2007-10-01

The Effects of Contention among Stations on Video Streaming Applications over Wireless Local Area Networks: an Experimental Approach

Nicola Cranley
Dublin Institute of Technology

Tanmoy Debnath
Dublin Institute of Technology, Tanmoy.Debnath@dit.ie

Recommended Citation

Cranley, N. & Debnath, T. (2007) The effects of contention among stations on video streaming applications over wireless local area networks: an experimental approach. *Information Technology and Telecommunications Conference, 2007 (ITT 2007),* Blanchardstown, Dublin, Ireland, October, 2007.

The Effects of Contention among stations on Video Streaming Applications over Wireless Local Area Networks- an experimental approach

Nicola Cranley, Tanmoy Debnath, Mark Davis
Communications Network Research Institute,
School of Electronic and Communications Engineering,
Dublin Institute of Technology,
Dublin 8, Ireland
nikki.cranley@cnri.dit.ie, tanmoy.debnath@cnri.dit.ie, mark.davis@dit.ie

Abstract

Multimedia streaming applications have a large impact on the resource requirements of the WLAN. There are many variables involved in video streaming, such as the video content being streamed, how the video is encoded and how it is sent. This makes the role of radio resource management and the provision of QoS guarantees extremely difficult. For video streaming applications, packet loss and packets dropped due to excessive delay are the primary factors that affect the received video quality. In this paper, we experimentally analyse the effects of contention on the performance of video streaming applications with a given delay constraint over IEEE 802.11 WLANs. We show that as contention levels increase, the frame transmission delay increases significantly despite the total offered load in the network remaining constant. We provide an analysis that demonstrates the combined effects of contention and the playout delay constraint have on the video frame transmission delay.

Keywords: Video Streaming, Multimedia, WLAN, Quality of Service

1. INTRODUCTION

Streaming multimedia over wireless networks is becoming an increasingly important service [1] [2]. This trend includes the deployment of WLANs that enable users to access various services including those that distribute rich media content anywhere, anytime, and from any device. There are many performance-related issues associated with the delivery of time-sensitive multimedia content using current IEEE 802.11 WLAN standards. Among the most significant are low delivery rates, high error rates, contention between stations for access to the medium, back-off mechanisms, collisions, signal attenuation with distance, signal interference, etc. Multimedia applications, in particular, impose onerous resource requirements on bandwidth constrained WLAN networks. Moreover, it is difficult to provide QoS in WLAN networks as the capacity of the network also varies with the offered load.

Packet loss and packets dropped due to excessive delay are the primary factors that have a negative effect on the received video quality. Real-time multimedia is particularly sensitive to delay, as multimedia packets require a strict bounded end-to-end delay. Every multimedia packet must arrive at the client before its playout time, with enough time to decode and display the contents of the packet. If the multimedia packet does not arrive on time, the playout process will pause and the packet is effectively lost. In a WLAN network, in addition to the propagation delay over the air interface, there are additional sources of delay such as queuing delays in the Access Point (AP), i.e. the time required by the AP to gain access to the medium and to successfully transmit the packet which may require a number of retransmission attempts.

Multimedia applications typically impose an upper limit on the tolerable packet loss. Specifically, the packet loss ratio is required to be kept below a threshold to achieve acceptable visual quality. For example, a large packet loss ratio can result from network congestion causing severe degradation of multimedia quality. Even though WLAN networks allow for packet retransmissions in the event of an unsuccessful transmission attempt, the retransmitted packet must arrive before its playout time or within a specified delay constraint. If the packet arrives too late for its playout time, the packet is effectively lost. Congestion at the AP often results in queue overflow, which results in packets being dropped from the queue. In this way, packet loss and delay can exhibit temporal dependency or burstiness [3]. Although, error resilient encoded video and systems that include error concealment techniques allow a certain degree of loss tolerance [4], the ability of these schemes to conceal bursty and high loss rates is limited.

