark. One can, of course, consider other performance metrics as optimization objectives, e.g., the mean error burst length. One could also consider combinations of metrics, e.g., mean loss rate \times mean burst length which tries to encompass the importance of loss rate and correlations. Due to space limitations, further discussion of these objective functions and their derivation [7] are omitted. These are also explored in [1] but in the context of the conventional GM. We do note that, when erasure codes are added, one might want to minimize the conditional mean information lost rate (MILR), which is the loss rate of the media data after the lost packet reconstruction process. However, an analytical derivation of this metric is complex. In Section 4, we obtain an optimum MILR from simulations and compare it with results based on other optimization objectives.




Fig. 8. Lag-1 autocorrelation with two heterogeneous paths.

4. Simulation experiments

In this section, we consider two types of experiments: type A, where servers send only data packets to the receiver, and type B, where an erasure code is used to reconstruct lost packets as much as possible. For all experiments, each data packet has a size of 1400 bytes with a packet group size of $k = 1000$. For each path, packet losses are emulated according to the FGM. The results presented below are obtained through simulation using CSIM[12] with optimal load distributions (α^*) obtained using the analytical formulations presented in Section 3.⁹ Each simulation is analogous to a 24 h CM stream.

4.1. Type A experiments: streaming without an erasure code

We consider a system with two heterogeneous paths wherein $\mathcal{F}_1(b) = 0.4 \times b$, $\mathcal{B}_1(b) = 21000/b$, $\mathcal{F}_2(b) = 0.0667 \times b$, and $\mathcal{B}_2(b) = 28500/b$. Tables 1 and 2 illustrate various performance metrics, such as the optimal splitting vector α^* , the achieved loss rate, the achieved lag-1 autocorrelation as well as the system's performance when we stream the data using the *best* single path and a round-robin approach (i.e., evenly spreading the workload without performing optimization).

We also consider improvements in the various performance measures when a third path is added, wherein $\mathcal{F}_3(b) = 0.2333 \times b$ and $\mathcal{B}_3(b) = 24750/b$. Note that this additional path does not possess the best packet loss characteristics; therefore, it simply provides greater path diversity for the data transmission process. Table 3 illustrates the corresponding performance metrics¹⁰ and Figs. 9 and 10 illustrate the corresponding conditional lost packet burst length probability mass functions, conditioned on there being

⁹ α^* only gives the relative traffic loading ratios among the M paths; in a real system, α^* still needs to be mapped to a packet sending pattern. Due to lack of space, we omit the specifics of this “quantization” process as described in [7].

¹⁰ Corresponding best path results are given in Tables 1 and 2.

Table 1
2-Paths optimization based on loss rate

Rate (pkts/s)	α_1^*	α_2^*	Loss rate at α^* (%)	Lag-1 at α^*	Best SP loss rate (%)	RR loss rate (%)
60	0.260	0.740	0.462	0.000017	0.831	0.949
120	0.260	0.740	1.808	0.011906	3.258	3.632
360	0.260	0.740	14.261	0.300698	23.273	22.603

Table 2
2-Paths optimization based on lag-1 autocorrelation

Rate (pkts/s)	α_1^*	α_2^*	Loss Rate at α^* (%)	Lag-1 at α^* (%)	Best SP lag-1	RR lag-1
60	0.093	0.907	0.627	0.000128	0.000469	-0.006212
120	0.347	0.653	2.039	-0.000272	0.129062	-0.022340
360	0.550	0.415	26.094	-0.093527	0.750808	-0.137788

Table 3
3-Paths optimization based on loss rate

Rate (pkts/s)	α^*	Loss rate at α^* (%)	Lag-1 at α^*	RR loss rate (%)	RR lag-1
60	[0.188, 0.542, 0.270]	0.247	0.000104	0.416	-0.000350
120	[0.190, 0.541, 0.269]	0.975	0.000172	1.604	-0.003258
360	[0.190, 0.541, 0.269]	8.156	0.092063	12.237	-0.025815

a loss,¹¹ for application sending rates of 120 and 360 pkts/s, respectively.¹² This is illustrated for 2- and 3-paths streaming and under different traffic splitting methods, e.g., “Loss 2p” refers to using 2-paths with load distribution computed using the loss rate based optimization method, “RR 3p” refers to using 3-paths with the round-robin approach, “Best SP” refers to using the best single path, etc.

