

Characterizing High-bandwidth Real-time Video Traffic in Residential Broadband Networks

Ramya Raghavendra, Elizabeth M. Belding

Department of Computer Science, University of California, Santa Barbara

{ramya, ebelding}@cs.ucsb.edu

Abstract—Users are generating and uploading multimedia content to the Internet at an unprecedented rate. Residential broadband networks, however, have low upload capacities and large packet latencies. Wi-Fi networks that are used to access the Internet can suffer from high packet losses and contention latencies. All of these factors can result in poor video quality for residential users. Using packet traces and active measurements from houses, we study video quality in residential scenarios. We analyze the primary factors that contribute to poor performance and compare the performance over both the wireless and the broadband hop. Our measurements show that the upload capacities on the broadband links restrict the video bitrate (and hence the resolution) that can be transmitted. Residential wireless networks, however, have much higher capacities than the broadband links and, despite being densely deployed, do not see extended periods of high utilization. Our measurements shed light on the video transmission quality that is typically achievable from residences and are used to characterize the reasons behind quality deterioration.

I. INTRODUCTION

The popularity of social networking and chat applications have resulted in an increase in the amount of data that users are generating and uploading to the web, compared to web traffic that was predominantly in the downlink direction. The delivery of multimedia content is growing at a tremendous rate, increasing 76% every year on average, with video communication and real-time traffic growth predicted to increase tenfold by the year 2013 [1].

This surge of media rich applications is leading to a “broadband access” gap¹, created by broadband access links, or the edges, that are not growing at the same rate as core routing, switching and transmission capacity. Upgrading access links is expensive and new technologies often take several years to deploy. As a result, access technology can vary dramatically from neighborhood to neighborhood, and even home to home in the same neighborhood.

There have been numerous measurement studies of Wi-Fi networks, broadband networks and the Internet in the past, but an overwhelming majority of them are based on TCP traffic. Video applications typically use UDP at the transport layer, and real-time applications further use protocols such as RTP and RTCP for streaming support. Hence, the conclusions from prior work do not apply directly when studying video streaming performance. Video traffic analysis has been studied over wireless LANs [2] as well as over the Internet [3]. These studies focus on the performance of video streaming

applications that can tolerate a high amount of initial delay due to buffering.

However, there are no comprehensive studies that characterize the performance of interactive video traffic over residential connections. Video streaming in such networks suffers from losses and delays at the wireless link, the access link, as well as the ISP network. In this paper, our goal is to characterize the performance of video traffic at the wireless and access links.

The common consensus among users is that high bitrate video quality suffers in residential networks. Various factors contribute to the degradation in video quality: wireless losses, congestion, poor uplink quality, queuing, large delays and jitters. It is important to understand what factors affect the video quality, and to what extent, in order to build robust video streaming techniques.

In this paper, we analyze the results of 24 hour long monitoring and real-time streaming experiments from eight residences that have broadband Internet connections. Using videos encoded with MPEG-4, and using RTP and UDP protocols for streaming, we study the properties of wireless and the end-to-end links in terms of the bandwidth available for streaming, loss and latency that packets experience, and the effect on streaming quality. We attempt to answer the following questions:

- 1) What video bitrate sessions can be sustained over typical residential networks?
- 2) What factors affect the interactive streaming quality?

The rest of the paper is structured as follows: In Section II, we provide a brief background on architecture of broadband and Wi-Fi networks considered. Section III describes the methodology of our measurement study. Section IV discusses the measurement results. In Section V, we review the related work in the area and we present our conclusions in Section VI.

II. BACKGROUND

In this section, we provide a brief background about the broadband and wireless networks over which the video transmission is evaluated in Section IV.

Residential broadband networks provide residences with ‘last mile’ Internet access. A large and rapidly growing percentage of residential Internet connections are via broadband DSL or cable modem technologies. A recent survey showed that 86% of residential users access broadband Internet through cable or DSL². Figure 1 shows the typical architecture of residential broadband networks. In both the cable and DSL networks, the

¹<http://www.networkworld.com/columnists/2010/020410-johnson.html>

²http://news.cnet.com/8301-1023_3-10450784-93.html




Fig. 1. Architecture of typical residential access networks.

core branches out into a number of central offices that support a cluster of geographically proximal customers. Despite having different physical layer technologies, the ISP core essentially splits into several regional headends (or DSLAMs in the case of a DSL network) and each headend connects a set of homes. A key difference between the cable and DSL access links is that DSL customers have a dedicated connection to the regional center, whereas cable users share the connection to the headend.