In IEEE 802.11b WLANs, the AP is usually the critical component that determines the performance of the network as it carries all of the downlink transmissions to wireless clients and is usually where congestion is most likely to occur. There are two primary sources of congestion in WLAN networks. The first is where the AP becomes saturated due to a heavy downlink load which results in packets being dropped from its transmission buffer and manifests itself as bursty losses and increased delays [5]. In contrast, the second case is where there are a large number of wireless stations contending for access to the medium and this results in an increased number of deferrals, retransmissions and collisions on the WLAN medium. The impact of this manifests itself as significantly increased packet delays and loss. For video streaming applications, this increased delay results in a greater number of packets arriving at the player too late for playout and being effectively lost. In this paper, we experimentally investigate this second case concerning the effects of station contention on the performance of video streaming applications.

The remainder of this paper is structured as follows. Section 2 provides an analysis of the video clips used during the experiments. Section 2.1 and 2.2 describe the experimental test bed and experimental results respectively. We focus on a single video content type and show in detail how the delay and loss rates are affected by increased station contention. We show the effects of contention on the performance of the video streaming application for a number of different video content types. We provide an analysis that shows how the play out delay constraint and the number of contending stations affect the video frame transmission delay. Finally we present some conclusions and directions for future work in section 3.

2. VIDEO CONTENT PREPARATION AND ANALYSIS

In the experiments reported here, the video content was encoded using the commercially available X4Live MPEG-4 encoder from Dicas. This video content is approximately 10 minutes in duration and was encoded as MPEG-4 SP with a frame rate of 25 fps, a refresh rate of one I-frame every 10 frames, CIF resolution and a target CBR bit-rate of 1Mbps using 2-pass encoding. Although a target bit rate is specified, it is not always possible for an encoder to achieve this rate. Five different video content clips were used during the experiments. DH is an extract from the film ‘Die Hard’, DS is an extract from the film ‘Don’t Say a Word’, EL is an extract from the animation film ‘The Road to Eldorado’, FM is an extract from the film ‘Family Man’, and finally JR is an extract from the film ‘Jurassic Park’. The video clips were prepared for streaming by creating an associated hint track using MP4Creator from MPEG4IP. The hint track tells the server how to optimally packetise a specific amount of media data. The hint track MTU setting means that the packet size will not exceed in the MTU size.

It is necessary to repeat the experiments for a number of different video content types since the characteristics of the streamed video have a direct impact on its performance in the network. Each video clip has its own unique signature of scene changes and transitions which affect the time varying bitrate of the video stream. Animated videos are particularly challenging for encoders since they generally consist of line art and as such have greater spatial detail.

TABLE 1 CHARACTERISTICS OF ENCODED VIDEO CLIPS

Clip	Mean Packet Size (B)	Mean Bit Rate (kbps)	Frame Size (B)		I-Frame Size (B)		P-Frame Size (B)		Peak-to-Mean Ratio
			Max.	Avg.	Max.	Avg.	Max.	Avg.	
DH	889	910	16762	4617	16762	7019	12783	812	3.63
DS	861	682	12734	3480	12734	6386	10600	713	3.66
EL	909	1199	27517	6058	27517	14082	14632	1587	4.54
FM	894	965	17449	4903	17449	10633	15078	1188	3.56
JR	903	1081	17299	5481	17299	8991	13279	1006	3.16

Table 1 summarizes the characteristics of the encoded video clips used during the experiments. The second column shows the mean packet size of the clip as it is streamed over the network and the third column shows the mean bit-rate of the video clip. The following columns show the maximum video frame size and the mean video frame size in bytes as measured over all frames, over I-frames only and P-frames only. Finally, the last column shows the peak-to-mean ratio of the video frames. It can be seen that despite encoding the video clips with the same video encoding parameters, the video clips have very different characteristics. Despite all the video clips being prepared with exactly same encoding configuration, due to the content of the video clips the mean and maximum I and P frames vary considerably in size.

