

# Application, Network and Link Layer Measurements of Streaming Video over a Wireless Campus Network

Feng Li, Jae Chung, Mingzhe Li, Huahui Wu, Mark Claypool, Robert Kinicki  
{lif,goos,lmz,flashine,claypool,rek}@cs.wpi.edu  
CS Department at Worcester Polytechnic Institute  
Worcester, MA, 01609, USA

**Abstract.** The growth of wireless LANs has brought the expectation for high-bitrate streaming video to wireless PCs. However, it remains unknown how to best adapt video to wireless channel characteristics as they degrade. This paper presents results from experiments that stream commercial video over a wireless campus network and analyze performance across application, network and wireless link layers. Some of the key findings include: 1) Wireless LANs make it difficult for streaming video to gracefully degrade as network performance decreases; 2) Video streams with multiple encoding levels can more readily adapt to degraded wireless network conditions than can clips with a single encoding level; 3) Under degraded wireless network conditions, TCP streaming can provide higher video frame rates than can UDP streaming, but TCP streaming will often result in significantly longer playout durations than will UDP streaming; 4) Current techniques used by streaming media systems to determine effective capacity over wireless LAN are inadequate, resulting in streaming target bitrates significantly higher than can be effectively supported by the wireless network.

## 1 Introduction

The combination of the decrease in price of wireless LAN access points (APs) and the increase in wireless link capacities has prompted a significant increase in the number of wireless networks in homes, corporate enterprise networks, and academic campus networks. The promise of up to 54 Mbps capacity<sup>1</sup> from a wireless AP means that users now expect to see applications such as streaming video that require high bitrates running seamlessly from wired media servers to wireless media clients.

Although much is already known about wireless LANs and the individual components of the wireless LAN environment that make the delivery of high-demand applications over wireless a challenge, there has been little effort to integrate measures of wireless link layer performance with streaming media application layer choices. Such knowledge can facilitate the redesign of streaming

---

<sup>1</sup> IEEE 802.11g

media systems to account for the trend towards a wireless last hop to clients. Moreover, a better understanding of the impact of wireless LAN transmission characteristics on streaming media is valuable to network practitioners concerned with providing adequate wireless LAN coverage and discovering trouble spots in network performance.

Previous work [8] has shown that streaming products such as RealNetworks and Windows Streaming Media make important decisions concerning the characteristics of the video stream prior to streaming the video to a client. These decisions are based on estimates of the underlying network characteristics obtained from network probes. However, it remains unclear which wireless channel characteristics, such as frame loss rate, signal strength, or link layer bitrate, are the most useful for streaming media strategies that improve the performance of a streaming video by adapting video transmission choices to current wireless network conditions.

A primary goal of this investigation is to correlate wireless link layer behavior and network layer performance with streaming media application layer performance. Application layer measurement tools [6] were combined with commercially available network layer measurement tools and publicly available IEEE 802.11 measurement tools to conduct wireless experiments and integrate the measurement results. Seeking to characterize the impact of wireless network conditions on streamed video performance, this active measurement study considers four hypothesis:

1. *Wireless LANs make it difficult for streaming video to gracefully adapt when network conditions degrade.* This investigation attempts to uncover specific characteristics of streams to poor locations that could trigger streaming server adjustments to improve video transmission quality. Increasing performance in poor locations is critical since a streaming wireless client with bad performance can negatively impact other wireless clients connected to the same AP [1].
2. *Videos encoded with multiple levels can stream better than videos encoded with only a single level when wireless LAN conditions are poor.* Commercial media encoders allow videos to be encoded with one or more target bitrate levels. When streaming, the server determines which encoding level to use based on feedback from the client regarding the client end-host network conditions. A video with multiple levels of encoding should make better use of a wireless LAN with limited capacity than a video with a single level of encoding.
3. *TCP is more effective than UDP for streaming video over wireless LANs.* Commercial media players typically let the client select the streaming transport protocol. UDP is often selected due to lower overhead and jitter. However, recent work [4, 5, 11] suggests TCP and TCP-like protocols can be at least as effective and potentially more effective at providing higher quality video to clients under poor network conditions.
4. *Current techniques used by streaming media systems to estimate available capacity to a wireless LAN client are inadequate for providing the best video*