In both experiments we observe that loss rate-based optimization can significantly reduce the packet loss rate, but it has a higher lag-1 autocorrelation than the lag-1 autocorrelation-based optimization method. Also, the lost packet burst length probability mass function under the lag-1 autocorrelation-based optimization method as well as under the round-robin approach is more skewed toward single packet losses—this should improve the visual quality of CM. When we move from a 2-paths system to a 3-paths system, in most cases, we observe that there is an improvement in terms of loss rate, lag-1 autocorrelation, and lost packet burst length probability mass function, when one uses the loss rate-based optimization method. Overall, we observe that both optimization methods result in better system performance (under several metrics) than best-path streaming or the round-robin approach.

4.2. Type B experiments: streaming with an erasure code

We consider the effects of an erasure code on the various performance measures. Since numerous erasure coding schemes exist we first give a brief explanation of the erasure code used here. We divide a

¹¹ We present the probability mass function rather than the probability distribution function, as we believe it depicts the results of the experiments better.

¹² When the sending rate is 60 pkts/s, almost all lost packet bursts are of length one under any of the mentioned traffic splitting methods; thus the corresponding figure is omitted.

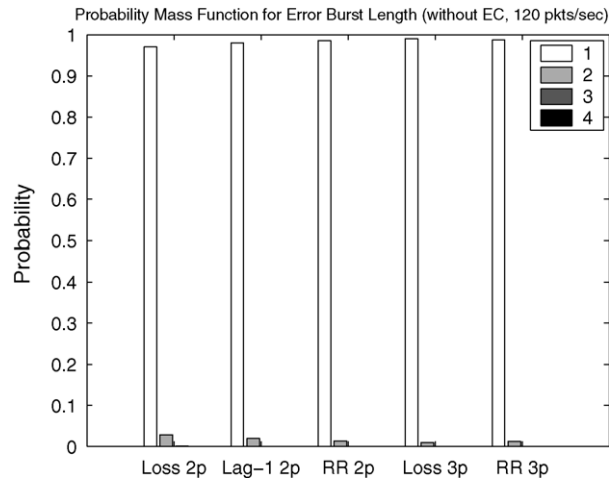


Fig. 9. Lost packet burst length probability mass functions (120 pkts/s).

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 an erasure code (EC) group. The encoding scheme is such that, if the number of lost packets within an EC group is less than or equal to $(n - k)$, then we can reconstruct the original k data packets within that EC group. In the following experiment, we consider an EC with parameters $k = 64$ and $n = 72$, i.e., the bandwidth requirements of the streaming application are increased by 12.5%. Note that this overhead is the same for single path and multi-path streaming. Optimal load distributions are obtained using the analysis in Section 3 but with an increased packet sending rate (due to the overhead).

Similarly to type A experiments, we first consider two heterogeneous paths wherein $\mathcal{F}_1(b) = 0.4 \times b$, $\mathcal{B}_1(b) = 21000/b$, $\mathcal{F}_2(b) = 0.0667 \times b$, and $\mathcal{B}_2(b) = 28500/b$. Tables 4 and 5 illustrate various perfor-

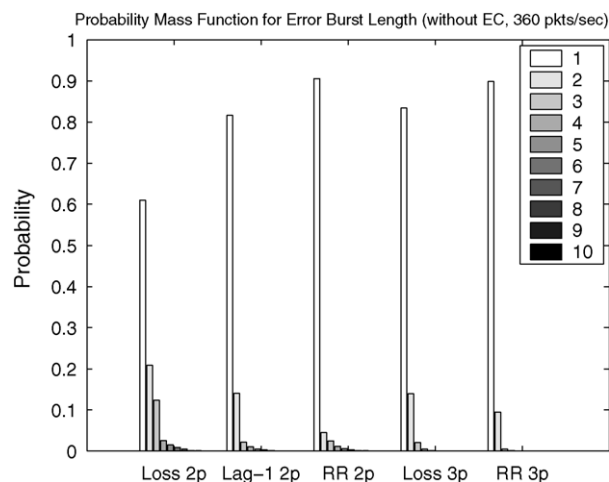


Fig. 10. Lost packet burst length probability mass functions (360 pkts/s).