2.1 EXPERIMENTAL TEST BED

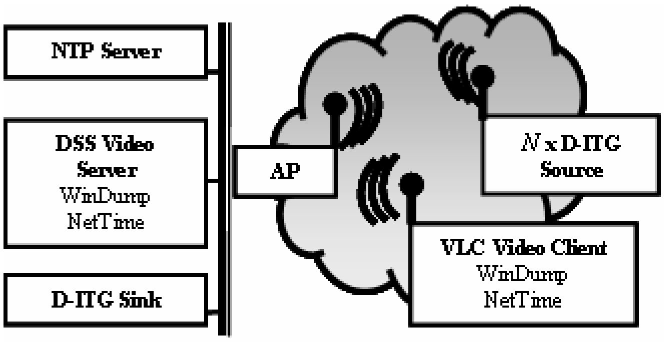

Fig. 1: Experimental Setup

To demonstrate the effects of station contention on video streaming applications, the video server was set up on the wired network and streamed the video content to a wireless client via the AP (Figure 1). The video streaming system consists of the Darwin Streaming Server (DSS) [6] acting as the video server and VideoLAN Client (VLC) [7] as the video client. DSS is an open-source, standards-based streaming server that is compliant to MPEG-4 standard profiles, ISMA streaming standards and all IETF protocols. The DSS streaming server system is a client-server architecture where both client and server consist of the RTP/UDP/IP stack with RTCP/UDP/IP to relay feedback messages between the client and server. The video client VLC allowed the received video stream to be recorded to a file for subsequent video quality analysis. Both the video client and server were configured with the packet monitoring tool WinDump [8] and the clocks of both the client and server are synchronised before each test using NetTime [9]. However, in spite of the initial clock synchronisation, there was a noticeable clock skew observed in the delay measurements and this was subsequently removed using Paxson's algorithm as described in [10]. The delay is measured here as the difference between the time at which the packet was received at the link-layer of the client and the time it was transmitted at the link-layer of the sender.

There are a number of wireless background load stations contending for access to the WLAN medium where their traffic load directed towards a sink station on the wired network. The background uplink traffic was generated using Distributed Internet Traffic Generator (D-ITG) [11]. The background traffic load had an exponentially distributed inter-packet time and an exponentially distributed packet size with a mean packet size of 1024B. To maintain a constant total background load of 6 Mbps, the mean rate of each background station was appropriately decreased as the number of background stations was increased.

2.2 RESULTS

Video streaming is often described as “bursty” and this can be attributed to the frame-based nature of video. Video frames are transmitted with a particular frame rate. For example, video with a frame rate of 25 fps will result in a frame being transmitted every 40ms. In general, video frames are large, often exceeding the MTU of the network and results in a several packets being transmitted in a burst for each video frame. The frequency of these bursts corresponds to the frame rate of the video [12].

In a WLAN environment, the bursty behaviour of video traffic has been shown to results in a sawtooth-like delay characteristic [13]. Consider, a burst of packets corresponding to a video frame arriving at the AP. The arrival rate of the burst of packets is high and typically these packets are queued consecutively in the AP’s transmission buffer. For each packet in the queue, the AP must gain access to the medium by deferring to a busy medium and decrementing its MAC back-off counter between packet transmissions. This process occurs for each packet in the queue at the AP causing the delay to vary with a sawtooth characteristic. It was found that the duration and height of the sawtooth delay characteristic depends on the number of packets in the burst and the packet size. This is to be expected since when there are more packets in the burst, it takes the AP longer to transmit all packets relating to this video frame.

To describe this sawtooth characteristic we have defined the Inter-Packet Delay (IPD) as the difference in the measured delay between consecutive packets within a burst for a video frame at the receiver. When there are no other stations contending for access to the medium, he IPD is in the range 0.9ms to 1.6ms for 1024B sized packets. This delay range includes the DIFS and SIFS intervals, data transmission time including the MAC Acknowledgement as well as the randomly chosen Backoff Counter values of the 802.11 MAC mechanisms contention windows in the range 0-31. This can be seen in Figure 2 where there is an upper plateau with 32 spikes corresponding to each of the possible 32 Backoff Counter values with a secondary lower plateau that corresponds to the proportion of packets that were required to be retransmitted through subsequent doubling of the contention window under the exponential binary backoff mechanism employed in the 802.11 MAC.