*performance*. Some commercial media players use packet-pair techniques [10] to estimate the capacity along the flow path prior to starting the streaming of the video to the client [8]. However, packet-pair was not designed for wireless networks where changes in transmission conditions cause mid-stream wireless capacity changes. By measuring frame errors and signal strength at the data link layer during wireless streaming experiments, changes in the wireless environment can be correlated with changes in video performance, and facilitate the development of better wireless capacity estimators.

## 2 Methodology

### 2.1 Tools

The unique aspect of this investigation is the concurrent use of measurement tools at multiple levels in the network protocol stack to evaluate streaming media performance over wireless LANs. This section discusses the tools employed in this study. For reference, the layer corresponding to each tool and examples of some of the performance measurements available from each tool are listed in Table 1.

**Table 1.** Measurement Tools

Layer	Tools	Performance Measures
Application	Media Tracker	Frame rate, Frames lost, Encoded bitrate
Network	UDP Ping, Wget	Round-trip time, Packet loss rate, Throughput
Wireless	Typeperf, WRAPI	Signal strength, Frame retries, Capacity

At the application layer, the WPI Wireless Multimedia Streaming Lab has experience measuring video client and server performance [4, 6, 8, 12]. An internally developed measurement tool, called *Media Tracker* [6], streams video from a Windows Media Server, collecting application layer data specific to streaming video including: encoding data rate, playout bitrate, time spent buffering, video frame rate, video frames lost, video frames skipped, packets lost and packets recovered.

For network layer performance measures such as round-trip time and packet loss rate along the stream flow path, *UDP ping*, an internally developed tool, was used. Preliminary experiments revealed that because the standard ICMP *ping* provided by Windows XP waits for the previous ping reply or a timeout before sending out the next ping packet, a constant ping rate could not be maintained in some poor wireless conditions where 10 second and longer round-trip times were recorded. Thus, a customized ping tool using application-layer UDP packets was built to provide constant ping rates, ping intervals configurable in milliseconds, and configurable ping packet sizes.


At the wireless data link layer, a publicly-available library, called *WRAPI* [2] was enhanced to collect information at the wireless streaming client that includes: signal strength, frame retransmission counts and failures, and information

about the specific wireless AP that handles the wireless last hop to the client. Additionally, *typeperf*, a performance monitoring tool built-in to Windows XP, collected processor utilization and network data including data received bitrate and the current wireless target capacity.

Although the above four tools were deployed concurrently on the wireless streaming client, baseline measurements indicated these tools consume only about 3% of the processor time. Given that streaming downloads consumed about 35% of processor time, the assumption is the measurement tools do not significantly effect the performance of the streaming downloads to the wireless clients.

## 2.2 Experiment Setup

This investigation conducts a series of experiments where video clips are streamed from a Windows Media Server over a wired campus network to a wireless streaming client at pre-determined locations in the WPI Computer Science Department building. As Figure 1 shows, the wireless portion of the WPI campus network is partitioned from the wired infrastructure. Thus, the assumption is that all video streams traverse the same network path except for the last two hops from a common exit off the wired campus LAN to a wireless AP and from the AP to the streaming client. The media server runs Windows Media Service v9.0 as part of the Windows Server 2003 Standard Edition, and the wireless client resides on a Dell laptop with a Centrino mobile CPU running Windows XP sp1 and an IEEE 802.11g wireless network adaptor based on the Broadcom<sup>2</sup> chipset. The WPI wireless LAN uses Airespace<sup>3</sup> APs and provides IEEE 802.11 a/b/g wireless service for all the experiments.



