Weather Forecasting - Predicting Performance for Streaming Video over Wireless LANs

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

ABSTRACT

The growth of wireless LANs has brought the expectation for high-bitrate streaming video to wireless PCs. However, it remains unclear how wireless channel characteristics impact the quality of streaming video sent over wireless LANs. This paper presents results from experiments that stream commercial video over a wireless campus network. By analyzing the streaming video quality and capturing wireless LAN characteristics across network and wireless link layers, "weather forecasts" are created such that selected wireless LAN performance indicators might be used to predict the streaming video quality. Furthermore, a quantified measurement of accuracy is presented to evaluate the effectiveness of individual weather forecasts. The paper evaluates six distinct weather forecasts four different streaming configurations including TCP and UDP streaming, and single and multiple-level encoded videos. The results show that the wireless Received Signal Strength Indicator (RSSI) and average wireless link capacity are the most accurate indicators to predict the performance of streaming video over wireless LANs. The weather forecast philosophy can be beneficial for adapting video streaming in wireless LAN environments.

Categories and Subject Descriptors: C.2.m [Computer-Communication Networks]: Miscellaneous

General Terms: Measurement, Performance, Design.

Keywords: Streaming Media, Wireless, IEEE 802.11.

1. INTRODUCTION

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 use the relationships between wireless link measurements to predict the performance of streaming media applications. Thus, predicting the performance of high-demand applications is analogous to weather forecasting. However, while meteorologists attempt to provide accurate weather predictions using well-known predictors, such as temperature and humidity, network practitioners do not yet have effective methods to forecast the performance for streaming video over wireless LANs as a function of reliable wireless LAN characteristics.

Our earlier work has shown that streaming products such as RealNetworks and Windows Streaming Media use network probes to provide estimates of the underlying network characteristics prior to making key decisions about the exact nature of the video stream sent over the network. However, current techniques do not adapt to wireless characteristics such as frame loss rate, signal strength, or link layer bitrate to protect the quality of video streams from bad wireless conditions.

A primary goal of this investigation is to correlate wireless link layer behavior and network layer performance with streaming video application layer performance. Application layer measurement tools [6] are combined with customized network layer measurement tools and publicly available IEEE 802.11 measurement tools to conduct wireless experiments and integrate the measurement results. Seeking the relationships between wireless network indicators and video performance, this study evaluates the effectiveness of several wireless network condition predictors for forecasting streaming video performance.

The remainder of the paper is organized as follows: Section 2 describes the methodology used to obtain video measurements on a wireless LAN; Section 3 presents the results from the experiments and describes how the weather reports are constructed; Section 4 depicts detailed wireless weather reports; and Section 5 summarizes the paper and presents possible future work.

2. METHODOLOGY

2.1 Tools

The strength of this investigation is concurrent use of measurement tools at multiple levels in the network protocol stack to seek the correlation between wireless transmission characteristics and the performance of streaming video. 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.

At the application layer, 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: video frame rate, encoded bitrate, playout bitrate, time spent buffering, frames lost, frames recovered, etc.

For network layer performance measures such as round-

NOSSDAV'05, June 13–14, 2005, Stevenson, Washington, USA.

Table 1: Measurement Tools

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

trip time and packet loss rate along the stream flow path, an internally developed tool, *UDP ping*, was used. Preliminary experiments revealed that a constant ping rate could not be maintained by the standard ICMP *ping* provided by Windows XP in some poor wireless conditions where 10 seconds and longer round-trip times were recorded. Thus, a customized ping tool using application-layer UDP packets was built to provide constant ping rates, configurable ping intervals and packet size.

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: Received Signal Strength Indicator (RSSI), frame retransmission counts and failures, and information about the specific wireless access point (AP) that handles the wireless last hop to the client. Additionally, *typeperf*, a performance monitoring tool built into Windows XP, collected processor utilization and network data including received bitrate and the current wireless capacity target.

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 on the test laptop. Given that the streaming videos consumed about 35% of the processor time, the assumption is the measurement tools do not significantly effect the performance of the streaming videos 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 runs 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[1] chipset. The WPI wireless LAN uses Airespace[2] APs and provides IEEE 802.11 a/b/g wireless service for all the experiments.