Table 4
2-Paths optimization based on loss rate (with erasure code)

Rate (pkts/s)	α_1^*	α_2^*	Information loss rate at α^* (%)	Lag-1 rate at α^*	Best SP information loss rate (%)	RR information loss rate
60	0.260	0.740	0	N.A.	0	0
120	0.260	0.740	0.001	0.146542	0.177	0.057
360	0.260	0.740	15.055	0.378475	26.754	26.128

Table 5
2-Paths optimization based on lag-1 (with erasure code)

Rate (pkts/s)	α_1^*	α_2^*	Information loss rate at α^* (%)	Lag-1 rate at α^*	Best SP lag-1	RR lag-1
60	0.158	0.842	0	N.A.	N.A.	N.A.
120	0.362	0.638	0.002	0.144608	0.400462	0.074382
360	0.416	0.584	20.567	0.051666	0.778310	-0.071064

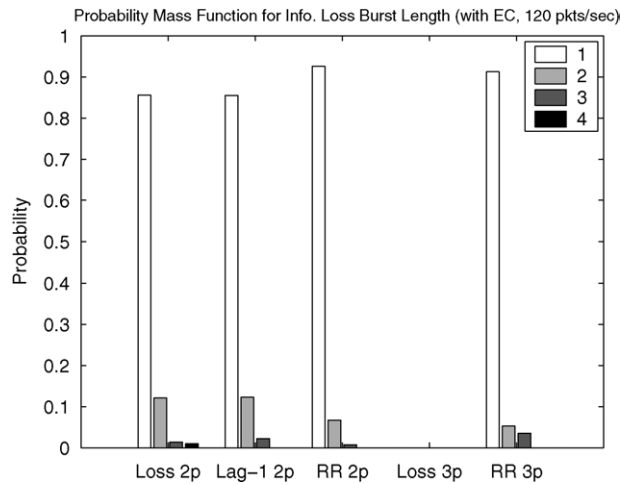


Fig. 11. Information packet loss burst length probability mass functions (120 pkts/s).

mance metrics, such as the optimal splitting vector α^* , the achieved information loss rate, the achieved lag-1 autocorrelation, as well as the system's performance when we stream the data under the *best* single path or using the round-robin approach. Information loss refers to packet loss after the reconstruction process.

Again, we consider the improvements in various performance measures when we employ an erasure code and an additional path wherein $\mathcal{F}_3(b) = 0.2333 \times b$ and $\mathcal{B}_3(b) = 24750/b$. Table 6 illustrates the relevant performance metrics,¹³ and Figs. 11 and 12 illustrate the corresponding conditional lost packet burst length probability mass function, conditioned on there being a loss, for application sending rates of 120 and 360 pkts/s, respectively.¹⁴ This is illustrated for 2- and 3-paths streaming and under different

¹³ Corresponding best path results can be found in Tables 4 and 5.

¹⁴ The 60 pkts/s case is omitted as all the missing data is reconstructed.




Fig. 12. Information packet loss burst length probability mass functions (360 pkts/s).

traffic splitting methods, e.g., “Loss 2p” refers to using 2-paths with load distribution computed using the loss rate based optimization method, “RR 3p” refers to using 3-paths with the round-robin approach, etc.

From above experiments, we observe that when one adopts an erasure code, in most cases, it can reduce the information loss rate but not the lag-1 autocorrelation. However, when the packet sending rate is high, employing an erasure code may have an adverse effect of increasing the loss rate (i.e., degrading the loss characteristics of paths, e.g., 2-paths streaming with 360 pkts/s and traffic splitting optimized based on loss rate). When an additional path is available, the workload (including the overhead) can be spread among more paths; this results in better information loss rate. We can also observe that our optimization methods result in significantly better system performance (under several metrics) than best path streaming or the round-robin approach. Moreover, we observe that optimization based on the loss rate may produce a lower information loss rate than the other optimization methods.

As mentioned in Section 3, one might want to minimize the mean information lost rate (MILR), which is the loss rate of the media data after the lost packet reconstruction process. Here we obtain an optimal load distribution based on the MILR metric by an essentially brute force search on the simulation results as an analytical derivation of this metric is quite complex. We do this only to illustrate the utility of the loss rate based optimization; we, of course, do not suggest the use of this brute force technique in real systems. For simplicity of presentation, we consider a 2-paths system. Path 1 corresponds to $\mathcal{F}_1(b) = 0.2333 \times b$ and

Table 6
3-Paths optimization based on loss rate (with erasure code)