Broadband networks are also known to have a large disparity in the upload and download speeds. In particular, the measurement study by Dischinger, *et al.* [4] shows that upstream bandwidths were less than 500 Kbps, even in cases where the downstream bandwidths exceeded 5 Mbps in the 600 residences that were measured. The ratio of downstream to upstream traffic is shown to be high, over a factor of 10 in the case of cable hosts.

Usage studies of broadband links, however, show that these links are typically used only a fraction of the time [5]. The per-user rate limiting of broadband networks ensures that users get a fair share of broadband access when the network usage is high, but when the network is not being used to its capacity, leaves a significant amount of bandwidth on the table.

In residential networks, a common method of accessing broadband networks is using the Wi-Fi network through the wireless access point (AP) deployed in a house. In this paper, we study the video streaming quality over the broadband network, as well as the wireless link on the first hop. Numerous past studies have shown the dense deployment of 802.11 APs in residential areas [6], [7]. Wireless signals often reach across homes, making it possible to sustain TCP and UDP connections with open APs in the neighborhood [8], [?].

III. MEASUREMENT METHODOLOGY

We analyze typical end-to-end performance of real-time video streams in residential networks. To do so, we assess the video performance first over the wireless link to the AP, and then over the access link to the broadband ISP.

Our measurement study is comprised of two parts. First, using passive sniffers, we study the typical wireless usage in homes. Then, using active probes we study the packet loss and latency that video traffic encounters in the wireless and broadband links.

The measurement setup and methodology are described below.

A. Passive Monitoring

In order to assess the wireless connectivity and usage in typical residences, we deploy packet sniffers in eight residences

with cable or DSL Internet connections. The data presented here are from residences in three different towns, four of them being from houses in apartment complexes and the other four being single houses. In general, apartment complexes see denser deployment of APs. The traces are collected for 24 hours at each residence using the packet sniffer tool Wireshark.

Our measurement setup is common to all locations. Two sniffers are deployed in each residence: one in the immediate vicinity of the AP and the other in an area where laptops are commonly used to access the Internet. The nodes are Linux-based laptops configured to be in monitor mode to record frames at the link layer. The sniffers record frames on the same 802.11g channel on which the AP was transmitting.

B. Active Measurements

In addition to passive measurements of wireless usage, we use active measurement sessions from each of the eight residences to study the link properties with respect to capacity, loss rate and latency.

The active measurements are performed in two phases. In the first phase, we probe the wireless and broadband link with packet trains of different rates, using packets of various sizes. The receiver records various properties including the number of packets received and the interpacket delay variation, and this information is used to infer the link properties.

In the second phase, we transmit an encoded video, and compare the received video stream with the transmitted video stream and report the effect on video quality. Video is streamed from a laptop connected to the residence's AP through an 802.11g link, and the AP is connected to the Internet through either a DSL or cable connection. The receiver is placed on a desktop that is connected to a high speed university network. The video is 30 seconds in length and encoded in MPEG-4 format. The video is compressed to different bitrates using ffmpeg. In order to mimic a real-time streaming session, the playout buffer is set to 200ms [9]. Table I lists the measurement details.

Length of each passive trace	24 hours
Number of residences	8
Amount of data collected	4 Gbytes
Number of streaming sessions	1000
Length of each session	30 seconds

TABLE I
DETAILS OF THE TRACE COLLECTED.

IV. MEASUREMENT RESULTS

A. Video Streaming

We measure the video streaming performance in residential neighborhoods. Figure 2 shows the results from 1000 streaming sessions of a 3 Mbps bitrate video using the setup described in § III-B. The x -axis is ordered in ascending order of the video quality. We use PSNR (Peak Signal-to-Noise Ratio) as the metric for video quality estimation. PSNR is the ratio between the maximum possible power of a signal and the power of corrupting noise that affects the fidelity of its representation. The PSNR of a video is well correlated with the perceived

