On Video Multicast over Wireless LANs

Sumathi Gopal,
WINLAB – Rutgers University
sumathi@winlab.rutgers.edu

Kumar Ramaswamy
Thomson Inc.
kumar.ramaswamy@thomson.net

Charles Wang
Thomson Inc.
charles.wang@thomson.net

Abstract

In this paper, we explore and analyze the feasibility of current generation 802.11b wireless networks to sustain video multicast applications. We assess bandwidth efficiency to transport packets of various sizes over 802.11b WLAN in Distributed Coordination Function (DCF) mode and derive corresponding theoretical maximum data rates. Due to the overhead associated with the different protocol layers, the efficiency of transport can vary widely with different packet sizes. By controlled WLAN experimentation the effect of this phenomenon on video traffic is studied. Application to MPEG4-AVC streams show that the packet overhead in 802.11 DCF mode causes significant bottleneck for simultaneous multicast streams. We conclude with several suggestions to alleviate this problem including operating in 802.11 PCF mode, using a statistical multiplexer and improving design of queues in the wireless access point.

Introduction

There is a proliferation of wireless local-area networks (WLANs) for data access. Reliable video distribution over WLANs is still a challenging problem to solve. Transport of video content is very sensitive to bandwidth variability, delay jitter, packet error rates and types of packets in which errors occur. Our studies indicate that packet errors in the WLAN end-to-end delivery system are affected not only by channel fades, but also by limited packet rates supported by the wireless channel, access point queues and by the nature of content itself. There is more susceptibility for video broadcast over wireless LAN, as there is no inherent feedback mechanism in the network.

Data traffic comprising of file downloads etc, uses fixed packet sizes and TCP - a reliable transport protocol that implements receiver-based as well as network-based flow control. With multimedia traffic, there are a few possible paradigms of transport. In a unicast download model, TCP/IP may still be used. Here buffering and consequent receiver-end delays are issues that need addressing. For multicast/broadcast applications, a common choice is the use of the RTP/UDP/IP stack[5]. The UDP/IP layers provide a best effort transmission service, with no guarantees of reliability or flow control. Further with retransmissions shut down in the link layer, there is minimum end-to-end delay but at the expense of reduced link reliability. Large video frames are fragmented at the RTP layer and form payloads of multiple packets. The RTP header provides a display timestamp and sequence number.

Several video compression standards have been developed to cater to video transmission. The most popular ones include H.263 and MPEG-2 that are tailored for different applications. Emerging standards include MPEG4-SP/ASP and MPEG4-AVC (also called H.264 or JVT). It is important to evaluate the efficacy of using efficiently compressed video over a lossy channel while employing forward error correction techniques (which implies an additional overhead) as opposed to using a hierarchical coding technique such as MPEG4-SP/ASP that is less coding efficient. In this paper, we focus in general on the carriage of an efficient coding format over WLANs.

Next we introduce some common characteristics of video compression. In particular, we will discuss those of MPEG4-AVC. There are several types of frames (Independent Decoding Frame- IDR, P and B) whose size is very content dependent and is influenced by a number of factors including resolution and other selected encoding parameters. We analyzed a representative set of MPEG4-AVC encoded video streams provided by [1]. All these video streams were QCIF at 30 *frames/sec* with a global quantization parameter 15. Some of their characteristics include:

- Wide variation in size of the different frame types (IDR frames may be up to 20 Kbytes while a B frame may be as small as a few tens of bytes)
- IDR frames may constitute almost half the compressed bit rate even though they constitute only 1/12 the number of frames

In the next section, we present details of WLAN experiments. Then we derive theoretical throughput limits for various packet sizes. Following that, we apply throughput limits to real video streams and present the main results of this paper. We conclude with several suggestions to alleviate the inherent problem of 802.11b bottleneck for video multicast.

WLAN Experiments

The main goal here was to understand how key parameters such as packet size, channel load and receiver location affect packet losses for video traffic. Figure 1 depicts the system of interest. However for our study it suffices to consider the setup depicted in Figure 2 with the set of encoders replaced with a single server generating traffic based on the nature of video data. The WLAN was an IEEE 802.11b network with a popular access point (AP) and wireless notebook adaptors. The experiments were conducted in an indoor office environment on weekends and holidays (to ensure repeatability). All sessions were multicast as per 802.11 MAC specifications, i.e. with no RTS/CTS exchanges, no MAC level ACKs and use of a fixed channel modulation.

