MEDIA STREAMING PERFORMANCE IN A PORTABLE WIRELESS CLASSROOM NETWORK

Xiaozhen Cao   Guangwei Bai   Carey Williamson
Department of Computer Science
University of Calgary
Calgary, AB, Canada
Email: {caox,bai,carey}@cpsc.ucalgary.ca

ABSTRACT

This paper presents empirical measurements of wireless media streaming traffic in an IEEE 802.11b wireless ad hoc network. The results show that the IEEE 802.11b WLAN can support up to 8 clients with good media streaming quality, with each client receiving a separate 400 kbps video stream and 128 kbps audio stream. With 9 clients, the WLAN is overloaded, and performance degrades for all clients. Finally, we demonstrate a "bad apple" phenomenon in wireless ad hoc networks, wherein a single client with poor wireless connectivity disrupts the media streaming quality for all clients sharing the WLAN.

KEY WORDS

Mobile Multimedia, Wireless Networks, Quality of Service

1 Introduction

The popularity of wireless LANs and the emergence of media streaming applications on the Internet jointly enable wireless multimedia streaming [1, 2, 3, 4, 5, 6]. The "wireless Web" is now part of our daily lives, in the classroom, the office, and the home [7, 8, 9, 10, 11]. Typical streaming applications include seminars, media events, sports, and entertainment applications. Educators can embrace wireless technologies to provide on-line access to lecture notes, demos, and other supplementary material in the classroom.

In this paper, we study wireless multimedia streaming performance in an IEEE 802.11b classroom area network. We use classroom measurements and laboratory tests to determine experimentally the achievable streaming performance for clients on an IEEE 802.11b WLAN. We also study the performance capabilities of a Darwin Streaming Server [12] running on a laptop computer with an IEEE 802.11b wireless interface. All laptops are configured in *ad hoc* mode. The clients access media content from the wireless streaming server. A wireless network analyzer is used to collect packet traces from the wireless channel.

Our experiments focus on issues such as number of wireless clients, wireless connectivity, protocol efficiency, and performance under overload. The results show that the IEEE 802.11b WLAN can support up to 8 clients, each receiving individual audio (128 kbps) and video (400 kbps) streams. The aggregate network load is approximately 4.6 Mbps. With 9 clients, the WLAN saturates, degrading performance for all clients. Finally, we demonstrate a "bad apple" phenomenon in wireless ad hoc networks: a single client with poor WLAN connectivity can disrupt the media streaming quality for *all* clients in the WLAN.

The remainder of this paper is organized as follows. Section 2 provides background information on IEEE 802.11b wireless LANs and multimedia streaming. Section 3 describes the experimental setup for our study. Section 4 and Section 5 present the measurement results. Finally, Section 6 concludes the paper.

2 Background

2.1 IEEE 802.11b Wireless LAN

The IEEE 802.11b WLAN standard [13] is a popular technology in the wireless LAN market today. "WiFi" (Wireless Fidelity) provides low-cost wireless Internet capability for end users, with data transmission rates of up to 11 Mbps at the physical layer.

The IEEE 802.11 standard allows two types of WLAN configurations. In *infrastructure mode*, all mobile stations in the WLAN communicate via an Access Point (AP) connected to the external Internet. In *ad hoc mode*, all the stations in the WLAN communicate directly with each other, without requiring an AP. Frames are addressed directly from sender to receiver using the MAC addresses in the frame header.

The IEEE 802.11b standard defines the channel access protocol used at the MAC layer, namely Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). It also defines the frame formats used at the data link layer: 128-bit preamble, 16-bit Start-of-Frame delimiter, 48-bit PLCP (Physical Layer Convergence Protocol) header, followed by a 24-byte LLC header, and a variable size payload, often used for carrying IP packets. Frames that are correctly received over the shared wireless channel are acknowledged by the receiver. Unacknowledged frames are retransmitted by the sender after a short timeout (a few milliseconds), using the same MAC protocol.

Our paper studies media streaming performance in ad hoc IEEE 802.11b networks. IEEE 802.11b solutions are widely available today, from many vendors. Many

802.11b-based networks have been installed in businesses and public areas. Price points for this technology are rapidly declining. However, the MAC mechanisms supported by IEEE 802.11b, namely Distributed Coordination Function (DCF) and Point Coordination Function (PCF), provide limited support for multimedia. In particular, they do not support multiple concurrent media streams well.