**Fig. 1.** WPI Campus Network

<sup>2</sup> <http://www.broadcom.com/>

<sup>3</sup> <http://www.airespace.com/>

Two distinct video clips, known as Coast Guard and Paris, were used in this study. Both clips were encoded to run at  $352 \times 288$  resolution and 30 frames per second. Both clips run for approximately two minutes.<sup>4</sup> Coast Guard is a high-motion video clip (5.4% skipped macro blocks) with a camera panning scene of a moving Coast Guard cutter. Paris is a low-motion video clip (41.2% skipped macro blocks) with two people sitting and talking with some rapid motion from two small objects in the scene.

Windows Media Server selects the streaming rate based upon the encoded bitrate of the layers in the video clip and an estimate of available capacity for the bottleneck link along the flow path. During this investigation, two distinct versions of each video were streamed to every client location: a single level version of the video encoded at 2.5 Mbps to stress the wireless link; and a multiple level version that includes eleven encoding layers such that the streaming server has the opportunity to do media scaling to dynamically choose the encoded clip to stream based on perceived network capacity. To compare the performance of standard streaming protocol choices, each of the four videos instances was streamed using TCP and repeated using UDP.

### 2.3 Experiment Design

At the beginning and the end of each experimental instance, the client downloaded a large file using *wget*, a publicly-available HTTP/FTP download application,<sup>5</sup> to estimate the effective throughput of a TCP bulk transfer. Thus, each experiment consisted of an initial bulk download, eight different video downloads (2 clips (Paris and Coast Guard)  $\times$  2 versions (Single Level and Multiple Level)  $\times$  2 transport protocols (UDP and TCP)) and a final bulk download. While each video was streamed, the client initiated UDP ping requests to determine round-trip time and packet losses. The UDP ping requests were 200 milliseconds apart, with 1350-byte packets for the single level video and 978-byte packets for the multiple level video. The choice of packet sizes came from the observation that 90% of the packets are 1350 bytes and 978 bytes for single level and multiple level video, respectively. While streaming, measurement data was also collected by WRAPI, typeperf and Media Tracker at the client side on a stationary laptop.

Clearly, wireless networking transmission performance is dependent on current network conditions. To reduce the variability in the network conditions, all the experiments were conducted during the Winter Break (December 23-25, 2004 and December 29-30, 2004) in the Computer Science Department on the WPI campus. During these testing periods, there was only occasional network activity and virtually no other wireless users in the Computer Science department. Each experiment was repeated five times at three distinct locations on three different floors in the Computer Science department. Thus the results come from a total of 45 experimental runs that include 360 video streams. On each floor, an AP was selected to interact with the client laptop. Then, preliminary experiments

<sup>4</sup> The median duration of video clips stored on the Internet [7] is 2 minutes.

<sup>5</sup> <http://www.gnu.org/software/wget/wget.html>

were conducted to find three laptop reception locations for each AP, representing good, fair, and bad reception locations. It turned out to be difficult to make a clear distinction between bad and fair locations due to high variability in the signal strength at fair and bad locations.

### 3 Results

#### 3.1 Data Collected

Ten data sets were removed from the 360 video streaming runs due to wireless connection failures that caused abnormal streaming terminations. Thus, 350 video instances (see Table 2) are included in the analysis of the results.

**Table 2.** Data Collected

	TCP Streaming	UDP Streaming	Total
Multiple Level Video	86	85	171
Single Level Video	89	90	179
Subtotal	175	175	350

Comparison of the two clips, Paris and Coast Guard, with analysis similar to the other experimental factors presented in Section 3.2-3.4, produced no statistically significant differences in performance. This suggests that the differences in motion between the low-motion Paris video and the high-motion Coast Guard video did not impact performance over a wireless network. Thus, all subsequent analysis combines the data obtained for both clips for each of the categories in Table 2.