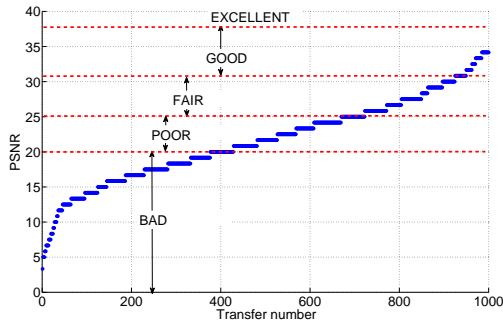


Fig. 2. PSNR values of video streaming sessions from eight residences. The sessions are arranged in ascending order of their PSNR value for visual clarity. The mapping between the PSNR range the MOS based user perceptions are indicated.

quality of video. The relationship between user perception expressed in Mean Opinion Score (MOS) and PSNR and is shown in Table II.

As shown in Figure 2, the video quality spans the entire range of quality from ‘poor’ to ‘good’. In the following subsections, we study the factors that affect the video quality. Specifically, we study the available bandwidth, packet latency, delay variation and packet loss and discuss the implication of these measurements on real-time video streaming.

MOS Rating of video quality	PSNR range
Excellent	> 37
Good	31 – 37
Fair	25 – 31
Poor	20 – 25
Bad	< 20

TABLE II
TABLE MAPPING THE MOS BASED USER PERCEPTION OF VIDEO QUALITY TO THE PSNR RANGE.

B. Bandwidth

It is estimated that Internet video is now approximately one-third of all consumer Internet traffic. Cisco projects that “video communications” traffic will increase tenfold from 2008 to 2013, and real-time video is growing in importance, owing to the popularity of Internet TV [1].

HD quality video transmission places a high demand for bandwidth. As an example, a common high definition resolution of 720p (1280x720 pixels) with a frame rate of 24 per second in raw format with YUV 4:2:2 color model (common for video) has an uncompressed bit rate of over 350 Mbps. Commercially available HD videos from Comcast DOCSIS 3.0, on-demand movies from Apple TV 2.0 and HD downloads from XBOX360 are highly compressed using compression codecs such as H.264 or VC-1 and have a bit rate in the range of 1.5 Mbps - 7 Mbps.

The average broadband speeds are shown to be in the order of 5 Mbps for downloads and 800 Kbps - 1.1 Mbps for uploads³. High resolution video upload and download is restricted by the broadband speeds. While streaming videos make use of

large buffers to counter the slow speeds, real-time streaming is limited in the amount of buffering that can be used and hence can be bottlenecked by broadband links. In the following sections, we discuss the amount of bandwidth available and the stability of the wireless and wired links.

1) *Available Capacity*: We measure the capacity available on the wireless link as well as the total end-to-end capacity. The capacity can be bottlenecked either at the wireless link, due to contention or poor quality links, or at the broadband uplink. Prior measurement studies have shown that the bottleneck in a broadband network is at the access hop, i.e, the last hop from the residence to the ISP network. Measuring the bandwidth available for a user helps predict the maximum bitrate for video streaming that can be sustained.

Figure 3(a) shows the available bandwidth on the broadband link. In a cable network, this refers to the portion of the shared link allocated to a user, while in a DSL network, it refers to the ISP’s cap on the user’s traffic rate. We did not observe any difference in the available broadband bandwidth between houses and apartments, but did observe a difference between the speeds of cable and DSL networks. While the DSL downstream bandwidths are typically higher than cable bandwidths, the gap between the upload and download bandwidths is much wider in cable networks than DSL networks. This implies that an interactive video application such as video chat is more constrained by the low uplink speed in a cable network.

Figure 3(b) shows the measured wireless medium utilization. Each boxplot represents the quartile distribution of utilization values observed across the house or apartment per channel. Figure 3(c) shows the bandwidth available on the wireless link from the experiment node to the AP deployed in the house. The x -axis denotes the location and channel. For example H1 refers to all the measurements that were from single house on channel 1. Similarly, A6 refers to measurements from an apartment on channel 6. In general, we could sustain higher rate flows in houses than in apartments, which we speculate is the result of denser wireless deployment and higher usage in apartments.

Figure 3(d) shows the PSNR obtained on streaming videos of different bitrates. The bitrates are shown on the x -axis and the corresponding PSNRs are plotted on the y -axis.