Currently, the IEEE 802.11e group is developing MAC improvements to support QoS sensitive applications. The IEEE 802.11e is under design and in the standardization process. A draft specification is available. It introduces two additional MAC modes: the Enhanced Distributed Coordination Function (EDCF) and the Hybrid Coordination Function (HCF), to enable a better mobile user experience and to make more efficient use of the wireless channel. Product availability will follow once 802.11e is finalized.

2.2 Multimedia Streaming

Streaming technology delivers multimedia content over a network from a server to a client in real time. The media is not downloaded to a viewer's hard drive. Rather, the media is played as the client receives it (except possibly for a short buffering delay). If the client wishes to play the media again, the streaming process is repeated.

An end-to-end streaming system requires a streaming media server and a client media player. Media clips can be created with production tools to convert audio, video, or animation to a digital format such as MPEG-4. Streaming servers such as the Darwin Streaming Server (Apple) or RealServer (RealNetworks) can deliver media clips to clients running MP4Player, RealPlayer, or QuickTime.

The main networking protocols used for multimedia streaming are Real-Time Streaming Protocol (RTSP) [14], Real-Time Control Protocol (RTCP) [15], and Real-Time Protocol (RTP) [15]. RTSP is a signalling protocol used to establish and manage a client/server streaming connection, including session initiation and media negotiation. The RTSP connection lasts throughout the media streaming session, in case the client wishes to pause, stop, rewind, or replay the media stream. RTP and RTCP are the protocols used to transmit and control the actual media data. RTP is a commonly-used protocol for real-time multimedia transport over IP networks. RTCP is an adaptive feedback control protocol for RTP.

The media streaming system used in this paper is the Darwin Streaming Server [12]. A media streaming session has three distinct phases:

- *Initialization*. The client requests a selected media file from the server, using RTSP over TCP. The server returns media format information to the client, which then issues a `setup` request to specify the protocols and ports for transmission. The server replies with the selected protocol, acknowledging the client's port numbers, and indicating port numbers for feedback sent by the client. Next, the client issues the `play` request. The server responds with "OK", plus synchronization information for the upcoming RTCP and RTP flows.

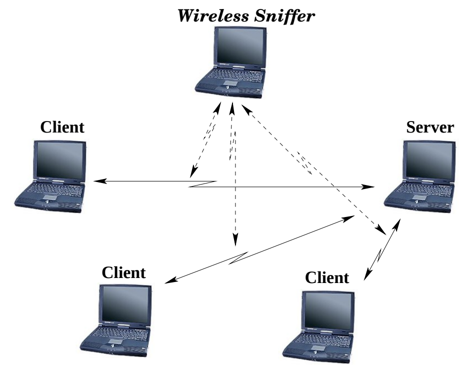

Figure 1. Experimental Setup for Streaming

- *Media Transmission*. RTCP and RTP work together, both running over UDP. The RTP protocol packetizes and sends the media data to the client (e.g., video to port 1024, and audio to port 1026). The RTCP info is sent to ports 1025 and 1027.

- *Session Termination*. When the streaming is finished, the server initiates a four-way TCP handshake to close the RTSP/TCP connection.

3 Experimental Methodology

3.1 Experimental Setup

In our work, we use an *ad hoc* WLAN as shown in Figure 1. The simple testbed consists of several wireless clients and a streaming server. In addition, we use a wireless network analyzer to monitor the wireless channel. Each laptop has a Cisco Aironet 350 Series Adapter for access to the IEEE 802.11b WLAN. The wireless cards operate in *ad hoc* mode.

The wireless clients run MP4Player [16] to access and play media content from the Darwin Streaming Server [12]. The media streaming session uses RTSP over TCP, RTP over UDP, RTCP over UDP, as well as the IEEE 802.11b MAC protocols. We run a specially instrumented Linux kernel on the streaming server to record packet arrivals, packet departures, and packet queueing events at the wireless network interface.

WLAN traffic measurements are collected using a wireless network analyzer. Its wireless network card operates in promiscuous mode, recording all activity on the wireless LAN (i.e., frame transmissions, MAC-layer retransmissions). Decoding of the captured traces enables protocol analysis at the MAC, IP, and TCP/UDP layers.

The MP4Player application provides summary information about video and audio playback rates. This information complements our network traffic measurements.