#### 3.2 Categorization

Figure 2 depicts the throughput obtained versus signal strength for all the streaming and bulk download instances. The streaming data and the bulk download data are separately fit with logarithmic functions. The root mean square value of the deviation of the data from the fitted function<sup>6</sup> are 0.49 Mbps and 1.47 Mbps for streaming throughput and bulk downloading throughput, respectively. Note, there is a “cliff” where throughput degrades suddenly when the signal strength is between -70 dBm and -80 dBm.


To provide a clearer picture of streaming video behavior, the experiments were classified by the average signal strength recorded for a download from the server to the instrumented video client. For the remainder of the analysis, the experiments are categorized in one of three distinct regions: “Bad” locations (less than -75 dBm); “Edge” locations (between -75 and -70 dBm); and “Good” locations (greater -70 dBm).<sup>7</sup> This classification facilitates focusing on understanding

<sup>6</sup> The *stdfit* reported from the *gnuplot* fit function.

<sup>7</sup> The variance in signal strength is about the same for both Good and Bad locations.



**Fig. 2.** Average Wireless Signal Strength versus Average Throughput




**Fig. 3.** CDF of Wireless Signal Strengths

the performance differences between the Good and the Bad locations. Figure 3 shows a cumulative distribution function (CDF) of the average signal strengths gathered and depicts the Good, Edge and Bad regions, with approximately 1/3 of the data points in each region.


### 3.3 Single Level Encoding versus Multi-Level Encoding

As described in Section 2.1, both video clips were encoded twice, once at a single, high-bandwidth encoded level and again with multiple encoded levels. Figure 4 and Figure 5 provide CDFs to compare the impact of the server having multiple encoding levels versus only a single encoding level for wireless streaming. These figures indicate that when the client is at a Good location, the number of encoded levels has little effect on the average video frame rate and the coefficient of variation of the video frame rate. Since a Good wireless connection can generally support both the single level and the highest level in the multiple level clip, the stream does not need to be scaled to a lower bitrate.

However, at Bad locations, multiple level encoding provides better streaming performance than single level encoding. More than 2/3 of the time, the multiple level clip has a higher frame rate than the single level clip, and the multiple level clip has a median frame rate of 22 frames per second compared to a median of 11 frames per second for the single level clip.



**Fig. 4.** Average Application Frame Rate for Multiple and Single Level Encoding




**Fig. 5.** Coefficient of Variation of Application Frame Rate for Multiple and Single Level Encoding

### 3.4 TCP Streaming versus UDP Streaming


Figures 6-10 provide CDFs to compare the impact of choosing TCP versus UDP when streaming videos to clients at Good and Bad wireless locations. These figures show that at Good wireless locations, the choice of TCP or UDP has little effect on the average and coefficient of variation of frame rate. However, Figure 6 demonstrates that at Bad wireless locations, streams received by TCP streaming clients have a higher median frame rate (24 fps) than streams received by UDP streaming clients (15 fps). Moreover, the TCP streams have a higher frame rate about 2/3 of the time. Similarly, in Figure 7 the TCP streams have a lower median variation in frame rate than the UDP streams, and for 2/3 of the Bad locations TCP streams have a lower variation in frame rate than the UDP streams.

TCP video streams may be able to achieve better application frame rates under Bad conditions than UDP because when the wireless layer loses data, TCP retransmits the data and allows it to be played. However, without built-in retransmissions, UDP does not automatically recover lost data. The inter-frame dependencies in video can cause loss rates as low as 3% to result in up to 30% of application frames being unplayable [3]. Figure 8 graphs the CDF of wireless layer retry fraction for upstream (from the client to the server) data and Figure 9 shows a CDF of network ping loss rates. Under Bad conditions, approximately 1/3 of all wireless layer frames need to be retransmitted and when the same wireless frame is retransmitted too many times, the wireless layer drops the





**Fig. 6.** Average Application Frame Rate for TCP and UDP Streams



**Fig. 7.** Coefficient of Variation of Application Frame Rate for TCP and UDP Streams

frame and this yields network ping packet loss. Under Bad wireless conditions, nearly 1/3 of the time the network loss rate is about 15%.