Rate (pkts/s)	α^*	Information loss rate at α^* (%)	Lag-1 at α^*	RR information loss rate (%)	RR lag-1
60	[0.191, 0.540, 0.269]	0	N.A.	0	N.A.
120	[0.190, 0.541, 0.269]	0	N.A.	0	N.A.
360	[0.190, 0.541, 0.269]	5.301	0.217488	12.256	0.012303

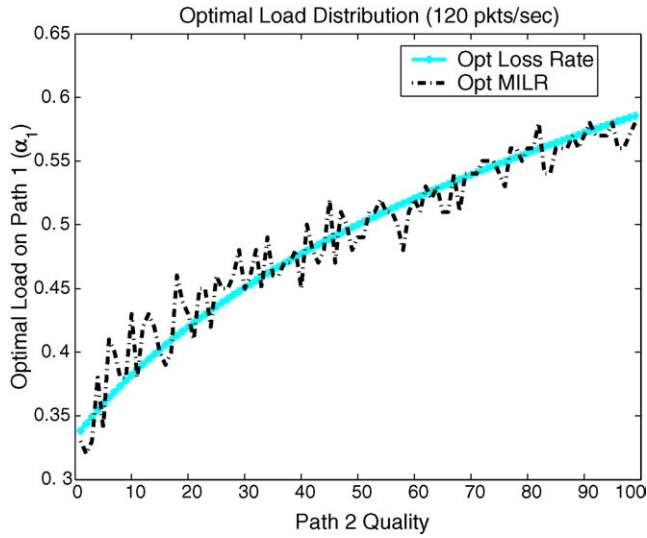


Fig. 13. Optimal load when minimizing MILR (120 pkts/s).

$\mathcal{B}_1(b) = 24750/b$. We vary the quality of path 2 and search for the optimal load distribution in simulation by testing all possible α_1 and picking the one with minimum resulting information loss rate. Let q denote the quality of path 2 ranging from 0 to 100 (0 meaning the best quality). The functions for the path 2 FGM are defined as: $\mathcal{F}_2(b) = (0.0667 + q/300) \times b$, $\mathcal{B}_2(b) = (28500 - q \times 75)/b$. When $q = 0$, these functions become: $\mathcal{F}_2(b) = 0.0667 \times b$, $\mathcal{B}_2(b) = 28500/b$ and are equal to the best path we used in the above experiments. When $q = 100$, these functions become: $\mathcal{F}_2(b) = 0.4 \times b$, $\mathcal{B}_2(b) = 21000/b$ and are equal to the worst path we used in the above experiments. Note that path 1 here has the same FGM




Fig. 14. Optimal load when minimizing MILR (360 pkts/s).

functions as the third path we used in the above 3-paths system. It has loss characteristics in-between the best path and the worst path. Figs. 13 and 14 show an optimal MILR-based load on path 1 when we vary the quality of path 2 for application sending rates of 120 and 360 pkts/s, respectively. The results of optimal load on path 1 when one optimizes the loss rate are also shown; they indicate that optimization based on loss rate produces a load distribution close to an MILR-based optimum. This suggests that (under erasure code use), packet loss rate based optimization (with an appropriately adjusted sending rate) is a reasonable approximation to minimizing MILR.

5. Prototype experiments

We implemented a multi-path streaming prototype [2], used here to study the performance metrics given in Section 2 as well as the resulting visual quality. Although the packet losses are still emulated using the FGM, the MPEG streams and their processing are real; hence we are able to illustrate the resulting visual quality in addition to the performance metrics. We stream a 150 s MPEG1 file requiring a 174.5 Kbps playback rate (≈ 170 pkts/s with a packet size of 1024 bytes); with $k = 64$ and $n = 72$, the packet sending rate is increased to ≈ 192 pkts/s. Two heterogeneous paths are used, with packet losses occurring according to their respective FGMs. FGM parameters for path 1 (with better loss characteristics), are: $\mathcal{F}_1(b) = 0.0667 \times b$, $\mathcal{B}_1(b) = 28500/b$, and FGM parameters for path 2 (with worse loss characteristics), are: $\mathcal{F}_2(b) = 0.4 \times b$, $\mathcal{B}_2(b) = 21000/b$, where $b = 192$ pkts/s. Using the loss rate based optimization approach in Section 3 we obtain $\alpha^* = [0.741, 0.259]$. Four cases with different traffic splitting vectors are studied: (Case 1) single path with *better* loss characteristics (using path 1), (Case 2) single path with *worse* loss characteristics (using path 2), (Case 3) dual path using *even* traffic splitting, and (Case 4) dual path using *optimal* traffic splitting vector α^* . Packet loss statistics are measured at the receiver throughout the streaming process. Video frames are transcoded to JPEG files to allow visual quality inspection. Table 7 gives the corresponding average statistics measured. Each column in Fig. 15 depicts a sequence of frames extracted from a particular test case; frames in the same row originate from the same original video frame prior to transmission. From these results, we make the following observations:

- (1) *Information loss rate significantly affects visual quality*: When we relate the information loss rate in Table 7 with the frame sequences' quality indication in Fig. 15, we find that higher information loss rate corresponds to poorer video quality. When the information loss rate is extremely high, for example 35.642% in Case 2, all the frames in the video sequence are damaged. When the loss rate is improved to 1.688% as in Case 1, some damaged frames can be found in some video segments. If

Table 7
Prototype experiments: average loss statistics

Test case	Loss rate before EC (%)	Lag-1 before EC	Average burst length before EC	Information loss rate (%)
1	6.569	0.393	1.762	1.688
2	35.630	0.363	2.439	35.642
3	6.955	-0.040	1.034	1.000
4	3.403	0.174	1.253	0.090

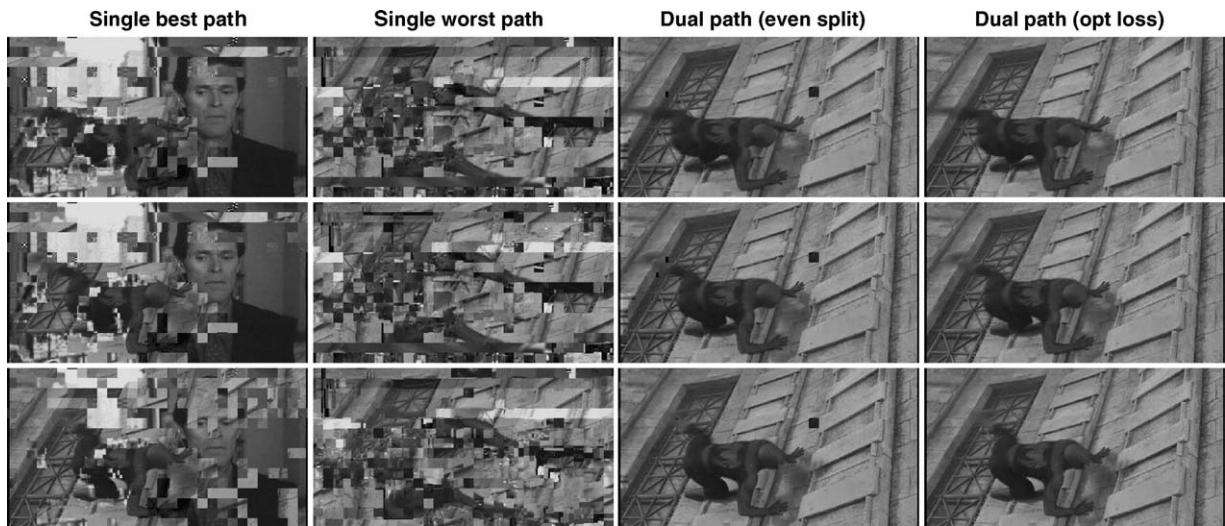


Fig. 15. Prototype experiments: visual quality of video frames.

some important data is lost (e.g., an I frame in the MPEG1 standard), scene mixing may also occur, as shown in the captured video output of Case 1. When the information loss rate is further improved to 1% in Case 3, video distortion happens less frequently and most of the distortion is in the form of “blocking effects”. When the loss rate is very low, e.g. 0.09% in Case 4, distortion is very rare and mostly unnoticeable to the human eye.