3.2 Media Traffic Characterization

The primary video clip used in our experiments is an 8-minute clip from the movie *Au Revoir Les Enfants* (1988). We digitized this clip, converting it from its VHS version to an MPEG-4 format. The specified media rates for the compression were 400 kbps for video and 128 kbps for audio. Analysis of the resulting clip showed an average video rate of 394 kbps (30 fps), with 128 kbps audio (43 fps).

Figure 2 shows the video frame size distribution for this media clip. Figure 2(a) shows the overall distribution for the 14,854 frames. The mean frame size is 1,641 bytes. About 90% of the frame sizes are between 1000 and 2000 bytes, while 1.4% of the frames are larger than 4 KB. The distribution for the 214 large frame sizes appears in Figure 2(b). These frames primarily represent I-frames in the MPEG-4 encoding, appearing periodically in the video trace every 90 frames. However, this structure varies occasionally in the trace, perhaps due to scene changes and different camera angles.

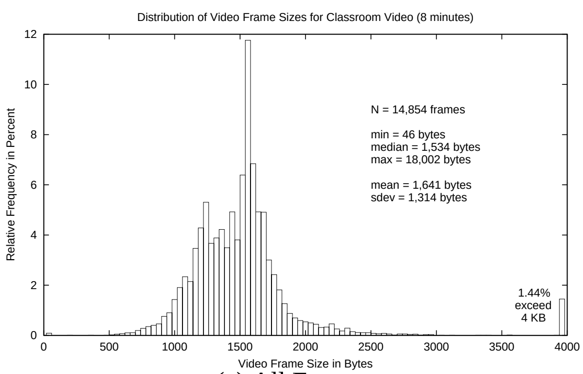

(a) All Frames

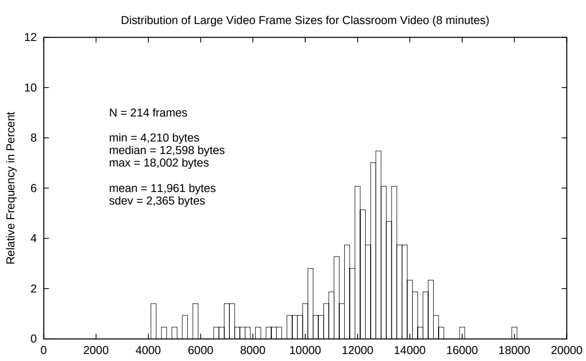

(b) I-Frames

Figure 2. Video Frame Size Distribution

At coarse-grain time scales (e.g., 30 seconds or more), the video is Constant Bit Rate (CBR), consistent with its compression specification. However, the video is Variable Bit Rate (VBR) at finer-grain time scales (e.g., 10 seconds or less). The VBR profile affects the performance of wireless media streaming.

4 Classroom Measurements

We conducted a live classroom test of our wireless media streaming technology in March 2004. With the cooperation of the French Department at the University of Calgary, we used a wireless ad hoc network to provide media streaming to students in a French cinematography class. There were 15 students in the class. We provided 8 client laptops for them to share while accessing content in the classroom.

The classroom experiment lasted 20 minutes. Three media clips were provided to the students: an 8-minute clip, a 30-second clip, and a 20-second clip. Students were instructed to view the 8-minute clip, and then complete a multiple-choice quiz about the film. Clients use a Web browser to download the quiz from the wireless Web server (the same laptop running the Darwin Streaming Server) in the classroom, answering the questions, and submitting the completed quiz to the Web server. Students could view the shorter clips when answering specific quiz questions.

The experiment was conducted successfully. All 8 clients viewed the film clip and completed the quiz. Student feedback regarding the use of the wireless streaming technology in the classroom was very positive.

Table 1 summarizes the network traffic measurement results from the live classroom experiment. During the 20 minute test, there were 469,778 packets sent successfully across the WLAN. Over 99% of these were UDP packets, for the media streaming. The TCP packets were initiated by RTSP streaming sessions, and by the HTTP transactions with the Web server. Wireless channel errors were negligible: 0.5% of the packets were MAC-layer retransmissions.

Table 1. Network Traffic Summary

Item	Value
Trace Duration	20 minutes
Total Packets	469,778
UDP Packets (99.3%)	466,331
TCP Packets (0.7%)	3,447
MAC-layer Retransmissions	2,041
CRC Errors	118

Figure 3 illustrates the WLAN usage during the classroom experiment. The UDP media streaming in Figure 3(a) dominates the activity. The small TCP traffic spikes in the first 200 seconds of Figure 3(b) show when each client started its RTSP streaming session. The corresponding jumps in UDP traffic are evident in Figure 3(a). When the 8-minute media clip completes, the clients initiate HTTP activity to the Web server. Several clients activate UDP media streaming while completing the quiz.