Two distinct video clips, one with high motion and the other with low motion, were used in this study. Both clips were encoded at 353×288 resolution and 30 frames per second with a duration of approximately two minutes.[3] Analysis of the two clips shows no statistically significant differ-

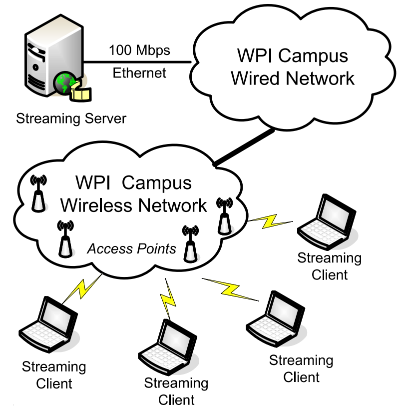

Figure 1: WPI Campus Network

ences in performance, so all subsequent analysis combines the measurement data obtained for both clips.

Broadly, there are two classes of videos stored on the Web, those with a single encoded bitrate level and those with multiple encoded bitrate levels [7]. With a single encoded bitrate level, a streaming media server is limited in its ability to adapt to changes in the network weather. An analogy is going outside without a coat; if it gets cold, there is no easy way to warm up. With multiple encoded bitrate levels, a streaming media server can adapt the video streamed to any changing network conditions. An analogy is going outside wearing multiple layers of clothing; if it gets warm, you can take off layers to remain comfortable.

Therefore, during this investigation, two distinct versions of each video were encoded using Windows Media Encoder v9.0: a single level version of the video encoded at 2.5 Mbps and a multiple level version that includes eleven encoding layers (2.5 Mbps, 1.0 Mbps, 700 Kbps, 500 Kbps, 300 Kbps, 250 Kbps, 128 Kbps, 93 Kbps, 45 Kbps, 32 Kbps, and 19 Kbps) such that the streaming server has the opportunity to do media scaling to dynamically choose the encoded clip to stream based on the network conditions.

As previous work has shown both UDP and TCP are both used for streaming [4], each of the four video instances was streamed using TCP and repeated using UDP to capture the effect of transport protocol choice on streaming performance.

2.3 Experiment Design

Each experiment consisted of streaming videos under eight different conditions (2 clips $\times$ 2 versions (Single Level and Multiple Level) $\times$ 2 transport protocols (UDP and TCP)) to a stationary, wireless laptop. While each video was streamed, the client initiated UDP ping requests to determine round-trip time and packet loss rates. 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.

On each floor of the building, an AP was selected to interact with the client laptop. It was found that the selected video clips could be played back at full-motion quality at

[1]http://www.broadcom.com/
[2]http://www.airespace.com/
[3]The median duration of video clips stored on the Web is about 2 minutes [7].

all locations where the RSSI was above $-65dBm$. At locations where the RSSI was less than $-65dBm$, the video performance was inconsistent. Thus, the experiments were designed to gather more data in areas where performance was inconsistent. A weather analogy is the need to be precise on the temperature near freezing to be able to predict if the precipitation is rain, sleet or snow, while prediction (of rain) is easy when the temperature is in the 40+ degree range. Preliminary experiments found three laptop reception locations for each AP, representing good, fair, and bad reception locations (as indicated by the display status in Windows XP).

Streaming performance over a wireless network depends upon the prevalent network conditions. To reduce the variability in the network conditions, all experiments were conducted during the University's winter break (December 23-25, 2004 and December 29-30, 2004). During the testing periods, there was only occasional network activity and virtually no other wireless users. Each experiment was repeated five times at the 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 stream runs.

2.4 Weather Forecast

Using a forecasting analogy emphasizes the importance of predicting streaming video quality (the weather) given measurements of current network conditions. For a specific quality metric (the weather *prediction*), the forecasting goal is to find a measurable network parameter (a weather *predictor*) that accurately estimates streaming video quality. However, quality metrics differ in their sensitivity to lower-layer predictors and streaming server choices such as transport protocol or encoding method. Thus, weather predictors will differ in their effectiveness in mapping to distinct prediction levels that yield a reliable forecast. This paper analyzes data from a streaming video measurement study to determine reliable lower-layer predictors for predicting the frame rate of Windows Streaming Media over a wireless LAN environment.

