Multi-path Streaming: Is It Worth the Trouble?

Leana Golubchik*
Computer Science Department,
ISI and IMSC
University of Southern California
Los Angeles, CA 90089
leana@cs.usc.edu.

John C.S. Lui
Computer Science & Engineering Dept
The Chinese University of Hong Kong
Shatin, Hong Kong
cslui@cs.cuhk.edu.hk

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, involving delivery of video and audio, is partly due to variations in delays as well as losses experienced 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, in our work 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 (such as video) destined for a particular receiver, i.e., this data is fragmented into packets and the different packets take alternate routes to the receiver. We 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: (a) reduction in correlation between consecutive packet losses which we believe will lead to improvements in the quality of delivered data, (b) increased throughput, and (c) ability to adjust to variations in congestion patterns on different parts of the network. In this work, our focus is on providing the 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 below). More 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 this work is a step in the right direction.

In general, the use of multiple paths in designing of distributed (over best-effort wide-area networks) continuous media applications requires consideration of a number of issues (please refer to [5] for details). However, in this work we limit the scope by focusing on (a) delivery of *pre-stored* video, (in contrast to delivery of "live" data), (b) *application-level* schemes (which are deployable today over the current Internet), where the path used between any pair of hosts is determined by a network-level routing algorithm (we note that 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 [7]); and (c) accomplishment of multiple paths to the same receiver by *distributing servers* across wide-area networks and streaming data from multiple senders simultaneously.

Our system is depicted in Figure 1, where any server can send any fraction of the continuous media data. More specifically, server

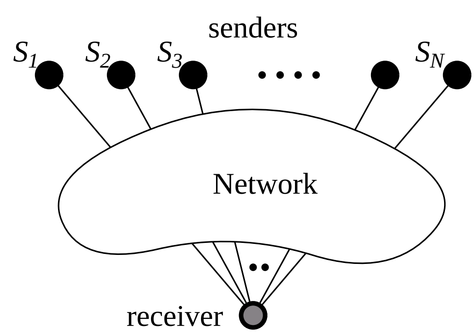

Figure 1: Continuous media system using multi-path streaming.

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.

Within the context of such a system, we focus on an analytical characterization of when a multi-path approach is beneficial, as compared to a single path approach. As in [2], we use a two-state Markov chain, known as the Gilbert model (refer to Figure 2), to model a path between a sender and the receiver (which is characterized by its bottleneck link). This model, allows for dependence in consecutive packet losses and should be a more accurate representation of the network than an independent loss model. Moreover, we use the following metrics in the evaluation of multi-path streaming benefits: (a) packet loss rate, (b) lag-1 autocorrelation of packet losses, and (c) burst length distribution. We focus on these metrics as we believe that the burstiness characteristics of the losses and the correlation between them (in addition to the loss rate) have a

*This work was partly done while the author was with the University of Maryland.

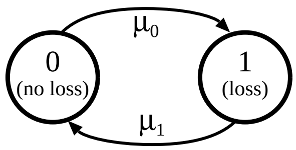

Figure 2: Gilbert model.

significant effect on the quality of the displayed continuous media (please refer to the Appendix). We also extend this analysis to information loss rate, i.e., we consider the resulting losses after an application of an erasure code, e.g., as in FEC [1] techniques.

Our results to date 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; we believe that all this should result in a higher viewing quality of the received continuous media. (Please refer to [4, 5] for details on current results.) These results can also 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.

Workshop discussion and other considerations

We note that one should also consider the potential costs or detrimental effects of multi-path streaming. For instance, multi-path 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. 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 multi-path streaming should also be considered. For instance, in our case, we employ detection of shared points of congestion [7] to improve the performance of our multi-path streaming system. Other approaches to multi-path streaming might require even more detailed information about the network (refer to [5]) which is likely to result in a need for more "intrusive" and complex measurements. Note that, scalability of such measurement schemes is an issue as well. Lastly, the applicability of the Gilbert model to real networks is an important consideration as well. Our future work includes Internet measurement studies which are intended to address this issue.

Acknowledgements

The authors are grateful to Bill Cheng, Don Towsley, and Edmundo de Souza e Silva for fruitful discussions. This research was supported in part by the NSF ANI-0070016 grant, the CCR-0113192 grant, the Integrated Media Systems Center, a National Science Foundation Engineering Research Center, Cooperative Agreement No. EEC-9529152, as well as by the Earmarked Grant and the Mainline Research Grant.

1. REFERENCES

[1] R. E. Blahut. *Theory and Practice of Error Control Codes*. Addison Wesley, January 1983.