In the first 10 minutes, when all 8 clients were streaming the 8-minute clip, the aggregate transmission rate was relatively stable around 4.6 Mbps. In the second half of the trace, the transmission rate varied due to the mixed HTTP and RTSP requests, and non-deterministic user behavior.

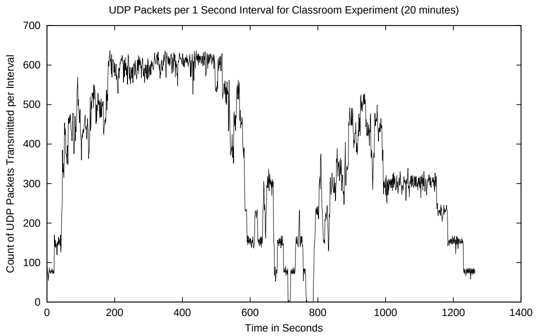

(a) UDP Traffic

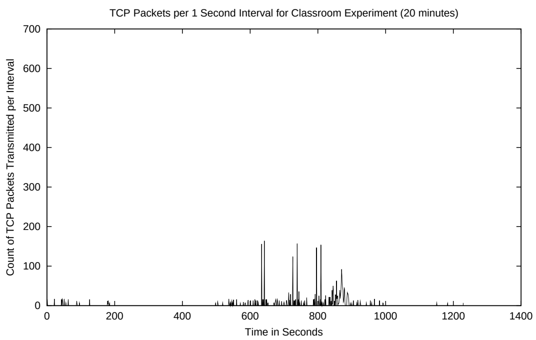

(b) TCP Traffic

Figure 3. Classroom WLAN Usage

The multimedia streaming clearly places the heaviest demand on the WLAN. Prior work shows that the achievable throughput on an IEEE 802.11b WLAN is typically 5-6 Mbps [7, 17]. Our achieved throughput is slightly lower. One reason is that neither UDP nor TCP use full-sized packets on the WLAN. For example, the UDP traffic is dominated by 1200-byte payloads. Larger packet sizes could improve protocol efficiency, by amortizing the WLAN channel access overhead.

These results show that a portable media server with an ad hoc IEEE 802.11b WLAN can support 8 clients with adequate media streaming quality. Aggregate network usage of 4.6 Mbps is achievable.

5 Additional Experiments

To better understand wireless media streaming behaviour, we conducted several additional experiments in our lab, under controlled conditions. These experiments used additional system instrumentation to provide more information about WLAN performance.

The additional experiments focused on multimedia content streaming only (i.e., no Web/HTTP transactions). We used the same system setup as described in Section 3. In the first experiment, we study how the number of clients affects the server's media streaming performance. In the second experiment, we explore the "bad apple" phenomenon.

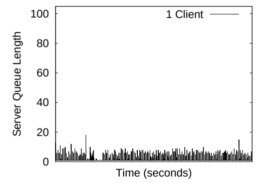

(a) 1 Client

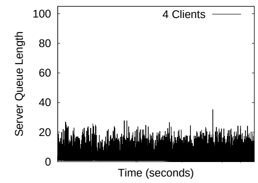

(b) 4 Clients

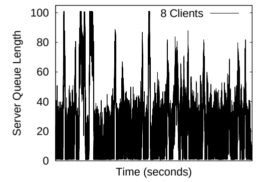

(c) 8 Clients

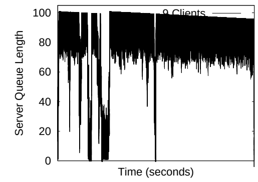

(d) 9 Clients

Figure 4. Server Link-Layer Queue Behaviour

5.1 Number of Clients

Earlier results indicated that the WLAN can adequately support 8 clients. Here we revisit this issue for 1, 4, 8, and 9 clients, studying the server-side WLAN bottleneck.

In the Linux kernel, a packet to be sent first enters a shared FIFO queue with other packets (if any) waiting for network transmission. We are interested in the queueing behaviour at the WLAN interface.