The first step in the forecasting process consists of dividing the numerical quality measurements for a weather prediction into three regions: *Good* (Sunny), *Edge* (Cloudy) and *Bad* (Rainy). Then, for each weather predictor, the ordered predictor samples are clustered into 10 equally-populated bins to determine the fraction of samples in each bin in the Good, Edge and Bad regions. Finally, a weather forecast map is created using the median predictor value in each bin and the fraction of Good, Edge and Bad sample points per bin.

An ideal forecasting map for streaming video weather is likely to make good predictions over the complete range of the predictor. To accurately distinguish weather regions, the map needs significant separation of the regions when viewed vertically. Namely, a weather map with large stretches of vertical overlap of the Good, Edge and Bad regions cannot accurately predict the video quality.

We define weather map effectiveness (E) as the fraction of the range of the weather predictor that is likely to produce accurate predictions:

$$E = \frac{R_{effective}}{R_{all}} \tag{1}$$

$R_{effective}$ is the range of the predictor that provides more than a 50% chance of having either Good or Bad performance. R_{all} is the observed predictor range using as endpoints the median of the first and last bins. Any predictor measurements less than the median of the first bin or greater than the median of the last bin are removed as outliers. For some wireless network predictors (e.g., round-trip time) the theoretical sampling space is infinite. Thus, this definition bounds the effective sample space to observed values minus a few outliers. E provides a method for a normalized comparison of the effectiveness of different weather predictors. Since an equal number of samples per predictor bin maintains reasonable sample density for computing the quality fractions, this approach provides a more accurate prediction across the predictor range than using the linear bins in a standard histogram.

An E of 1 indicates a perfect indicator where the weather map provides effective predictions (more than 50% chance of having Good quality or more than 50% chance of having Bad quality) over the entire practical range of the weather predictor. An E of 0 implies a useless indicator whereby the map does not effectively predict the weather for any portion of the weather predictor's practical range.

3. RESULTS

Due to wireless connection failures that resulted in abnormal terminations, ten data sets were removed from the set of 360 streaming runs. Thus, 350 video streaming runs are included in the analysis.

3.1 Weather Predictors

The weather predictors used in this research are all measurements taken from our tools and include: the physical layer Received Signal Strength Indicator (RSSI), wireless link capacity, MAC layer retry fraction, IP loss rate, round-trip time (RTT) and throughput.

Aguayo et al.[1] suggest that signal strength alone is not an accurate indicator of performance for some wireless applications. Figure 2 presents the relationship in this study between wireless connection capacity and wireless RSSI, with a second order best-fit polynomial curve for reference. Since the wireless network capacity adapts based on RSSI, the strong relationship shown with RSSI is not surprising. Conversely, Figure 3, shows that upstream wireless layer retry fraction is not strongly correlated with the RSSI since the retry fraction is also affected by the network traffic load.

3.2 Weather Prediction

For weather prediction, the average frame rate, one of the fundamental measures of video performance, is used as the measure of video quality. The standard frame rate for full-motion video is 24 to 30 frames per second (fps). At these speeds, the human eye perceives movement as continuous, without seeing individual frames. A common frame rate for computer video that approximates full-motion video is 15 fps. To most people, a 15 fps video flows smoothly, although for some videos, it will not appear quite as fluid as it would at a higher frame rate. A video looks choppy if the frame rate is lower than 15 fps. Using these guidelines, video quality is partitioned into three distinct regions: Bad (less than 15 fps), Edge (between 15 and 24 fps) and Good (more than 24 fps). Figure 4 shows the cumulative distribution function (CDF) of the average frame rates for all the

