

WPI-CS-TR-05-03

Feb. 2005

Weather Forecasting - Predicting Performance for Streaming
Video over Wireless LANs

by

Mingzhe Li
Feng Li
Mark Claypool
Robert Kinicki

Computer Science
Technical Report
Series



WORCESTER POLYTECHNIC INSTITUTE

Computer Science Department
100 Institute Road, Worcester, Massachusetts 01609-2280

Weather Forecasting - Predicting Performance for Streaming Video over Wireless LANs

Mingzhe Li, Feng Li, Mark Claypool, Robert Kinicki
{lmz,lif,claypool,rek}@cs.wpi.edu
Computer Science Department
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 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 over the space of 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.

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 find the relationships between wireless link measurements and the performance of streaming media applications. Therefore, 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 WLAN characteristics.

Previous work [10] 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 weather.

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 [7] 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.

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

At the application layer, the WPI Wireless Multimedia Streaming Lab has experience measuring video client and server performance [3, 7, 10, 12]. An internally developed measurement tool, called *Media Tracker* [7], 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-trip time and packet loss rate along the stream flow path, an internally developed tool, *UDP ping*, was used. Prelim-

inary 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 streaming downloads consumed about 35% of the 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 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 Broadcom¹ chipset. The WPI wireless LAN uses Airespace² APs and provides IEEE 802.11 a/b/g wireless service for all the experiments.

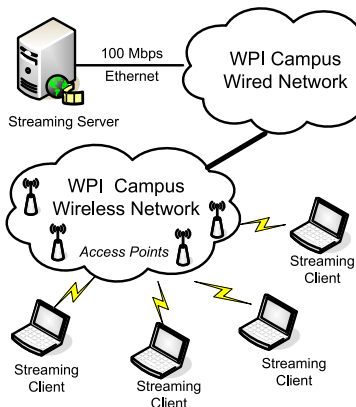


Figure 1: WPI Campus Network

¹<http://www.broadcom.com/>

²<http://www.airespace.com/>

Two distinct video clips, one with high motion and the other with low motion, were used in this study. Both clips were encoded to run 353×288 resolution and 30 frames per second with a duration of approximately two minutes.³ However, an earlier analysis of the two video clips concluded that these two clips did not significantly impact performance over a wireless network.

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 [8]. With a single encode 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 can adapt the video streaming to network condition changes. 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 streamed to every client location: a single level version of the video encoded at 2.5 Mbps 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 the network weather.

As previous work has shown both UDP and TCP are frequently used for streaming [5], 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 side.

On each floor 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 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 natural weather analogy is the need to be precise on the temperature near freezing to be able to predict rain/sleet/snow, while prediction 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 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 con-

³The median duration of video clips stored on the Web is about 2 minutes [8].

ducted 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 around. 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

The term “weather forecast” is used to emphasize the importance of predicting (forecasting) the quality of a streaming video (the weather) given measurements of a network condition. For a given quality metric (the weather *prediction*), the mission of a weather forecast is to use a measurable network parameter (the weather *predictor*) to predict the quality of streaming video. However, different predictions may be more sensitive to certain predictors than others, and may vary under different network environments, such as transport protocols or encoding method. Thus, one weather forecast may not be equally effective for other quality metrics or network environments. This study seeks good weather predictors for one specific weather prediction (Windows Streaming Media frame rate) in a predefined network environment (a wireless LAN).

There are a number of measurable network parameters that can be used as weather predictors. For instance, the wireless Received Signal Strength Indicator (RSSI) is one of the most widely used wireless performance indicators, therefore, can naturally be selected as a weather predictor to predict the streaming quality in wireless network.

This WLAN weather forecasting proceeds in three steps. First, divide the numerical quality measurement for a weather prediction into three regions, namely *Good* (Sunny), *Edge* (Cloudy) and *Bad* (Rainy). Second, after selecting a weather predictor, divide the predictor samples into 10 bins with an equal number of samples and determine the fraction of sample points in the Good, Edge and Bad regions for each bin. Finally, a map is draw to show the weather forecast by using the median value of the predictor in each bin and the fractions of Good, Edge and Bad.