- (2) *Introducing EC may cause adverse effects*: Improper adding of EC may not improve the resulting information loss rate, e.g., in Case 2, the loss rate before the EC operation is less than the information loss rate. Adding an EC increases the loading along a congested path which may degrade the resulting loss characteristics.
- (3) *Use of EC is more effective under multi-path streaming*: By reducing the lag-1 auto-correlation and shortening the error burst length, the multi-path technique enhances the error correction capability of an EC. For example, the loss rate before an EC for Case 1 (single best path) is 6.569% and the loss rate before an EC for Case 3 (dual path with even split) is 6.955%. Using the best single path in Case 1 can give a lower loss rate *before* the EC process. However, *after* the error correction process, information loss rate is 1.688% for Case 1 and 1% for Case 3, i.e., the reconstruction process recovers more packet losses in the dual path case as compared to the single path case. The visual quality, which is strongly related to the information loss rate, is also better in Case 3 than in Case 1. It shows that simply using the best single path for video streaming may *not* be a good approach.
- (4) *A path with worse loss characteristics can be used to improve the overall performance*: Although path 2 is a worse path than path 1, it can still be used to share a fraction of the workload to improve the overall quality of the received video, i.e., the two multi-path test cases give better performance than the two single path test cases.
- (5) *Using an optimal traffic splitting vector can result in better performance*: The lowest loss rate achieved in Case 4 indicates that simply splitting the traffic evenly between paths may not result in the best use of multiple paths.

6. Conclusions

We considered the problem of providing QoS in streaming pre-stored CM using an application-layer multi-path streaming approach. An advantage of this approach, as compared to approaches that require support of lower layers and resource reservation schemes, is that the complexity of QoS provision can be pushed to the network edge. Hence, we can improve the scalability and deployability of a streaming application and at the same time provide a certain QoS level. Our past work evaluated the potential benefits of multi-path streaming using a conventional Gilbert model. Evidence presented here indicates that this model may not be sufficient, and hence we proposed a functional Gilbert model (FGM) which is more expressive in capturing the dependency between an application's sending rate and the loss characteristics of a path. We showed that under the FGM and with homogeneous paths, any valid traffic splitting has a packet loss rate no worse than a single path streaming approach. We then focused on optimal traffic splitting approaches. We presented results based on two optimization objectives, with and without the use of erasure codes. Moreover, no matter whether an erasure code is added or not, we observed that both optimization methods result in significantly better system performance (under several metrics) than single best path streaming or the simple round-robin approach. Finally, we have also implemented a prototype multi-path streaming system which we used to illustrate the significant merits of the the multi-path streaming approach in improving the visual quality of CM delivery.

Acknowledgments

This research was funded in part by the NSF ANI-0070016, NSF EIA-0091474, NSF CCR0113192, and RGC grants, as well as the Croucher Foundation Scholarship and the Okawa Research Award. It has also been funded in part by the Integrated Media Systems Center, a National Science Foundation Engineering Research Center, Cooperative Agreement No. EEC-9529152. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect those of the National Science Foundation.

References

- [1] B. Abdouni, W.C. Cheng, A.L.H. Chow, L. Golubchik, W.-J. Lee, J.C.S. Lui, Multi-path streaming: optimization and performance evaluation, in: Proceedings of the SPIE Conference on Multimedia Computing and Networking, vol. 5680, San Jose, California, January, 2005, pp. 216–227.
- [2] B. Abdouni, W.C. Cheng, A.L.H. Chow, L. Golubchik, J.C.S. Lui, Picture-perfect streaming over the internet: is there hope?, *IEEE Commun. Mag. Special Issue Proxy Support Stream. Internet* 42 (8) (2004) 72–79.
- [3] John Apostolopoulos, Mitchell Trott, Path diversity for enhanced media streaming, *IEEE Commun. Mag. Special Issue Proxy Support Streaming Internet* 42 (8) (2004) 80–87.
- [4] A. Begen, Y. Altunbasak, O. Ergun, Fast heuristics for multi-path selection for multiple description encoded video streaming, in: Proceedings of the IEEE Conference on Multimedia and Expo, Baltimore, MD, July, 2003, pp. 517–520.
- [5] J.-C. Bolot, S. Fosse-Parisis, D. Towsley, Adaptive FEC-based error control for Internet Telephony, in: Proceedings of INFOCOM, 1999.
- [6] J.-C. Chen, S.-H. Chan, V. Li, Multi-path routing for video delivery over bandwidth-limited networks, *IEEE J. Selected Areas Commun. Special Issue Design, Implement. Anal. Commun. Protocols* 22 (10) (2004) 1920–1932.