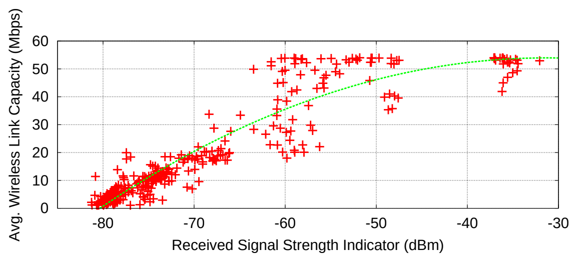

Figure 2: Average Wireless Capacity versus RSSI

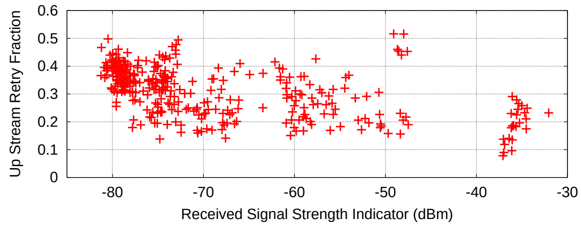

Figure 3: Upstream MAC Layer Retry Fraction versus RSSI

experimental runs, with arrows depicting the Good, Edge and Bad regions.

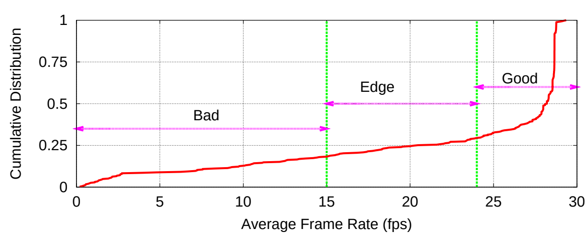

Figure 4: CDF of Frame Rate

The coefficient of variation (CoV) of the frame rate was also considered as a predictor of streaming video quality, but analysis showed CoV to provide the same prediction effectiveness as average frame rate. Analysis using alternate video quality metrics, such as buffering count, media scaling count, and video image quality is left as future work.

4. ANALYSIS

All the analysis presented uses the average video frame rate for weather prediction.

Figure 5 shows a forecasting weather map where RSSI is the weather predictor. The (unlabeled) horizontal illustrations above all maps are a visual histogram of RSSI samples that indicate the data sample density. The Good (Sunny) and Bad (Rainy) regions are separated by the Edge (Cloudy) quality area.

This weather map can be used for weather forecasting as follows: If the RSSI is -60 dBm, there is a 100% chance for Sunny weather (24-30 fps). If the RSSI is -75 dBm, there is a 75% chance of Sunny weather, about a 20% chance of Cloudy weather (15-23 fps) and a 5% chance of Rainy weather (less than 15 fps). If the RSSI is -80 dBm, it is likely to Rain.

The lack of a large vertical overlap between the three areas implies RSSI is a good predictor of average video frame rate. In the RSSI range from -80 dBm to -36 dBm, the only region that does not provide clear predictions of Good or Bad performance is between -79 dBm and -78 dBm. An RSSI lower than -79 dBm forecasts Rain is likely (the probability of a Bad frame rate is 50+%), while an RSSI higher than -78 dBm forecasts likely Sunny weather (the probability to get a Good frame rate is 50+%). The region where RSSI is greater than -68 dBm strongly forecasts Sunny weather.

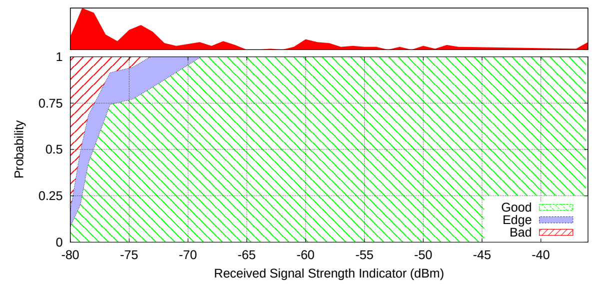

Figure 5: Frame Rate Prediction by RSSI

Average wireless capacity is the predictor for the weather map in Figure 6. Similar to the previous result, the lack of significant vertical overlap in the map suggests average wireless capacity is also an effective predictor of frame rate. In the sampling range from 0 to 54 Mbps, an average wireless capacity greater than 5 Mbps forecasts a high likelihood of Good weather, while a capacity greater than 18 Mbps always forecasts Good weather. Given the maximum encoding bit rates of 2.5 Mbps for the videos used in the experiments, the performance degradation in the region between 2.5 to 5 Mbps is not only due to capacity, but may be due to the variance of link capacity. Figures 7 and 8 demonstrate that even with high average link capacity, the variation in capacity can be high enough to degrade the video frame rate. The link capacity variance may cause upper layer congestion. In the case of TCP streaming, the sender might reduce to a lower sending rate, while a UDP stream may suffer from bursty packet drops as the AP queue fills.