The CDF for round-trip times in Figure 10 demonstrate that UDP packets suffer significant delays. Since the CDF of network ping packet loss rates measured for these UDP streams do not rise nearly as swiftly as the round-trip times in Figure 10, the conjecture is that the downstream wireless AP queues are large. Previous experience with Windows Streaming Media UDP streams [6, 8] suggests that excessively high average round-trip times occur when the initial UDP streaming stage uses a high data rate to fill the playout buffer. In Bad wireless situations, the downstream AP queue grows excessively long and the AP is never able to drain the queue since the actual wireless layer capacity is limited by degraded capacity and wireless layer retries.

In the presence of loss, the TCP stream may take longer to play out the same length video due to retransmissions. Severe loss causes TCP timeouts that delay video playout further. Figure 11 illustrates this behavior where total application playout duration (including buffering and playout) has been normalized by dividing it by the encoded (real-time) playout duration. In this figure, a normalized duration of one<sup>8</sup> indicates that the clip playout was the same length as the encoded duration, while a 2 implies the clip took twice as long to play as the encoded duration. At Bad locations, TCP streaming can take significantly longer to playout than UDP streaming. For pre-recorded clips, it is not unreasonable

<sup>8</sup> Note, the data points are all above one since the playout invariably includes at least one, initial buffering stage of about 10 seconds.




Fig. 8. Wireless Layer Retry Fraction for Upstream Traffic




Fig. 9. Network Loss Rates for TCP and UDP Streams




Fig. 10. Network Round-Trip Times for TCP and UDP Streams




Fig. 11. Normalized Duration of Application Playout for TCP and UDP Streams

to consider a stream duration extended by more than 10% to be unacceptable to users. Using this criteria, approximately 40% of the TCP Bad streams in Figure 11 are unacceptable.

## 4 The Challenges of Streaming over Wireless

Upon connection, video servers select an encoded send bitrate based on client feedback on network performance. Past work [8] indicates Windows Streaming Media uses a packet-pair technique [10] to estimate the bottleneck link capacity on the streamed path. Near the “cliff” in wireless performance, it is likely that a client will indicate an optimistically high average capacity that causes the video server to select a high encoding level. Figure 12 captures this phenomenon via a scatter-plot of the average encoding rate versus average wireless capacity both averaged over the duration of the video run. Points below the diagonal represent runs where the average encoding rate chosen for streaming is below the average capacity reported by the wireless network.

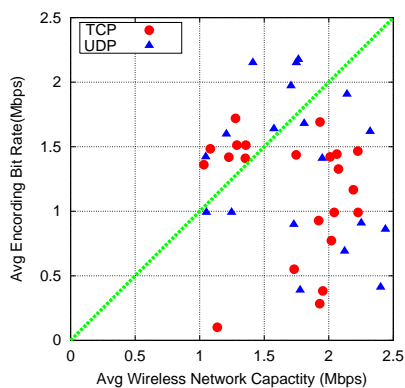
A conservative measure of effective capacity is the *TCP-Friendly* rate, namely, the data rate does not exceed the maximum rate of a conformant TCP connection under the same network conditions. The TCP-Friendly rate,  $T$  Bps, for a connection can be computed by [9]:

$$T = \frac{s}{R\sqrt{\frac{2p}{3}} + t_{rto}(3\sqrt{\frac{3p}{8}})p(1 + 32p^2)} \quad (1)$$


with packet size  $s$ , round-trip time  $R$  and packet drop rate  $p$ . TCP retransmission timeout  $t_{rto}$  is set to four times round-trip time by default. For each video clip for each run, Equation (1) is used to compute the TCP-Friendly rate ( $T$ ), using a packet size ( $s$ ) of 1350 bytes for the single level video and 978 bytes for the multiple level video, and the loss rate ( $p$ ) and round-trip time ( $R$ ) obtained from the corresponding ping samples.