- [7] Alix L.H. Chow, Leana Golubchik, John C.S. Lui, Adam W.-J. Lee, Multi-path streaming: optimization of load distribution. Technical report, Computer Science Department, University of Southern California, 2005.
- [8] H. Chu, K. Nahrstedt, Dynamic multi-path communication for video traffic, in: Proceedings of the Hawaiian International Conference on System Science, Hawaii, January, 1997.
- [9] L. Golubchik, J.C.S. Lui, T.F. Tung, A.L.H. Chow, W.-J. Lee, G. Franceschinis, C. Anglano, Multi-path continuous media streaming: what are the benefits? *Perform. Eval.* 49 (2002) 429–449.
- [10] http://tequila.usc.edu/iml/mppath/papers/mppath_burstloss_exp.pdf, importance of considering loss correlations in multi-path streaming.
- [11] <http://www.isi.edu/nsnam/ns/>, The Network Simulator-ns-2.
- [12] <http://www.mesquite.com/>, CSIM18.
- [13] T. Nguyen, P. Mehra, A. Zakhor, Path diversity and bandwidth allocation for multimedia streaming, in: Proceedings of ICME, July, 2003.
- [14] T. Nguyen, A. Zakhor, Multiple sender distributed video streaming, *IEEE Trans. Multimedia* (2004) 315–326.
- [15] T. Nguyen, A. Zakhor, Distributed video streaming over internet, in: Proceedings of the SPIE Conference on Multimedia Computing and Networking, San Jose, California, January, 2002.
- [16] Thanh Nguyen, Avidah Zakhor, Path diversity with forward error correction (pdf) system for delay sensitive applications over the internet, in: Proceedings of IEEE InfoCom, San Francisco, California, March, 2003.
- [17] D. Rubenstein, J. Kurose, D. Towsley, Detecting shared congestion of flows via end-to-end measurement, *IEEE/ACM Trans. Network.* 10 (3) (2002) 381–395.



A.L.H. Chow was born in Hong Kong. He received his BEng in Computer Engineering and a MPhil in Computer Science and Engineering from the Chinese University of Hong Kong in 2001 and 2003, respectively. He is currently pursuing a PhD degree, supported by the Croucher Foundation Scholarship, in the Department of Computer Science at the University of Southern California. His research interests are in data networking and distributed multimedia systems.



L. Golubchik is an Associate Professor in the Computer Science and EE-Systems Departments at USC. Prior to that she was an Assistant Professor and then an Associate Professor at the University of Maryland at College Park and an Assistant Professor at Columbia University. Leana received her PhD from UCLA in 1995. She is a past Chair (2003–2005) and Vice Chair (2001–2003) of ACM SIGMETRICS (2003–2005) and a member of its Board of Directors (1999–2001, 2005–2007). She was a guest co-editor for special issues of IEEE TKDE, the Parallel Computing Journal, and the International Journal of Intelligent Systems, a program co-chair of the 2001 Joint ACM SIGMETRICS/Performance Conference and MIS'99, and a program committee member of numerous conferences. Leana received several awards, including the NSF CAREER award, the Okawa Foundation award, and the IBM and NSF Doctoral Fellowships. She is a member of the IFIP WG 7.3 and Tau Beta Pi.



J.C.S. Lui was born in Hong Kong. He received his PhD in Computer Science from UCLA. After his graduation, he joined the IBM Almaden Research Laboratory/San Jose Laboratory and participated in various research and development projects on file systems and parallel I/O architectures. He later joined the Department of Computer Science and Engineering at the Chinese University of Hong Kong. For the past several summers, he has been a visiting professor in Computer Science Departments at UCLA, Columbia University, University of Maryland at College Park, Purdue University, University of Massachusetts at Amherst and Universit degli Studi di Torino in Italy. His research interests span both in system and in theory/mathematics. His current research interests are in theoretic/applied topics in data networks, distributed multimedia systems, network security, OS design issues and mathematical optimization and performance evaluation theory. John received various departmental teaching awards and the CUHK Vice-

Chancellor's Exemplary Teaching Award. He is an associate editor in the Performance Evaluation Journal, member of ACM, a senior member of IEEE and an elected member in the IFIP WG 7.3. John was the TPC Co-chair of ACM Sigmetrics 2005. His personal interests include films and general reading.



W.-J. Lee received his BS degree from the Department of Computer Science and Information Engineering at the National Taiwan University in 1993, and his MS degree from the Department of Computer Science at the New York University in 1998. He worked at Bell Laboratories Research, Lucent Technologies, from 1998 till 2000. He is currently pursuing a PhD degree in the Department of Computer Science at the University of Maryland at College Park. His research interests include systems simulation and performance evaluation, computational and systems biology. He is a member of the ACM, the IEEE and the ISCB.