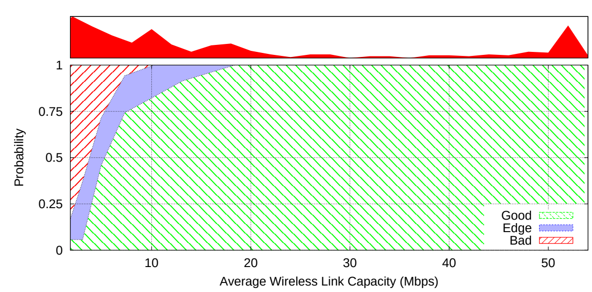

Figure 6: Frame Rate Prediction by Average Wireless Capacity

Figure 9 provides a forecasting weather map using the wireless layer retry fraction as the predictor. As the wireless layer retry fraction increases over the 16% to 44% range, the probability of Good weather slowly decreases. Moreover, the vertical overlap between Good, Edge and Bad over much of the x-axis suggests wireless layer retry fraction is not an effective predictor of video frame rate.

IP packet loss rate is the predictor for the weather map in Figure 10. As with wireless retry fraction, the IP packet

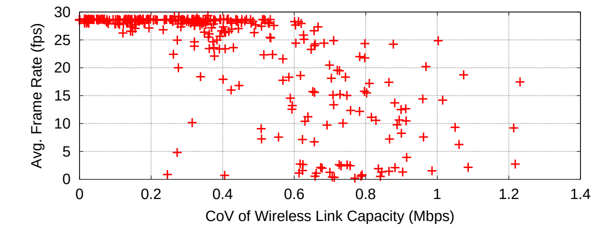

Figure 7: Frame Rate versus Coefficient of Variation of Wireless Link Capacity

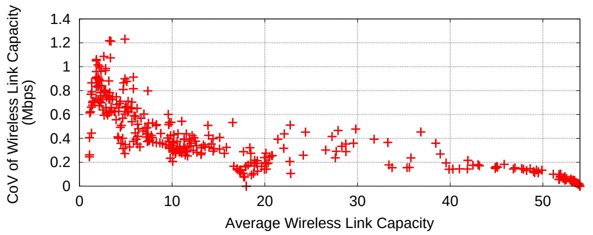

Figure 8: Coefficient of Variation of Wireless Link Capacity versus Wireless Link Capacity

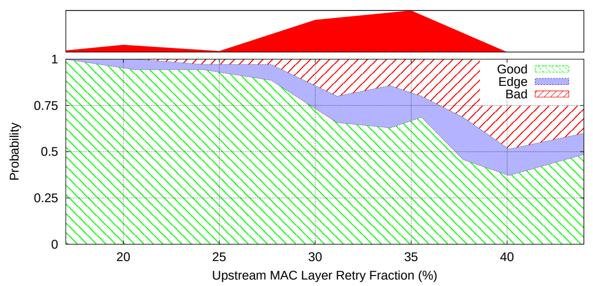

Figure 9: Frame Rate Prediction by Upstream Wireless Layer Retry Ratio

loss rate is not effective for forecasting video frame rate. Only when the loss rate is under 2% or over 16% is a single forecast likely.

Note, the fact that the IEEE 802.11 data link layer retransmits lost frames up to 7 times [3] significantly reduces the number of lost data link frames and also lowers the number of IP packet losses. Comparing the wireless layer retry histogram (the thin, horizontal illustration at the top of each figure) at the top of Figure 9 with the IP packet loss rate histogram at the top of Figure 10, one sees that the density of the samples has shifted from 25%-40% for wireless retries down to less than 10% for IP packet loss rate.