Figure 13 shows a scatter-plot of the average encoding rate and average wireless network capacity both averaged over the video duration. Points above the diagonal line represent video runs in which the average encoding rate chosen for streaming are above the average effective capacity that can be supported by the wireless network. The preponderance of points above the diagonal line suggest the video streaming rate chosen is quite often higher than the capacity that the wireless network can effectively support. This results in the application streaming rate being too high to be supported by the network. Under such cases, when the video is streamed over UDP, the result is a reduced frame rate and when the video is streamed over TCP, the result is a longer playout duration.

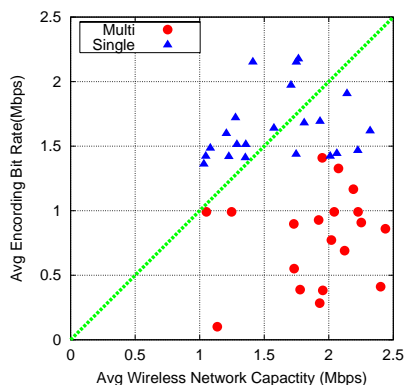
Videos encoded with multiple levels provide modest performance improvement by enabling the video streaming rate to more easily adapt to the effective network capacity after streaming has commenced. This is depicted in Figure 14 and Figure 15 which show scatter-plots similar to Figure 12 and Figure 13, respectively, but broken down by multiple and single encoding levels. In Figure 15, the cluster of points in the bottom left corner of the graph are cases where the



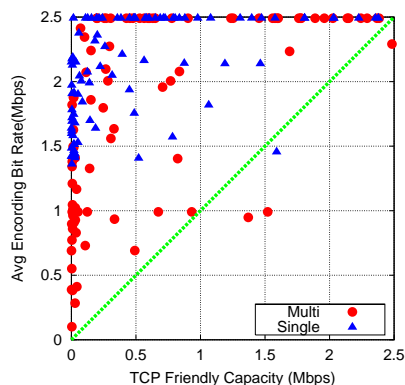
**Fig. 12.** Average Application Encoding Rate versus Wireless Capacity for TCP and UDP Streams



**Fig. 13.** Average Application Encoding Rate versus TCP-Friendly Capacity for TCP and UDP Streams



**Fig. 14.** Average Application Encoding Rate versus Wireless Capacity for Multiple and Single Level Stream



**Fig. 15.** Average Application Encoding Rate versus TCP-Friendly Capacity for Multiple and Single Level Stream

multiple level clips are able to stream at an average encoded rate closer to the capacity that the wireless network can effectively support.

This data suggests a need for more effective techniques to estimate the effective capacity for wireless networks to facilitate better choices for video encoding and streaming rates.

## 5 Conclusions and Future Work

This investigation reinforces the notion that IEEE 802.11 wireless networks can support streaming of high-quality video to wireless clients at high signal strength reception locations. Under such good conditions, nearly all video clips in this

study played out at high frame rates. Moreover, server choices of multiple versus single encoding levels and TCP versus UDP streaming did not significantly impact performance at good locations. However, these experiments produce a noticeable cliff such that throughput drops off suddenly when signal strength degrades below -75dBm. The bad wireless environment for those experiments at the bottom of this cliff can be characterized by nearly 50% more retries for wireless MAC layer frames and median packet loss rates over 5%. Under such bad wireless environments, multiple level videos adapt better to volatile wireless conditions than videos encoded only at a single level. Under bad conditions, multi-level videos consistently had higher frame rates with a median of 24 application frame per second, approximately double the median application frame rate of single level videos.

At bad client locations, TCP streamed videos usually recorded higher frame rates than UDP streamed videos. For the TCP streams, the median of 24 application frames per second, was approximately 50% higher than the UDP median. The conjecture is that TCP retransmissions reduce network packet loss rates that yield more playable frames than UDP when wireless conditions are bad. Unfortunately, this higher TCP frame rate comes at a price, significantly longer video playout durations. Nearly 20% of the TCP streamed videos to bad client locations had the two-minute video clip produce four minute playout durations. Approximately 40% of these TCP videos had playout durations considered to be intolerable. While UDP streams also experienced extensions in playout durations under bad conditions, only 25% of the UDP durations were intolerable and no UDP playout reached a doubled duration in this investigation.