An ideal map that could be used with some accuracy to predict the streaming video weather over a wireless terrain should have vertically separated Good, Edge and Bad regions. This allows a measurement of the weather predictor on the axis to accurately predict the streaming video performance for that condition. A weather map that has vertical overlap of the Good, Edge and Bad regions means measurement of the weather predictor in the overlap region cannot always be used to accurately predict video streaming performance for that condition.

The effectiveness (E) of a weather map is defined as the fraction of the range of the weather predictor that can be used to accurately produce predictions:

$$E = \frac{R_{effective}}{R_{all}} \quad (1)$$

$R_{effective}$ 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 range of the predictor using as endpoints the median of the first bin and the median of the last bin. Any predictor sample values less than the median of the first bin and greater than the median of the last bin

are removed as outliers. For some wireless network predictors, such as round-trip time, the theoretical sampling space is infinite. Thus, this definition of a practical range bounds the sampling space to observed values minus a few outliers. This definition of E then provides a method of normalizing across different predictors, such as round-trip time and signal strength, allowing comparison of the effectiveness of different weather maps. Note that having an equal number of samples in each bin maintains a reasonable density of samples for computing the quality fractions, thus providing a more accurate prediction across the predictor range than might uniform sized bins as occurs in a typical histogram.

The value of E ranges between 0 and 1. An E of 1 implies a perfect indicator, which means that 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, which means that the weather map does not provide effective predictions for any part of the practical range of the weather predictor.

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.

as discussed in Section 2.1. At the wireless layer, the physical layer Received Signal Strength Indicator (RSSI) is an easily accessible and widely used performance indicator provided by most wireless card drivers and operating systems.

The wireless connection capacity which adapts from 1, 2, 5.5 ... up to 54 Mbps and the wireless layer errors (and subsequent frame retransmissions) also impact performance.

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.

From the end host point of view, the measurable wireless layer retry fraction is on the upstream and might differ from the downstream retry fraction from Access Point (AP) to the end host. Since typical APs do not provide a means to measure the downstream wireless layer retry fraction, the assumption is the contention and error status on the wireless channel is symmetric.

Network layer measurements, such as the IP packet loss rate and round-trip time (RTT) are also selected as weather predictors. In addition, the TCP-Friendly rate, which combines the effects of both packet loss rate and round-trip time is also used as a weather predictor. The TCP-Friendly rate,




Figure 2: Average Wireless Capacity versus RSSI




Figure 3: Upstream MAC Layer Retry Fraction versus RSSI

T Bps, for a connection can be computed by [11]:

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

with packet size s in bytes, round-trip time R in seconds and packet drop rate p . The TCP retransmission timeout t_{rto} is set to four times round-trip time by default. For each video clip for each run, Equation (2) 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.

Finally, the throughput measured on the end host is also used as a weather predictor. Previous research [6] shows that, for the same throughput, videos encoded with multiple levels often achieve better performance than videos encoded with a single level. Using throughput as an predictor may provide accurate weather maps for videos encoded with single versus multiple levels.

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 this speed, 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 a cumulative distribution function (CDF) of the average frame rates for all the experimental runs, with arrows depicting the Good, Edge and Bad regions.

A high average frame rate video can still appear choppy




Figure 4: CDF of Frame Rate




Figure 5: Coefficient of Variation of Frame Rate versus Frame Rate

if the variance in the frame rate is high. Therefore, a second weather prediction, the coefficient of variation (CoV) of the frame rate was analyzed. However, Figure 5 shows the coefficient of variation of frame rate strongly related to average frame rate. The CoV of frame rate remains high when the average frame rate is low, while a higher average frame rate usually implies a lower CoV. This corresponding relationship results in similarly accurate weather maps and the CV of frame rate is not considered as a separate weather prediction.

The coefficient of variation (CoV) of the frame rate was also considered for weather prediction but analysis showed CoV to be highly correlated with average frame rate. Analysis of other weather predictions 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 6 shows a forecasting weather map where the weather predictor is RSSI. The (unlabeled) horizontal illustration above the figure is a visual histogram of RSSI samples that indicates the data sample density for the weather maps. The Good (Sunny) and Bad (Rainy) regions are separated by the Edge (Cloudy) quality area.

The 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 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 likely Rain (the probability of Bad




Figure 6: Frame Rate Prediction by RSSI




Figure 7: Frame Rate Prediction by Average Wireless Capacity

frame rate is 50+%), while an RSSI higher than -78 dBm forecasts likely Sun (the probability to get the Good frame rate is 50+%). The region where RSSI is greater than -68 dBm strongly forecasts Sunny weather.