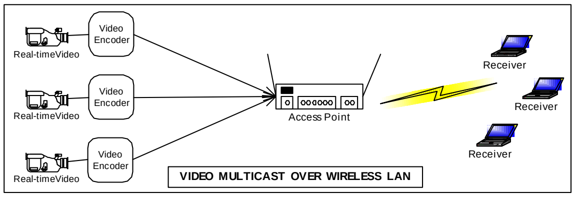

Figure 1

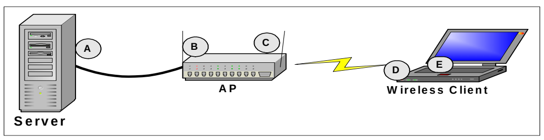

Figure 2: Wireless Network Setup

Test Parameters and Configuration

We developed the test bench and software tools required for network performance monitoring. Each session comprised of 100,000 packets, all of the same size and sent at a constant rate. The experiments were designed around the following parameters:

(a) Channel rates (ChBR): 2 and 11 Mbps (b) Three locations w.r.t. the AP- 2F(one floor below), CR (conference room 10*mts* away) and LOS (line of sight) (d) Channel load (ChLoad): Physical-level constant bit rate traffic of 25% and 50% of ChBR (e) Size of packets (application level) (PS): 16, 32, 64, 128, 256, 512, 1024 byte packets, representing video packet sizes. (f) Application level bit rate (ABW): This was derived from PS, ChBR and ChLoad using:

$$ABW = \frac{PS}{PS + Overhead} * ChLoad$$

where *Overhead* is the overhead due to UDP, IP, LLC, MAC and Physical layer headers summing to 94 bytes.

WLAN Experiment Results

Effect of packet size on packet loss

Figure 3 plots packet loss % vs. packet size for different channel loads at a line of sight location w.r.t the AP. Due to the vicinity, loss of data due to a lossy channel was not the primary contributor. At 11*Mbps* channel rate, over 95% of the packets in a session with small packets were lost. Setting probes at various points in the network as shown in Figure 2 and analyzing AP logs with standard network management tools, we found that majority of the packet losses were due to the overflow of queues in the AP. Analysis revealed this buffer overflow problem to be due to delays incurred by

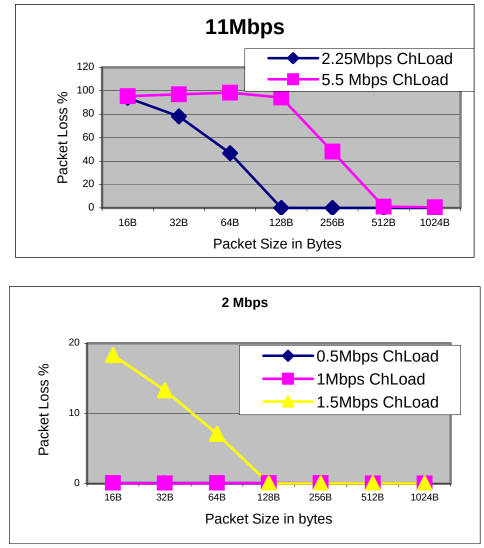

Figure 3: Packet loss at various packet sizes and channel loads

Probe Point	16B	32B	64B	128B	256B
A: Packets sent by Server	100%	100%	100%	100%	100%
D and E	2.91%	2.48%	1.98%	8.74%	53.00%
Packets dropped by AP	97.09%	97.52%	97.02%	91.26%	47%
AP Reported loss: Reading at C minus Reading at B	91.1%	92.53%	95.88%	91.20%	46.80%

Table 1. Packet Loss Profile for Varying Packet size at various probe locations shown in Figure 2

802.11b MAC and Physical layer overheads. In the DCF operating mode of 802.11b, the average backoff duration per packet and the Distributed Interframe Space (DIFS) together amount to 350μs, in the absence of collisions. The random backoff window doubles in case of a collision. Further the 802.11b physical layer protocol (PLCP) header and preamble are transmitted at 2*Mbps* channel rate (1*Mbps* for backward compatibility), instead of the previously assumed *11Mbps*, causing additional delays tabulated in Table 2.

Thus physical and MAC layer overheads in 80211b limit the outgoing packet rate at the AP. When the incoming packet rate exceeded the outgoing packet rate, packets got buffered in the AP queues and were subsequently discarded as reported in AP logs. Further analysis of this situation is presented in the next section.

Effect of channel load on packet errors

While operating within theoretical max throughput limits, there was no correlation between packet losses and the operating channel load, for any packet size.

Figure 4 shows one plot of experimental runs. Thus we may infer that the error characteristic of the wireless network does not vary with instantaneous bit rates of video traffic.