2) *Bandwidth Stability*: We evaluate the short-term stability of the uplinks by measuring the available bandwidth at 100ms intervals. Understanding the link stability is important in the design of a video streaming solution since loss or late delivery of a single frame can cause disruption for a perceivable length of time, depending on the type of frame lost. Figure 3(e) shows the variation of available bandwidth from a single residence over time at the wired and the wireless links over a five minute interval chosen. When we examine the trace over the entire measurement duration, we see that a high degree of variation is seen in broadband networks, especially in cable networks where residences share the connection to the headend. Wireless networks show variation due to the changing environment and changes in interference and contention from other users and networks or non 802.11 devices that share the spectrum.

The conclusions from the bandwidth study are as follows.

- Wireless networks see high utilization for short periods of

³http://cwafiles.org/speedmatters/state_reports_2009/CWA_Report_on_Internet_Speeds_2009.eps

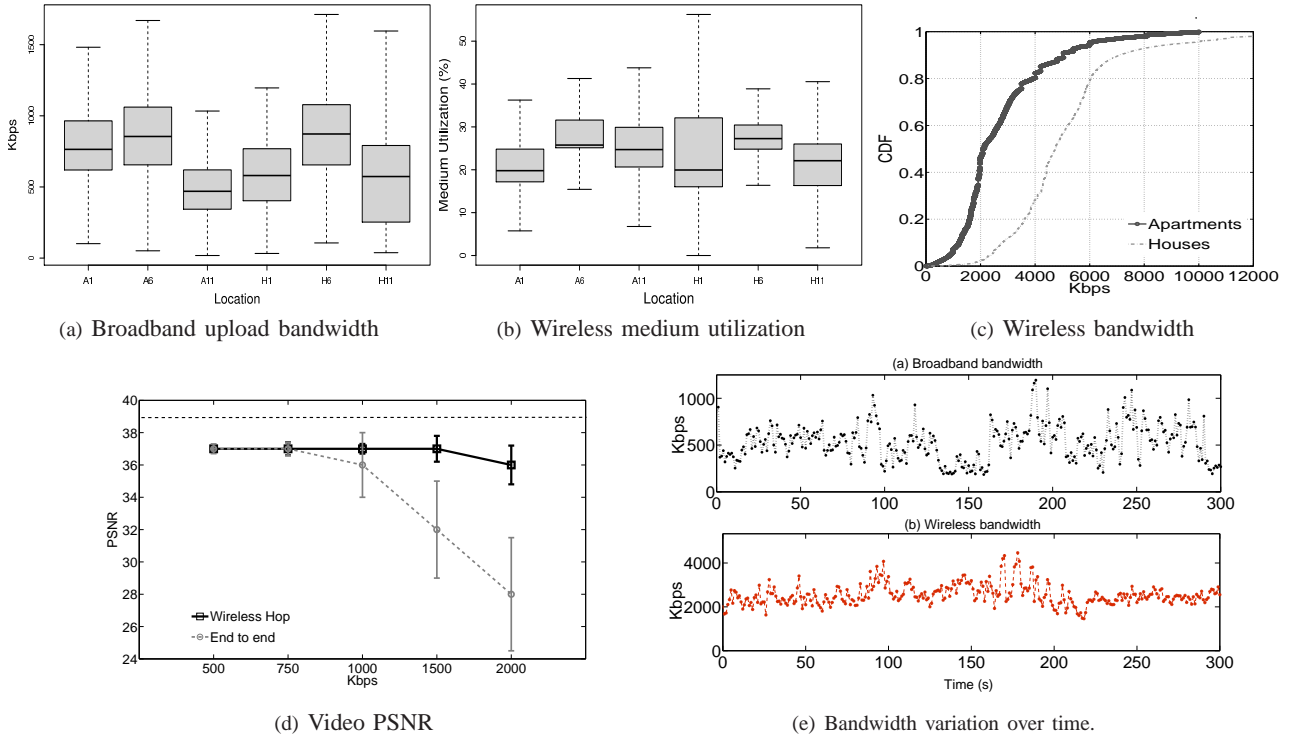


Fig. 3. Available bandwidth on wired and wireless links and resulting video PSNR when streaming 3 Mbps video on these links. Results are from passive and active measurements from eight residences.

time, but for the most part, are not heavily utilized.

- Broadband uploads witness slow speeds that are much lower than the bit rates needed by high-definition videos.
- Broadband links have a high degree of variation in the available bandwidth.