This investigation also considered round-trip time as a weather predictor for forecasting performance for TCP and UDP video streams separately. Due to space constraints, the weather maps cannot be shown, but the results imply that round-trip time is not a good choice as a weather predictor for average frame rate. Similarly, throughput was analyzed as a weather predictor for both multiple and single level videos and was also shown to be ineffective in forecasting the wireless weather. However, while not presenting the respective weather maps, a closer look below the surface is shown with a few appropriate scatter plots.

Figure 11 presents scatter points for throughput in different streaming setups: multiple level TCP streaming, multiple level UDP streaming, single level TCP streaming, and single level UDP streaming. Each graph has a best-fit line for visual reference and to allow observation of the y-intercepts at the left edge of the lines. Comparing Figure 11(a)-11(b) to Figure 11(c)-11(d), one sees that multiple level encoded video sustains frame rates of 10+ fps even for very low throughput, while single level encoded video has frame rates near 0 fps at the same throughput. Conversely, reviewing frame rates for the different protocols shows TCP streaming maintains a higher average frame rate than UDP streaming for low throughput. However, TCP streaming also suffers from long buffering times and a higher frequency of re-buffer events as shown in [5].

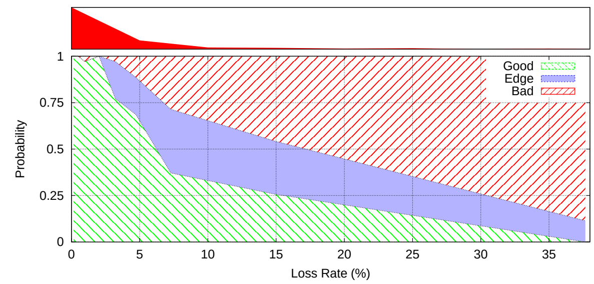

Figure 10: Frame Rate Prediction by IP Packet Loss Rate

These experiments include four distinct video configurations that are analyzed with weather maps based on the six distinct weather predictors. The number of samples in each experimental category is shown in Table 2. Table 2 indicates that encoding video with multiple levels (versus only a single level) results in fewer Bad frame rates, with many Bad rates having been moved into the Edge frame rate region. Furthermore, TCP streaming provides slightly more Good frame rates overall and for multiple level encoding than does UDP streaming.

Table 2: Experiments Categorized by Frame Rate

Setup		Good	Edge	Bad	Total
Multiple level	TCP	73	5	8	86
	UDP	50	25	10	85
	Subtotal	123	30	18	171
Single level	TCP	62	7	20	89
	UDP	62	3	25	90
	Subtotal	124	10	45	179
All	TCP	135	12	28	175
	UDP	112	28	35	175
	Subtotal	247	40	63	350

The weather maps for all of the configurations and predictors are not included in this paper due to lack of space, but more can found in [8]. A summary of the four categories and corresponding effectiveness measurement, E (Equation 1), are provided in Table 3 (sorted in decreasing order of effectiveness). The weather map of the predictors with bold E value in the table are analyzed in this paper.

From Table 3, RSSI and average wireless link capacity are effective predictors of video performance for all streaming setups. Predictors such as round-trip time and throughput are effective for single level encoded video but are ineffective for multiple level encoded video. Finally, forecasting performance for videos encoded with a single level is easier than for videos encoded with multiple levels. This is likely