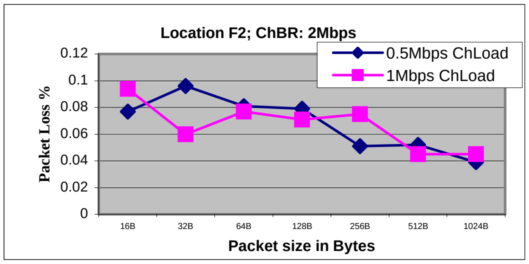

Figure 4: Losses at various channel loads

Effect of location on packet losses

Signals experience path loss, and we normally expect packet loss to increase with distance. However signal attenuation by intermediate objects could cause higher losses. The location based comparison below is for a ChLoad of 1*Mbps* while operating at 2*Mbps* ChBR. We only considered sessions operating within maximum theoretical throughput levels (see next section). Two locations were considered – 2F and CR. Neither was at line of sight. 2F was closer in distance to the AP despite a ceiling separation. CR was on the same floor as the AP, at a distance of about 10 meters, but with many intermediate objects (doors, cabinets etc). Higher packet losses were observed at CR than at 2F as shown in Figure 5.

Further, at any location, consecutive packets were seldom lost implying that for the particular controlled experimentation scenario, channel fades seldom lasted longer than the duration of a packet.

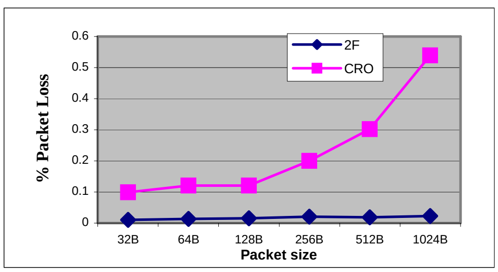

Figure 5: Loss at different locations

Maximum Application Level Throughputs

In the previous section, we analyzed AP buffer overflow as due to limited packet rates supported by 802.11b protocol in DCF mode. Jun et. al.[2] derived theoretical limits for MAC-level throughputs for various packets sizes over 802.11. We extend that analysis to derive application level bit rate limits for multicast packets (no MAC-level ACKs) sent over RTP/UDP/IP stack, for various packet sizes. Table 2 lists overheads incurred for each multicast packet transmitted at 11*Mbps* ChBR, in the absence of collisions or channel fades.

Let $T_{overhead}$:overhead duration per packet; T_p :time taken to transmit a packet PS Bytes; ε_{PS} : Bandwidth efficiency for a PS byte packet; $TMT_{app}(PS)$: Theoretical maximum application throughput for PS byte packet. The following equations hold:

$$T_{overhead} = T_{DIFS} + T_{backoff} + T_{phy+preamble} + T_{MAC} + T_{LLC+IP+UDP} + T_{RTP}\ \mu s$$

$$T_{PS} = 8 * PS * 11\ \mu s$$

$$\varepsilon_{PS} = T_{PS}/(T_{overhead} + T_{PS})$$

$$TMT_{app}(PS) = \varepsilon_{PS} * 11\ \text{Mbps}$$

Overhead	Duration (µs)
DIFS	50
Average backoff with Min congestion window	310
Physical layer: short Preamble(144bits/2Mbps) + PLCP header (48bits/2Mbps)	96
MAC header + FCS duration (8*34bytes/11Mbps)	24.73
LLC+IP+UDP headers duration (8*(8+20+8)bytes/11Mbps)	26.18
RTP header duration (8*12bytes/11Mbps)	8.73
Total overhead duration per packet $T_{overhead}$	**515.63**

Table 2. Per packet Time Overhead for 802.11 DFS

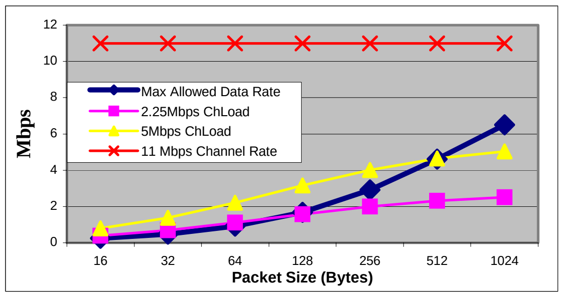

Figure 6. Maximum Application level Throughputs

Figure 6 plots $TMT_{app}(PS)$ along with actual channel loads for the ChLoad levels considered earlier. ε_{PS} ranges from a mere 2% for 16 byte packets to 60% for 1024 byte packets. For the ChLoads considered earlier, the application-level data rates in many instances exceed TMT_{app}. Comparing Figure 6 to Figure 3 it is clear that significant packet losses are incurred by exactly those sessions.

Multiple MPEG4-AVC Stream Transport over 802.11b

While most of the previous analysis focused on fixed packet sizes, compressed video in practice consists of a wide range of packet sizes. To evaluate how this affects performance, we consider transport of multiple MPEG4-AVC streams over 802.11b. Assume the arriving packet rate (via Ethernet interface) at the AP during a specific 1*sec* window is *N* packets/sec. The actual information bits (video data bits) should be transmitted within the duration:

$T_{video\text{-}bits}$ = 1,000,000 - (N * $T_{overhead}$) µs, that remain after accounting for overheads. At 11*Mbps* channel rate, at most :

$$B_{video\text{-}bits} = (T_{video\text{-}bits}\,\mu s * 11Mbps)$$

bits may be transmitted (in time $T_{video\text{-}bits}$).