C. Packet Latency

In contrast to data transmission, which is usually not subject to strict delay constraints, real-time video requires bounded end-to-end delay. In real-time video streaming, video frames are played as they are received and packet delay variation, or jitter, degrades the perceptual quality of the video. If a frame arrives late, the players freeze the most recently seen image. When the next frame arrives, it is displayed briefly to preserve the timing for the subsequent frame. A video packet that arrives beyond its playback deadline is useless and can be considered lost.

In order to characterize the extent of packet latencies incurred on broadband networks, we study the end-to-end packet transmission latency (when the link is not saturated) and the queuing delay incurred by sending saturating probes. Packet transmission latency is measured by sending small probe packets and measuring the RTT, shown in Figure 4(a). Figure 4(b) shows the upstream queuing delay by using saturating ICMP probes and measuring the variation in RTT. There is a significant variation in the delays observed, and the variation can be attributed to the difference in cable and DSL uplink properties [10].

D. Packet Delay Variations

We next examine the variation in packet arrival times. In order to compensate for variation in packet arrival times, video players employ buffering at the receiver. Buffering can mask

delays and jitter and render a relatively smooth video, but it introduces an overall delay in the video playback. Real-time video comes with the constraint that buffering is limited; in fact, the maximum delay tolerable for video communications is 200 ms. We plot the jitter measurements from our setup in Figure 5. We define jitter as the difference between the 10th and the 90th percentile value of the RTT values from our latency measurements.


The implications of this analysis are the following:

- Packets experience considerable, but varying, delays on the access link. Some broadband links can have large queues that make real-time traffic infeasible.
- The video receiver should have a playout buffer that accounts for the large broadband jitter values. Since the maximum end-to-end delay is known, a real-time streaming system should choose APs that are likely to deliver a packet in time, taking into account the link delays and jitter.

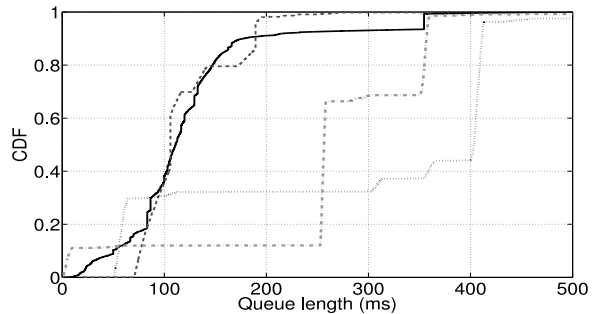
E. Packet Loss

We measure the extent of packet loss and the burst lengths of those losses. Loss burst lengths are important to video applications since players can typically mask losses of a few bytes. However, losses that last for a long duration can result in perceivable video quality degradation such as a frozen frame [11].

We measure the wireless link loss rate using the packet traces from the sniffer below the AP and correlating the number of video packets transmitted with the number ACKed by the AP. The losses on the broadband link are computed by correlating the packets transmitted by the AP with the ones received at the



(a) Packet latency.



(b) Broadband queuing delay.

Fig. 4. Packet delays and inter-packet delay variation (jitter). ‘Apartments’ and ‘Houses’ refer to measurements on the wireless hop and ‘DSL’ and ‘Cable’ refer to the wired-side measurements. The queuing delay is measured on the broadband hop with four representative houses shown here.




Fig. 5. Packet arrival delay variation.

end destination. Figure 6(a) shows the loss rate on the wireless and broadband links. As can be seen from the figure, the loss rates are low, in the order of 1-2%. The low wireless loss rate indicate that the autorate selection algorithms are able to select correct bitrates.

The number of packets consecutively lost, which we call the *loss length*, is an important factor in video transmission. Video encoding schemes are built to tolerate losses, and work well when the losses are random. Long bursts of losses cannot be masked by clever encoding schemes and will result in frame freezing.

We plot a CDF of loss lengths in wired and wireless links between the source and the destination in Figure 6(b). The broadband networks show low overall loss length. On average, two packets are lost in succession. The wireless network shows a more bursty loss length, with loss lengths in the order of tens of packets on average. This variation in loss length behavior affects the choice of error recovery techniques for streaming video in residential networks over both the wireless and broadband links.