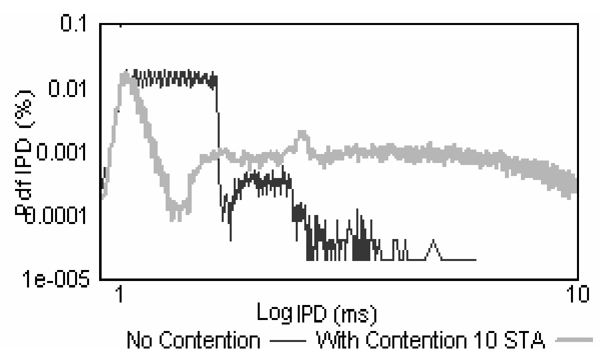

Fig. 2: PDF of the IPD with and without contention

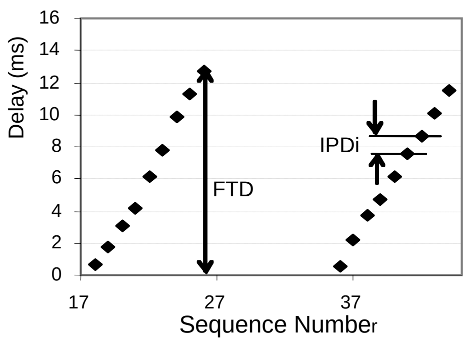

Fig. 3: IPD and FTD Relationship

As contention levels increase, all stations must pause decrementing their Backoff Counter more often when another station is transmitting on the medium. As the level of contention increases, it takes longer to win a transmission opportunity and consequently the maximum achievable service rate is reduced which increases the probability of buffer overflow. In these experiments, the nature of the arrivals into the buffer remains constant, i.e. only the video stream is filling the AP's transmission buffer with packets, but by varying the number of contending stations we can affect the service rate of the buffer and thereby its ability to manage the burstiness of the video stream. This can be seen in Figure 2 where there is a long tail in the distribution of IPD values for the 10 station case. In this case, 10 wireless background traffic stations are transmitting packets to the wired network via the AP's receiver. The aggregate load from these stations is held constant as the number of background stations is increased.

For video streaming applications, not only is the end-to-end delay important, but also the delay incurred transmitting the entire video frame from the sender to the client. A video frame cannot be decoded or played out at the client until all of the constituent video packets for the frame are received correctly and on time. For this reason, in our analysis we also consider the video Frame Transmission Delay (FTD), i.e. the end-to-end delay incurred in transmitting the entire video frame and is related to the number of packets required to transmit the entire video frame and the queuing delay in the AP buffer for the first video packet in the burst to reach to head of the queue. Figure 3 shows the relationship between the IPD and FTD for two consecutive video frames. In our analysis, we also consider the loss rate and the Playable Frame Rate (PFR). The PFR is inferred by using the statistical techniques described in [14]. The loss rate corresponds to packets that have failed to be successfully received as well as those packets that have been dropped as a result of exceeding the Delay Constraint (Dc). If packets arrive too late exceeding Dc, these packets are effectively dropped by the player since they have arrived too late to be played out.

2.2.1 THE EFFECTS OF CONTENTION ON STREAMED VIDEO

In this section, we experimentally demonstrate the effects of contention on video streamed applications. We shall begin by focusing on a single video clip DH being streamed from the wired network via the AP to a wireless client. This particular clip was chosen since it is representative of a typical non-synthetic video stream. Table 2 presents the mean performance values for the video clip DH over the test period with increased contention. It can be seen that the mean delay, loss rate, FTD and IPD increase with increased contention. In this work we have set the Dc to 500ms which is the delay constraint for low latency real-time interactive video.

TABLE 2 MEAN PERFORMANCE VALUES FOR DH CLIP WITH INCREASED CONTENTION (DC = 500MS)

Performance Metric	0STA	3STA	4STA	5STA	6STA	7STA	8STA	9STA	10STA
Mean Delay (ms)	10.43	29.62	30.97	37.91	63.63	105.75	174.91	311.71	395.27
FTD (ms)	11.50	36.62	37.96	45.39	71.76	115.61	186.05	325.01	406.83
IPD (ms)	1.24	3.73	3.75	3.97	4.34	4.82	5.27	5.66	5.95
Mean Loss Rate (Dc > 500ms)	0.00	0.01	0.01	0.03	0.08	0.15	0.23	0.34	0.41
PFR (fps)(Dc > 500ms)	25.00	25.00	23.00	21.83	19.04	16.91	14.02	10.51	9.92