Average wireless capacity is the predictor for the weather map in Figure 7. Similar to the previous result, the lack of 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 8 and 9 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.

Multiple level encoding can benefit from the adaptive streaming rate with correctly adapting the sending rate to the available capacity [6].

Figure 10 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 used as the predictor for the weather map in Figure 11. As with wireless retry fraction, the IP packet loss rate is not effective for forecasting video frame




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




Figure 9: Coefficient of Variation of Wireless Link Capacity versus Wireless Link Capacity




Figure 10: Frame Rate Prediction by Upstream Wireless Layer Retry Ratio

rate. Only when the loss rate is under 2% or over 16% is a single forecast likely.

This is perhaps explainable considering the predicted TCP-Friendly data rate. A 2% loss rates with a corresponding observed average round-trip time of 16.53 ms, provides a TCP-Friendly rate of 5.2 Mbps, sufficient to transfer the 2.5 Mbps encoded video under TCP. For UDP streaming, the 2% lost packets can be recovered by the Windows Media Service retransmissions, and even ignored for data rate reduction.

Note, that since the IEEE 802.11 data link layer retransmits lost frames up to 7 times [4], this 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 10 with the IP packet loss rate histogram at the top of Figure 11, 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.

Figure 12 and Figure 13 depict weather maps for forecasting video frame rate using the round-trip time as a predictor, for streaming over TCP and UDP, respectively. In both figures,

This investigation also considered using round-trip time as a weather predictor for forecasting TCP and UDP streams




Figure 11: Frame Rate Prediction by IP Packet Loss Rate




Figure 13: Frame Rate Prediction by Round-Trip Time for UDP Streaming




Figure 12: Frame Rate Prediction by Round-Trip Time for TCP Streaming




Figure 14: Frame Rate Prediction by TCP-Friendly Rate

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 also shown to be ineffective in forecasting the wireless weather. However, while not presenting the respective weather maps, we choose to look below the surface a bit to review the appropriate scatter plots.

For streaming over TCP and UDP, the predicted frame rate decreases as the round-trip time increases. However, the decreasing trend is slow, resulting in significant vertical overlap among the Good, Edge and Bad regions. Thus round-trip time not effective for forecasting video frame rate. Notice that the practical range of the round-trip time observed for for UDP streaming is up to 1285 ms, while the practical range observed for TCP streaming is only about 144 ms. This increased round-trip time for UDP streaming might be caused by a combination of a large AP queue size that fills due to the unresponsive UDP sender [6]. For both TCP and UDP streaming, only round-trip times less than 10 ms have a high probability of Good frame rates.

The practical throughput range for both the multiple and single level videos are bounded by the highest encoding bit rate at 2.5 Mbps. From the range of vertical overlap, throughput is not effective in forecasting video frame rate. However, a throughput close to the upper bound at 2.5 Mbps provides a 100% likelihood of a Good frame rate for both multiple and single level videos. In addition, for multiple level encoded videos, a 1 Mbps throughput provides a relatively higher probability of having a Good frame rate than for a single level video. This is because the multiple level video includes an encoding bitrate at 1 Mbps, encoded at 30 fps and a lower quantization.

Figure 14 depicts a weather map for forecasting video

frame rate using the TCP-Friendly rate as a predictor. The TCP-Friendly rate, computed using Equation 2, goes to infinity as the loss rate decreases. Since the maximum encoded rate used for streaming is 2.5 Mbps, the practical maximum for the TCP-Friendly is set to 2.5 Mbps to allow comparison of the TCP-Friendly weather maps with the other weather maps. Over the range from to 2.5 Mbps, the TCP-Friendly rate is a reasonable predictor of video frame rate. Moreover, the probability of having a Good frame rate is high even when the TCP-Friendly rate is less than the video encoding rate, 2.5 Mbps. From the results in [6], video stream will tend to select an encoding bit rate higher than the TCP-Friendly rate when streaming video. This results in a TCP-Unfriendly bitrate when using UDP, and buffer underflows (but still periods of Good frame rates) when using TCP.

The average buffer count is 1.59 and 1.18 for TCP and UDP, respectively.

Figure 15 and Figure 16 depict weather maps for forecasting video frame rate using the throughput measured on the end host as a predictor, for multiple level encoding and single level encoding, respectively.