Figure 7 shows the percentage of packets that are lost and delayed on the wired and wireless networks. We count a packet as lost if the packet was not received at the destination, and as delayed if it reached the destination too late for the packet to be useful (>200ms in our case). While wireless networks have high loss rates, the delays on broadband networks increase as the load increases.

F. Discussion

As the popularity of multimedia applications and the amount of video traffic generated continues to increase, the measurement study shown here points out critical constraints that exists in the present day networks, i.e the capacity and latency of the broadband links. The upload bandwidths available are in the order of 1 Mbps, which is much below what is required for a HD video of 720p or 1080p resolutions. The bandwidth available can also vary, specifically in cable networks where the medium is shared and this will pose a problem for video encoders that scale the encoding rate based on available bandwidth [12], [13]. The latency and jitter measurements imply that players need to include a large playout buffer, which affects the real-time traffic. The queuing delays that are seen during high load conditions can introduce delays in the order of several hundreds of milliseconds up to a second, which can severely degrade an interactive video quality.

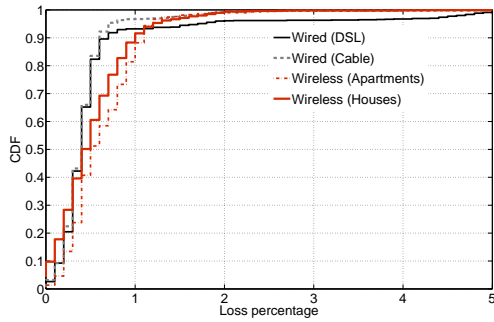
V. RELATED WORK

There have been several studies characterizing video quality streaming over WLANs and Internet. We discuss the most relevant pieces of work here.

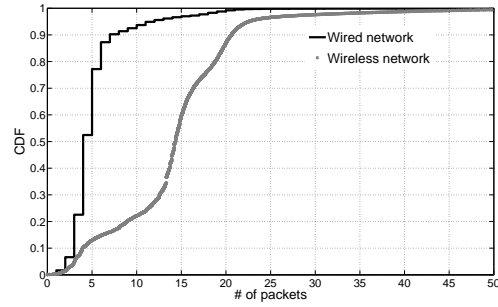
Loguinov *et al.* study the performance of low bitrate Internet video streaming [3]. Their experiments consisted of streaming MPEG-4 videos from homes using dial-up connections. The QCIF videos (176x144) used in the streaming experiments were of bitrates 14 Kbps and 54 Kbps. Using extensive experiments from 600 homes lasting for 7 months, this work characterizes the packet loss rates, loss lengths and delays experienced on dial-up connections. Video streaming performance in the context of content distribution and peer-to-peer systems have also been studied [14], [15].

Residential broadband networks have received attention in the recent years. Dischinger *et al.* study the characteristics of residential cable and DSL networks [4]. The study involved experiments using broadband hosts in North America and Europe to measure the allocated upstream and downstream bandwidths and packet latencies using TCP and ICMP probes. Our work expands these measurements in the context of video streaming quality measurements.

The work by Majumdar *et al.* [2] represents a class of work [16] wherein packet loss in wireless networks are modeled



(a) Loss rates



(b) Loss lengths

Fig. 6. CDF of Packet loss and delay from streaming 30 second videos in eight residences. There were 1000 streaming sessions in total.

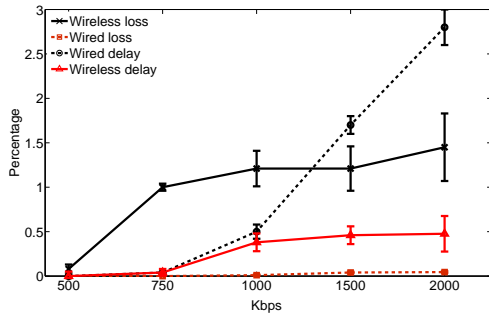


Fig. 7. Percentage of video packets lost or delayed.

and forward error correction (FEC) techniques and Automatic Repeat reQuest (ARQ) algorithms are proposed to improve multicast and unicast real-time video quality streaming. It is important to characterize the loss rates and delays in residential wireless and broadband networks so that suitable error correction techniques can be designed.