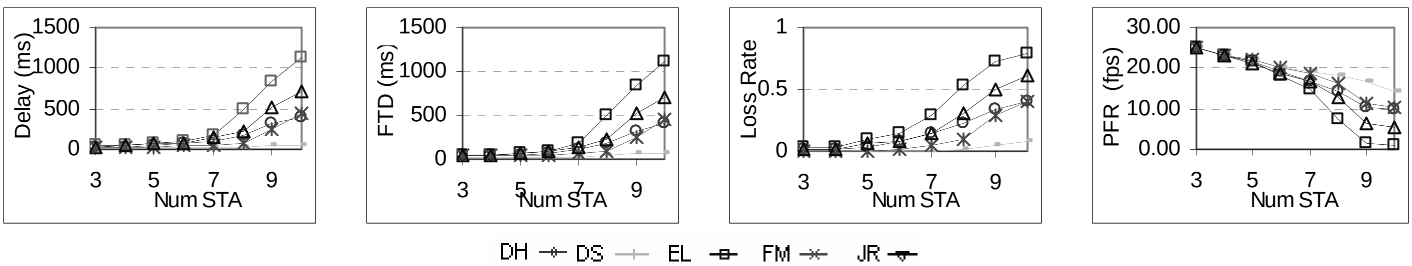

Fig.4 Mean values for a number of video clips for a fixed total offered uplink load with increased number of contributing stations (a) Mean Delay, (b) Mean FTD, (c) Average loss rate with a Dc of 500ms, (d) Inferred PFR with a Dc of 500ms.

It can be seen that when there are no background contending stations, the mean packet delay is about 10ms. As the number of contending stations increases from 3 to 7 to 10, the mean delay increases to 30ms, to 100ms to 400ms respectively. This can be explained from the growing tail of the IPD distribution as shown in Figure 2. As the number of contending stations is increased from 3 to 7 to 10 stations with a Dc of 500ms, the mean loss rate including packets dropped due excessive delay is increased from 1% to 15% to 41% respectively. This in turn affects the ability of the codec to decode the video frames since there is increased likelihood that packets will not arrive within the given delay constraint.

The experiment was repeated for the other video clips all encoded with the same encoding configuration but having different content complexity characteristics. Figure 3 shows the mean performance metrics for different content types with increased contention. For all content types, it can be seen that the mean packet delay and FTD increases with increased contention as shown in Figures 4(a) and 4(b). Figure 4(c) shows the mean loss rate over the test period for each of the video clips where it can be seen that there is a dramatic increase in the mean loss rate when the number of contending stations exceeds 7 stations when a delay constraint of 500 ms is imposed on the system which results in an even greater impact of the contention on performance. Figure 4(d) shows the PFR that is statistically inferred from the packet loss and delay. Apart from the impact of contention, Figures 4(a)-4(d) also highlight the impact of the video content where it can be seen that the animation clip *EL* is the most severely affected by increased contention whilst the clip *DS* is the least affected. The high complexity of the animation clip *EL* is due to frequent scene cuts and line art within the scene that affects the burstiness of the encoded video sequence since much more information is required to encode the increased scene complexity.

2.2.2 ANALYSIS

In this section we shall generalize the results presented in the previous section to account for all content types and a given delay constraint. For video streaming applications, there is a tradeoff between acceptable delay and tolerable packet loss. A delay constraint imposes an upper limit on this tradeoff since the lower the delay constraint, the greater the probability of packets being dropped due to exceeding the delay constraint.