Figure 4 shows the server-side link-layer queue dynamics for 1, 4, 8, and 9 clients. Figure 4(a) shows the results for a single client. There is low delay at the server's WLAN interface queue, with typically 1-10 packets pending. With 4 clients (Figure 4(b)), the occupancy of the server's link-layer queue increases. The default queue capacity in Linux is 100 packets. With 8 clients (Figure 4(c)), the queue occasionally reaches this limit, though the system still runs well.

The system becomes unstable with 9 clients. The WLAN is the bottleneck. The queue increases dramatically (see Figure 4(d)), and packet losses occur. The large queue increases the delay for all packets, increasing the risk of late-arriving packets being discarded at the player. This problem affects all the clients in the system since they share the same server queue. As a result, 9 clients experience poor playback performance.

Table 2 summarizes the playback performance in terms of video rate, audio rate, displayed frames, skipped frames, and average lateness. Ideally, the remote video and audio playback should match the local playback rates. Missing frames are an indication of quality degradation, as are skipped frames and increasing lateness.

The results in Table 2 show that remote playback works well for up to 8 clients, but degrades for 9 clients. Lateness increases, and more video frames are skipped. The effective media playback rate decreases.

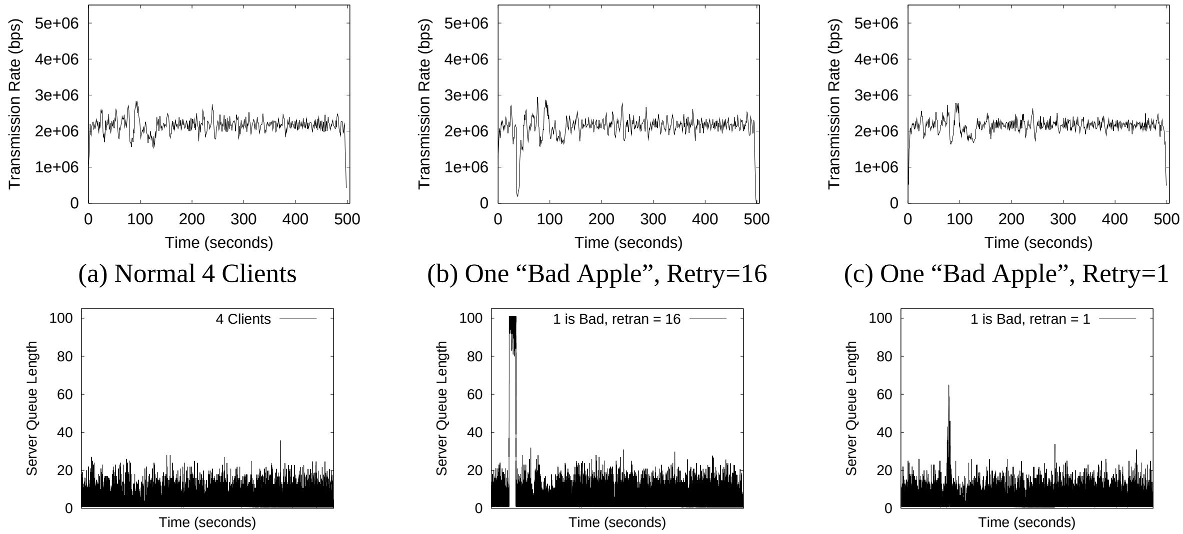

Figure 5. Network Usage (top) and Queue Behaviour (bottom) showing the "Bad Apple" Phenomenon

Table 2. Media Player Performance with Different Numbers of Clients

Number of Clients	1	4	8	9
Video rate (fps)	29.96	29.96	29.87	27.36
Audio rate (fps)	43.07	43.09	43.03	39.76
Displayed frames	14,854	14,854	14,809	13,567
Skipped video frames	0	0	4	5
Avg lateness (sec)	1	2.5	3.75	24.67

5.2 The "Bad Apple" Phenomenon

Bai *et al.* [1] reported a "bad apple" phenomenon for wireless media streaming. When one mobile client has poor wireless connectivity in the WLAN, this client can disrupt media streaming quality for *all* of the other clients in the WLAN. In other words, the "bad apple spoils the batch". The disruption persists until the "bad apple" rejoins the WLAN with adequate network connectivity, at which point the media streaming sessions (if they survive) resume.

In this section, we study the "bad apple" phenomenon more closely, to identify its root cause, and explore possible solutions. We conduct this experiment in our laboratory with 4 clients, placing only moderate load on the WLAN. During the experiment, we dislodge the wireless network card for one of the clients, and observe the results. We then quickly replace the wireless network card, and continue the streaming experiment.