The result visually confirms early that the media scaling of Windows Media Service with multiple encoding level video can improve the frame rate performance in the wireless environment [6].

Commented out by Mingzhe for space saving

Figure 17 graphs 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 take special note of the y-intercepts at the left edge of the lines. Comparing Figure 17(a)-17(b) to Figure 17(c)-17(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




Figure 15: Frame Rate Prediction by Throughput for Multiple Level Encoding Video




Figure 16: Frame Rate Prediction by Throughput for Single Level Encoded Video

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 [6].

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 rate region. Furthermore, TCP streaming provides slightly more Good frame rates overall and for multiple level encoding than 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 be found in [9]. 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 for all streaming setups. While one predictor might perform well for one setup but poorly for another setup, 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 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 characteristics. Nearly 400 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, an 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 encoding levels 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-buffer count and image quality need to be investigated.

Finally, weather forecasting for other commercial streaming applications, such as Real Media and QuickTime, is also a rich area for future work.


6. REFERENCES




(a) Multiple Level TCP Streaming



(c) Single Level TCP Streaming



(b) Multiple Level UDP Streaming



(d) Single Level UDP Streaming

Figure 17: Comparison of Throughput versus Average Frame Rate

- [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] 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.
- [4] I. C. S. L. M. S. Committee. IEEE 802.11, 1999 Edition, Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications.
- [5] 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.
- [6] 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.
- [7] 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.
- [8] 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).
- [9] 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.
- [10] J. Nichols, M. Claypool, R. Kinicki, and M. Li. Measurements of the Congestion Responsiveness of Windows Streaming Media. In *Proceedings of the 14th NOSSDAV*, June 2004.
- [11] 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.
- [12] Y. Wang, M. Claypool, and Z. Zuo. An Empirical Study of

RealVideo Performance Across the Internet. In *Proceedings of the ACM SIGCOMM IMW*, pages 295 – 309, San Francisco, California, USA, Nov. 2001.

Appendix

Weather forecast for all experiments setups




Figure 18: Frame Rate Prediction by Throughput




Figure 19: Frame Rate Prediction by Round Trip Time

Weather forecast for Mutiple level video




Figure 20: Frame Rate Prediction by Receive Signal Strength Indicator




Figure 21: Frame Rate Prediction by Average Wireless Capacity




Figure 22: Frame Rate Prediction by Upstream MAC layer Retry Ratio




Figure 23: Frame Rate Prediction by IP Packet Loss Rate

Weather forecast for Single level video




Figure 27: Frame Rate Prediction by Receive Signal Strength Indicator




Figure 24: Frame Rate Prediction by Round Trip Time




Figure 28: Frame Rate Prediction by Average Wireless Capacity




Figure 25: Frame Rate Prediction by TCP Friendly Rate




Figure 29: Frame Rate Prediction by Upstream MAC layer Retry Ratio




Figure 26: Frame Rate Prediction by Throughput

Weather forecast for TCP Streaming




Figure 30: Frame Rate Prediction by IP Packet Loss Rate




Figure 34: Frame Rate Prediction by Receive Signal Strength Indicator




Figure 31: Frame Rate Prediction by Round Trip Time




Figure 35: Frame Rate Prediction by Average Wireless Capacity




Figure 32: Frame Rate Prediction by TCP Friendly Rate




Figure 36: Frame Rate Prediction by Upstream MAC layer Retry Ratio




Figure 33: Frame Rate Prediction by Throughput




Figure 37: Frame Rate Prediction by IP Packet Loss Rate

Weather forecast for UDP Streaming




Figure 41: Frame Rate Prediction by Receive Signal Strength Indicator




Figure 38: Frame Rate Prediction by Round Trip Time




Figure 42: Frame Rate Prediction by Average Wireless Capacity




Figure 39: Frame Rate Prediction by TCP Friendly Rate




Figure 43: Frame Rate Prediction by Upstream MAC layer Retry Ratio




Figure 40: Frame Rate Prediction by Throughput




Figure 44: Frame Rate Prediction by IP Packet Loss Rate




Figure 45: Frame Rate Prediction by Round Trip Time




Figure 46: Frame Rate Prediction by TCP Friendly Rate




Figure 47: Frame Rate Prediction by Throughput