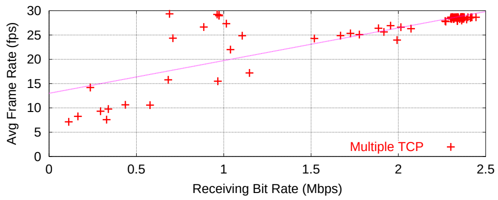

(a) Multiple Level TCP Streaming

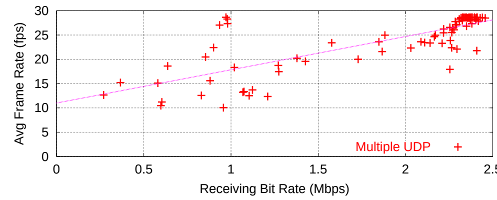

(b) Multiple Level UDP Streaming

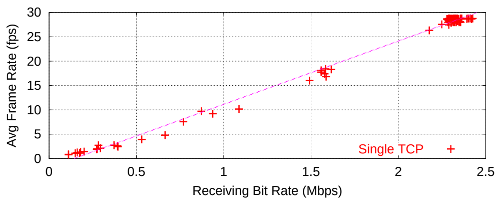

(c) Single Level TCP Streaming

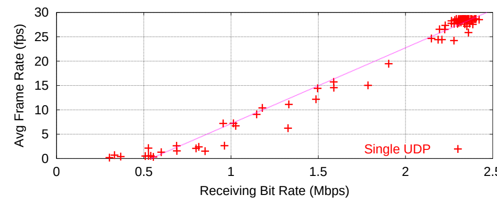

(d) Single Level UDP Streaming

Figure 11: Comparison of Throughput versus Average Frame Rate

because a video with multiple levels of encoding may adapt better to the network weather and yield Good performance, while in the single level encoding case, there is only Bad weather.

Table 3: Effectiveness of Weather Maps

Predictor	All	Multiple	Single	TCP	UDP
RSSI	**0.98**	0.96	0.99	0.99	0.96
Capacity	**0.97**	0.95	0.99	0.97	0.94
Retry rate	**0.75**	0.76	0.81	0.79	0.59
Loss rate	**0.71**	0.69	0.98	0.79	0.89
RTT	0.54	0.35	0.85	0.83	0.94
Throughput	0.47	0.31	0.82	0.59	0.66

5. CONCLUSIONS

This study uses streaming wireless experiments to investigate the relationship between streaming video performance and wireless network behavior. 360 videos were streamed in carefully designed experiments over multiple access points and multiple network conditions to accurately capture performance for wireless locations where streaming is a challenge.

The main analysis vehicle was generation and interpretation of weather maps to forecast streaming video performance. A quantifiable measure of effectiveness is presented allowing comparison of the value of individual weather maps. By considering weather maps for six distinct predictors in four different experimental setups, this research makes several key contributions.

First, the wireless RSSI and average wireless capacity are effective predictors of video frame rate. Second, even predictors that are not effective for forecasting video performance often provide weather maps that have regions of accurate performance prediction. For example, IP packet loss rate predicts high video frame rates when loss rates are less than 2%. Third, the effectiveness of individual predictors varies for different video configurations. For example, multiple level encoding improves video performance over single level encoding for poor wireless conditions, and TCP streaming improves frame rates compared with UDP streaming in the same regions. These findings can improve rate adaption schemes for streaming video over dynamic wireless LAN environments.

Future research includes incorporating knowledge derived from the weather maps into a dynamic video system. Additional weather maps can be developed based on combined weather predictors, such as RSSI and retries or even retries and IP packet loss. Weather maps with different predictions, such as buffering time, re-buffering events and image quality need to be investigated. Additional future work may include investigating performance when there is link contention and competing traffic in wireless networks.

6. REFERENCES

[1] D. Aguayo, J. Bicket, S. Biswas, G. Judd, and R. Morris. Link-level Measurements from an 802.11b Mesh Network. In *Proceedings of ACM SIGCOMM*, Portland, OR, USA, Sept. 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] I. C. S. L. M. S. Committee. IEEE 802.11, 1999 Edition, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.

[4] J. V. der Merwe, S. Sen, and C. Kalmanek. Streaming Video Traffic: Characterization and Network Impact. In *Proceedings of the 7th WCW*, Boulder, CO, USA, Aug. 2002.

[5] F. Li, J. Chung, M. Li, H. Wu, M. Claypool, and R. Kinicki. Application, Network and Link Layer Measurements of Streaming Video over a Wireless Campus Network. In *Proceedings of the 6th PAM*, Boston, Massachusetts, USA, Apr. 2005.

[6] M. Li, M. Claypool, and R. Kinicki. MediaPlayer versus RealPlayer – A Comparison of Network Turbulence. In *Proceedings of the ACM SIGCOMM 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] M. Li, F. Li, M. Claypool, and R. Kinicki. Weather Forecasting - Predicting Performance for Streaming Video over Wireless LANs. Technical Report WPI-CS-TR-05-03, CS Department, Worcester Polytechnic Institute, Feb. 2005.