Figure 5 shows the results from this experiment. The top row of graphs plots the aggregate WLAN usage during the experiment, while the bottom row of graphs plots the server's link-layer queue occupancy. The leftmost column of graphs is for the normal 4-client case. The middle column of graphs is for the 4-client case with one "bad apple". The rightmost column of graphs considers the "bad apple" case with one of our proposed solutions.

Figure 5(a) shows the normal case. The network load is approximately 2 Mbps, and the server's queue fluctuates between 0 and 20 packets.

Figure 5(b) shows the "bad apple" scenario. The network load is approximately 2 Mbps most of the time, but there is a sharp decline in effective utilization during the network outage for the disconnected client. The outage manifests itself in the server's queue, which rapidly fills at the time of the network anomaly, causing some packet losses. The queue stabilizes when the WLAN connectivity for the "bad apple" is restored a few seconds later.

The "bad apple" phenomenon is explained as follows. The FIFO queue at the server's WLAN interface contains an arbitrary interleaving of server-generated packets for different clients. At some point in time, the front packet in the queue is destined to the disconnected client. The IEEE 802.11b protocol tries in vain to send this packet to the client, retransmitting repeatedly, with a random delay between each attempt. No MAC-layer ACK is received, and retransmissions continue until the maximum retry limit (default 16) is reached. At this point, the packet for the "bad apple" client is discarded, and the server's WLAN interface tackles the next waiting packet in the FIFO queue.

In essence, this problem is a transient manifestation of Head Of Line (HOL) blocking, at the MAC layer. We suspect that the same phenomenon would occur in IEEE 802.11e. All pending packets in the queue are blocked while the front packet undergoes retransmissions. The effective service rate of the queue diminishes. Since the media server continues to generate packets for the streaming clients, the queue fills and overflows. When the "bad apple" reconnects, the service rate of the WLAN queue returns to normal, and the backlog dissipates.

Table 3. Statistical Summary of “Bad Apple” Phenomenon

Scenario	Normal 4 Clients	1 Bad, Retry=16	1 Bad, Retry=1
Video rate (fps)	29.96	29.69	29.77
Audio rate (fps)	43.09	41.84	42.90
Avg skipped video frames (per client statistics)	0 (0/0/0/0)	36 (37/34/27/46)	2 (0/0/0/7)
Avg displayed frames (per client statistics)	14,854 (14,854/14,854) (14,854/14,852)	14,524 (14,545/14,534) (14,560/14,455)	14,774 (14,774/14,777) (14,771/14,125)
Avg lateness (sec)	2.5	14.5	5.75

Several possible fixes for this problem are possible. One approach would be to use multiple queues at the wireless network interface, with one queue for each client. However, the queues alone are not sufficient. A scheduling discipline is also needed to arbitrate amongst the queues, giving precedence to well-connected clients, and minimal service to the “bad apple”. A second approach would be to use separate transmission channels (frequencies, in Hz) for each client, but this has hardware implications on IEEE 802.11b transmitters and receivers. A third approach, and one that we propose as an interim solution, is to limit the number of MAC-layer retransmissions.

To test the latter solution, we configured the server’s wireless network card to allow at most one MAC-layer retransmission of each frame. The results from this configuration are shown in Figure 5(c). Only a small queue response is seen at the server at the time of the network outage, and the system recovers quickly.

Table 3 provides a statistical summary of the “bad apple” phenomenon, and the effectiveness of our solution. This zero-cost solution solves the “bad apple” problem for wireless media streaming, but at the risk of introducing unreliable wireless delivery for *both* TCP and UDP packets. Our future work will investigate a more satisfactory solution for this problem.

6 Conclusions

This paper presents an empirical study of wireless media streaming performance in an IEEE 802.11b wireless ad hoc network. Experiments are conducted in a classroom environment with students, as well as in a research lab under controlled conditions.

The experimental results show that an IEEE 802.11b WLAN can support up to 8 clients with good media streaming quality. In our experiment, each client receives a 400 kbps unicast video stream and a 128 kbps unicast audio stream, producing an aggregate network load of approximately 4.6 Mbps. With 9 clients, the WLAN is overloaded, and performance degrades for all clients. We also demonstrate the “bad apple” phenomenon: a single client with intermittent wireless connectivity can disrupt media streaming quality for all clients sharing the WLAN.