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

[3] D. Le Gall. MPEG: a Video Compression Standard for Multimedia Applications. *Communications of the ACM*, April 1991.

[4] Leana Golubchik, John C.S. Lui, Tak Fu Tong, Lik Hang Chow, W.J. Lee, G. Franceschinis, and C. Anglano. Multi-path continuous media streaming: What are the benefits? *Performance Evaluation (presented at Performance 2002)*, 2002.

[5] Leana Golubchik, John C.S. Lui, Tak Fu Tong, Lik Hang Chow, W.J. Lee, G. Franceschinis, and C. Anglano. Multi-path continuous media streaming: What are the benefits? Technical Report CS-TR-2002-01, Department of Computer Science and Engineering, The Chinese University of Hong Kong, April 2002.

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

[7] 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, 2002.

Appendix – Visual Quality of Data

Here we 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 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 [6]. Hence, we illustrate the effects of bursty losses on the quality of the resulting video (and specifically on the signal to noise ratio) using the following experiment.

Experiment (Effect of Correlated Bursty Losses on Video Quality) : In this experiment, we drop 2% of the frames from video $\mathcal{V}$. These 2% losses are introduced in a variety of "patterns", e.g., the dropped frames can be evenly spaced throughout video $\mathcal{V}$, or they can be more bursty. The details of which frames are dropped, given a particular drop pattern as identified by the burst length, are given in the first two columns of Table 1. Moreover, in evaluating the quality of the resulting video $\mathcal{V}$, we use a common error concealment scheme to make up for a dropped frame. Specifically, a dropped frame is replaced by the previous frame which is successfully received. For example, frame i replaces frames $i+1, i+2, \cdots, i+k$ if frame i is received successfully and frames $i+1, \cdots, i+k$ are lost.

For each possible frame loss pattern, we measure the quality of the received video by computing the peak signal-to-noise ratio (PSNR) as follows. (Note that, a larger value of PSNR implies a higher quality of the video.) In general, for a video of l frames where each frame consists of $m \times n$ pixels, (each containing an RGB value[1] with each of the three colors represented by 8-bits), the PSNR is calculated using the following expression (in dB):

$$SNR_{peak} =$$
$$10 \times \log_{10} \frac{255^2}{\left(\frac{\sum_{i=1}^{m} \sum_{j=1}^{n} \sum_{k=1}^{l} \sum_{c=1}^{3} (P_1(i,j,k,c) - P_2(i,j,k,c))^2}{3 \times m \times n \times l}\right)}.$$

[1] Information about the three colors, red, green, and blue.

Error Burst Length	Lost Frames Numbers	PSNR (dB)
1	25+k*50 where k$\in$ $\{0, 1, \cdots, 29\}$	39.107 dB
2	$\{$50,51$\}$ + k*100 where k$\in$ $\{0, 1, \cdots, 14\}$	38.015 dB
3	$\{$74,75,76$\}$ +k*150 where k$\in$ $\{0, 1, \cdots, 9\}$	31.325 dB
5	$\{$123,124,125,126,127$\}$ +k*200 where k$\in$ $\{0, 1, \cdots, 5\}$	30.433 dB
15	$\{$368,369,...,381,382$\}$ +k*750 where k$\in$ $\{0, 1\}$	28.407 dB
30	$\{$736,737,...,764,765$\}$	29.942 dB

Table 1: Peak signal-to-noise ratio (PSNR) for various bursty loss patterns.

where $P_s(i, j, k, c)$ is the pixel value at coordinate (i, j) of k-th video frame (of stream s, $s = 1, 2$) and color channel c where $c = 1, 2, 3$, for red, green, and blue, respectively. In our experiment, the values of m,n, and l are 352, 240 and 1500, respectively. The source video in this experiment is using MPEG-1 NTSC settings [3] where each frame is 352×240 (with 29.97 frames per second), hence the values of m and n above. Also, we use approximately the first 50 seconds of this video for this experiment, hence the value of l above. Values for P_1 are obtained from the frame sequence resulting after the drop-and-conceal process while values for P_2 are obtained from the original video frames of $\mathcal{V}$.

Table 1 gives the PSNR values for the different burst patterns. We can observe that given the same amount of information loss (e.g., 2% in our experiment), the PSNR metric can be significantly lower for the more bursty loss patterns, and hence is potentially the perceptual quality of the video. Thus, we believe that burst length distribution and correlations between losses are the right metrics for evaluating the goodness of a streaming approach as they directly reflect on the quality of the received video.