If the video bit rate exceeds $B_{video\text{-}bits}$ (1*sec* window considered), packets get buffered at the AP, and either incur queuing delays or are discarded upon buffer overflow. Thus ***packet rate*** is an important performance indicator. The MTU size used for packetization by the RTP layer determines the packet rate in each video stream arriving at the AP.

Example: We consider the Tempete reference video[1], compressed in MPEG4-AVC(CIF, QP=25). The average bitrate per stream is 1.9Mbps. The second and third streams are 0.1 sec and 0.2 sec delayed versions of the first and there is no traffic shaping.

Figure 7 depicts the performance of the three multiplexed video streams with different MTU size for 8 consecutive 1*sec* windows. With an RTP MTU size of 892 bytes, the video bit rate exceeds the permitted rate in several 1*sec* windows. This depicts the situation when incoming packet rate at the Ethernet interface of the AP exceeds the departure rate at the radio interface.

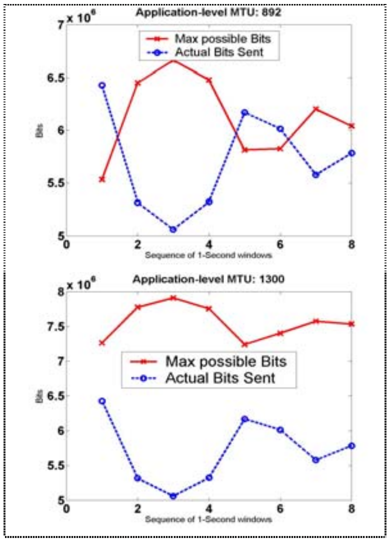

Figure 7. Video Throughput : Varying MTU (892 & 1300 bytes)

When the RTP MTU size increased to 1300 bytes, number of packets constituting each frame reduced, resulting in fewer packets in the 1*sec* window. Thus lesser time is spent transmitting packet overheads, allowing for more video bits instead as depicted.

Conclusions and Future Work

We have clearly demonstrated that multicast of multiple video streams over a wireless channel operating in DCF mode is faced by a significant bottleneck due to 802.11b MAC and physical layer overheads in DCF mode. Bandwidth efficiency gained by producing small B-frames is lost due to large overhead incurred while transmitting that packet over the WLAN. The following are possible solutions to alleviate the problem:

1. *802.11 PCF mode:* The DCF mode analysis showed that channel contention and collision avoidance mechanisms (DIFS, contention window, backoff etc) result in over 50% of the overhead. We identify Point Coordination Function (PCF) mode to be a strong candidate for video broadcast. Here the access point periodically polls each sender for data and allocates time slots when each sender can exclusively transmit. In the simplistic situation of video broadcast from the AP with no feedback of any kind, the AP may maintain a monopoly on the channel resulting in an extremely bandwidth efficient operation. No time is spent on overheads except for the PCF Inter frame Slot (PIFS) in between frames. This can be chosen based on the channel condition, as the frame error rate is proportional to frame size.

2. *Smart AP Queues:* We inferred from the WLAN experiments that AP queues implemented a simple DropTail method of packet discards. If APs are utilized to carry video traffic, smarter queues need to be implemented that prioritize IDR frames over P and B frames. This can improve received video quality considerably.

3. *Statistical Multiplexer:* Alternately a statistical multiplexer may be introduced between the AP and the set of encoders that could signal the latter to adapt output packet rates based on queuing status at the AP.

We intend to report on results of these solutions in future work.

Acknowledgements

We would like to thank Ivan Seskar and Dipankar Raychaudhuri of WINLAB, Rutgers University, and Alan Stein and Karel VanDoorsellaer of Thomson Inc. for their help and insights at various stages.

References

[1] http://peach.eas.asu.edu/, "Video Traces for Network Performance Evaluation"

[2] J.Jun, P. Peddabachagari, M. Sichitiu. "Theoretical Maximum Throughput of IEEE 802.11 and its Applications", Second International Conference on Network Computing and Applications, April 2003.

[3] IEEE 802.11, 1999 Edition (ISO/IEC 8802-11: 1999) Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications

[4] IEEE 802.11b-1999 Supplement to 802.11-1999,Wireless LAN MAC and PHY specifications: Higher speed Physical Layer (PHY) extension in the 2.4 GHz band

[5] "RTP: Real Time Protocol" Schulzrinne, H., S. Casner, R. Frederick, and V. Jacobson, RFC 1889.