The “bad apple” phenomenon can seriously degrade multimedia delivery in wireless environments. We presented a simple solution to solve the “bad apple” problem by limiting MAC-layer retransmissions. Finding a better solution for the “bad apple” phenomenon remains as future work. We are also investigating multicast protocols for efficient wireless multimedia streaming.

Acknowledgements

Financial support for this research was provided by NSERC (Natural Sciences and Engineering Research Council), iCORE (Informatics Circle of Research Excellence), CFI (Canada Foundation for Innovation), and TRLabs (Telecommunications Research Laboratories).

The authors are grateful to Professor Daniel Maher and Professor Brian Gill for allowing us to use their French 217 students as test subjects for our experiments. Grant McGibney and Simon Arseneault from TRLabs helped digitize the media clips used in this study, and Nayden Markatchev helped coordinate the experiments.

The authors thank the anonymous IMSA reviewers for their comments on an earlier version of this paper. The media clips and quiz used in this study are available from `http://www.cpsc.ucalgary.ca/~caox/`

References

[1] G. Bai and C. Williamson, “The effects of mobility on wireless media streaming performance”, *Proc. IASTED Intl. Conf. on Wireless Networks and Emerging Technologies (WNET)*, Banff, Canada, 2004, 596-601.

[2] T. Kuang and C. Williamson, “Hierarchical analysis of RealMedia streaming traffic on an IEEE 802.11b wireless LAN”, *Computer Communications, 27*, 2004, 538-548.

[3] A. Majumdar, D. Sachs, I. Kozintsev, K. Ramchandran and M. Yeung, “Multicast and unicast realtime video streaming over wireless LANs”, *IEEE Trans. Circuits and Systems for Video Technology, 12*(6), 2002, 524-534.

[4] A. Muir and J. Garcia-Luna-Aceves, “Supporting real-time multimedia traffic in a wireless LAN”, *Proc. ACM/SPIE Multimedia Computing and Networking Conf. (MMNC)*, San Jose, USA, 1997.

[5] W. Pattara-atikom, P. Krishnamurthy, and S. Banerjee, “Distributed mechanisms for quality of service in wireless LANs”, *IEEE Trans. Wireless Communications*, *10*(3), 2003, 26-34.

[6] K. Xu, Q. Wang, and H. Hassanein, “Performance analysis of differentiated QoS supported by IEEE 802.11e enhanced distributed coordination function (EDCF) in WLANs”, *Proc. IEEE GLOBECOM*, San Francisco, USA, 2003.

[7] G. Bai, K. Oladosu, and C. Williamson, “Performance issues for wireless web servers”, *Proc. Intl. Workshop on Mobile and Wireless Ad Hoc Networking (MWAN)*, Las Vegas, USA, 2004, 59-65.

[8] G. Bai and C. Williamson, “Simulation evaluation of wireless web performance in an IEEE 802.11b classroom area network”, *Proc. 3rd IEEE Intl. Workshop on Wireless Local Networks (WLN)*, Bonn, Germany, 2003, 663-672.

[9] D. Kotz and K. Essien, “Analysis of a campus-wide wireless network”, *Proc. ACM MOBICOM*, Atlanta, USA, September 2002, 107-118.

[10] D. Tang and M. Baker, “Analysis of a local-area wireless network”, *Proc. ACM MOBICOM*, Boston, USA, 2000, 1-10.

[11] C. Williamson and N. Kamaludden, “Network traffic measurements of a wireless classroom network”, *Proc. 16th Intl. Conf. on Wireless Communications*, Calgary, Canada, 2004, 561-570.

[12] Apple Open Source Projects, *Darwin Streaming Server*, `http://developer.apple.com/darwin/`

[13] IEEE, IEEE 802.11, “Wireless LAN medium access control (MAC) and physical layer (PHY) specifications”, 1999.

[14] RFC 2326, *Real Time Streaming Protocol*, `http://www.faqs.org/rfcs/rfc2326.html`

[15] RFC 1889, *RTP, A Transport Protocol for Real-Time Applications*, `http://www.faqs.org/rfcs/rfc1889.html`

[16] MPEG, *MPEG4IP: Open Source, Open Standards, Open Streaming*, `http://www.mpeg4ip.net/`

[17] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, “Performance anomaly of 802.11b”, *Proc. IEEE INFOCOM*, San Francisco, USA, 2003, 836-843.