The effective capacities reported by the wireless MAC layer are significantly below the capacity the wireless MAC layer is expected to support, and the measured encoding rate for the streaming video, while lower than the wireless capacity, is higher than the effective capacity. The use of multiple encoding levels in a video clip partially alleviates this problem, but significant improvements to streaming performance under bad wireless conditions may require new techniques that identify and adapt to challenging wireless transmission situations.

Understanding packet and frame burst loss behavior is also critical to improving multimedia streaming encoding mechanisms designed to protect, correct or conceal video frame errors. Unfortunately, our tool set was unable to capture error bursts across layers. Developing measurement techniques to capture error bursts during real streaming events remains an important item for future research. Another missing component to improving the strategies used by video servers to adjust to volatile wireless network conditions is a better understanding of when and how a video server decides to do media scaling. Ongoing research is to measure the media scaling reaction of media players to changes in wireless network conditions.

Finally, two other commercial applications, Real Media and QuickTime, are also major contributors to streaming Internet traffic. However these servers probably behave differently than Media Player and investigations with customized measurement tools for these two application suites are also possible future work.

## References

1. G. Bai and C. Williamson. The Effects of Mobility on Wireless Media Streaming Performance. In *Proceedings of Wireless Networks and Emerging Technologies (WNET)*, pages 596–601, July 2004.
2. A. Balachandran and G. Voelker. WRAPI – Real-time Monitoring and Control of an 802.11 Wireless LAN. Technical report, CS at UCSD, 2004.
3. J. Boyce and R. Gaglianella. Packet Loss Effects on MPEG Video sent over the Public Internet. In *Proceedings of ACM Multimedia*, pages 181–190, Bristol, U.K., Sept. 1998.
4. J. Chung, M. Claypool, and Y. Zhu. Measurement of the Congestion Responsiveness of RealPlayer Streaming Video Over UDP. In *Proceedings of the Packet Video Workshop (PV)*, Nantes, France, Apr. 2003.
5. A. Goel, C. Krasic, K. Li, and J. Walpole. Supporting Low Latency TCP-Based Media Streams. In *Proceedings of the Tenth International Workshop on Quality of Service (IWQoS)*, May 2002.
6. M. Li, M. Claypool, and R. Kinicki. MediaPlayer versus RealPlayer – A Comparison of Network Turbulence. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 131 – 136, Marseille, France, Nov. 2002.
7. M. Li, M. Claypool, R. Kinicki, and J. Nichols. Characteristics of Streaming Media Stored on the Web. *ACM Transactions on Internet Technology (TOIT)*, 2004. (Accepted for publication).
8. J. Nichols, M. Claypool, R. Kinicki, and M. Li. Measurements of the Congestion Responsiveness of Windows Streaming Media. In *Proceedings of the 14th ACM International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV)*, June 2004.
9. J. Padhye, V. Firoiu, D. Towsley, and J. Kurose. Modeling TCP Throughput: A Simple Model and Its Empirical Validation. In *Proceedings of ACM SIGCOMM*, Vancouver, British Columbia, Canada, 1998.
10. R. Prasad, M. Murray, C. Dovrolis, and K. Claffy. Bandwidth Estimation: Metrics, Measurement Techniques, and Tools. *IEEE Network*, November-December 2003.
11. B. Wang, J. Kurose, P. Shenoy, and D. Towsley. Streaming via TCP: An Analytic Performance Study. In *Proceedings of ACM Multimedia*, Oct. 2004.
12. Y. Wang, M. Claypool, and Z. Zuo. An Empirical Study of RealVideo Performance Across the Internet. In *Proceedings of the ACM SIGCOMM Internet Measurement Workshop (IMW)*, pages 295 – 309, San Francisco, California, USA, Nov. 2001.