VI. CONCLUSION

We present the methodology and analyze the results of measurement study of real-time video streaming experiments. First, we answer the question as to what is the expected video streaming quality over residential links. We show that video streaming quality can range from ‘poor’ to ‘good’, and then examine the factors that contribute to video quality deterioration. We study the properties of wireless and end-to-end links in residential networks in terms of the bandwidth available for streaming, loss and latency that packets experience and the effect on streaming quality. We find the uplink bandwidth in broadband networks is typically insufficient to stream HD video streams. Further, the high latency that can be experienced on these networks can make real-time communication infeasible. The measurements presented in this work can serve as a guide on what video resolutions will be supported, and the buffer sizes needed for residential real-time video applications.

REFERENCES

[1] “Cisco Visual Networking Index: Forecast and methodology 2008-2013,” http://www.cisco.com/en/US/solutions/collateral/ns341/ns525/ns537/ns705/ns827/white_paper_c11-481360.pdf.

[2] Abhik Majumdar and Daniel Grobe Sachs and Igor V. Kozintsev and Kannan Ramchandran, “Enhanced Ethernet for Data Center: Reliable, Channelized and Robust,” in *15th IEEE Workshop on Local and Metropolitan Area Networks*, Jun. 2007.

[3] D. Loguinov and H. Radha, “Measurement Study of Low-bitrate Internet Video Streaming,” in *Proceedings of the 1st ACM SIGCOMM Workshop on Internet Measurement*, 2001.

[4] M. Dischinger, A. Haeberlen, K. P. Gummadi, and S. Saroiu, “Characterizing residential broadband networks,” in *Proc. ACM SIGCOMM Internet Measurement Conference*, San Diego, CA, USA, Oct. 2007.

[5] M. Ihmig and P. Steenkiste, “Distributed Dynamic Channel Selection in Chaotic Wireless Networks,” in *13th European Wireless Conference*, Paris, France, Apr. 2007.

[6] A. Akella, G. Judd, S. Seshan, and P. Steenkiste, “Self-management in chaotic wireless deployments,” in *Proc. ACM Mobicom*, Cologne, Germany, Sep. 2005.

[7] D. Han, A. Agarwala, D. G. Andersen, M. Kaminsky, K. Papagiannaki, and S. Seshan, “Mark-and-Sweep: Getting the ‘inside’ scoop on neighborhood networks,” in *Proc. Internet Measurement Conference*, Vouliagmeni, Greece, Oct. 2008.

[8] V. Bychovsky, B. Hull, A. K. Miu, H. Balakrishnan, and S. Madden, “A measurement study of vehicular internet access using in situ Wi-Fi networks,” in *Proc. ACM Mobicom*, Los Angeles, CA, Sep. 2006.

[9] “Audio/Video Streaming over 802.11,” www.ieee802.org/802_tutorials/..../video%20over%20802%2011%20Tutorial-final.ppt.

[10] M. Dischinger, A. Mislove, A. Haeberlen, and K. P. Gummadi, “Detecting bittorrent blocking,” in *Proc. Internet Measurement Conference*, Vouliagmeni, Greece, Oct. 2008.

[11] J. Boyce and R. Gaglianello, “Packet loss effects on MPEG video sent over the public internet,” in *ACM Multimedia*, 1998.

[12] “Adaptive Video Streaming Home Page,” <http://nms.lcs.mit.edu/projects/videoem/>, 2001.

[13] N. Feamster, “Adaptive delivery of real-time streaming video,” Master’s thesis, Massachusetts Institute of Technology, May 2001, Winner of the MIT EECS William A. Martin Memorial Thesis Award.

[14] H. Yu, D. Zheng, B. Y. Zhao, and W. Zheng, “Understanding user behavior in large-scale video-on-demand systems,” in *In Proc. EuroSys*, 2006.

[15] H. Yin, X. Liu, T. Zhan, V. Sekar, F. Qiu, C. Lin, H. Zhang, and B. Li, “Design and deployment of a hybrid CDN-P2P system for live video streaming: experiences with LiveSky,” in *Proc. ACM Multimedia*, 2009.

[16] B. W. Wah, X. Su, and D. Lin, “A survey of error-concealment schemes for real-time audio and video transmissions over the Internet,” in *Proc. Int’l Symposium on Multimedia Software Engineering*. Taipei, Taiwan: IEEE, Dec. 2000, pp. 17–24.

[17] Vouliagmeni, Greece, Oct. 2008.