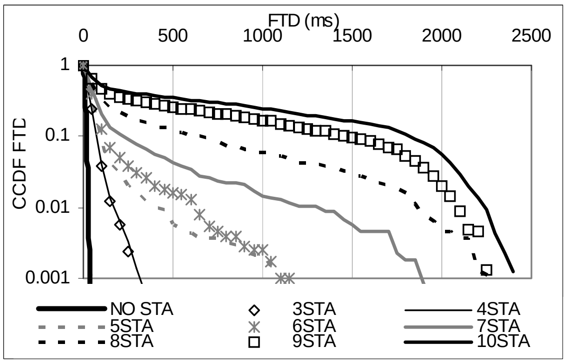

Fig.5 Generalized distribution of the FTD with increased contention

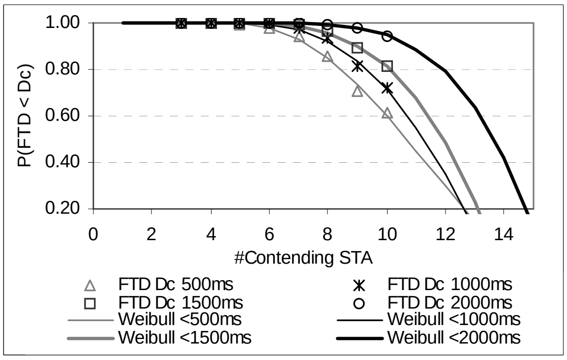

Fig.6 Fitted Weibull Distribution to CDF of FTD with Dc

TABLE 3: CDF OF FTD BELOW THE PLAYOUT DELAY CONSTRAINT, DC

	Number of Contending Stations							
Dc (ms)	**3STA**	**4STA**	**5STA**	**6STA**	**7STA**	**8STA**	**9STA**	**10STA**
500	1.000	1.000	0.994	0.984	0.957	0.877	0.740	0.653
1000	1.000	1.000	0.996	0.998	0.986	0.942	0.832	0.752
1500	1.000	1.000	1.000	1.000	0.994	0.971	0.903	0.836
2000	1.000	1.000	1.000	1.000	1.000	0.995	0.980	0.945
2500	1.000	1.000	1.000	1.000	1.000	1.000	1.000	1.000

In our analysis we focus on the FTD since all or most of the packets belonging to a video frame packet burst must be received in order for the video frame to be decoded on the client device.

Figure 5 shows the Complementary Cumulative Distribution Function (CCDF) of the FTD averaged over all content types with an increasing number of contending stations. For example consider a video streaming application with a Dc of 500ms, it can be seen that with 4 contending background stations, the FTD is always less than 500ms. However with 6, 8, and 10 background contending stations, statistically 2%, 12% and 35% of video frames will have an FTD that exceeds a Dc of 500ms. The statistical distribution of the FTD has been summarized in Table 3 which presents the CDF of the FTD for different values of Dc and with an increased number of contending stations. It can be seen that when there are 10 contending stations, with a Dc of 500ms 65% of video frames will arrive within this upper delay bound whereas 95% of video frames will arrive within a Dc of 2000ms. Figure 6 shows a plot of the fitted Weibull distribution to the probability of the FTD arriving within a given Dc with an increased number of contending stations. The Weibull distribution fit had a correlation coefficient of over 99.5% in all cases. The shape and scale parameters are related to the number of contending stations and the delay constraint of the video. This distribution can be used to provide statistical FTD guarantees by a resource management system to perform admission control to assess the impact of station association on the video streaming applications. Furthermore, adaptive streaming systems can use the statistical characterization of FTD to adaptively dimension the playout buffer on the client device or to adapt the number of packets per video frame i.e. the bitrate of the video stream based on current contention load conditions since by reducing the number of packets per video frame, the FTD is reduced.

3. CONCLUSIONS

In this paper, we have experimentally investigated the effects of station contention on streaming video over IEEE 802.11b WLAN networks. Video is a frame-based media where video frames are transmitted from the server to the client at regular intervals that is related to the frame rate of the video. In general, several packets are required to transmit the video frame. The video frame cannot be decoded at the client until all the packets for the video frame have been received. In this way, loss and delay have a serious impact on the performance of video streaming applications. Loss can occur due to packets reaching their retransmission limit following repeated unsuccessful attempts and packets that are dropped due to incurring excessive delays resulting in them arriving too late to be decoded.

Through experimental work, we have demonstrated that as the number of contending stations increases, while a maintaining a constant total offered load, the video streaming application experiences increased delays. These delays are due to the 802.11b MAC mechanism where stations must contend for access to the medium. As the number of stations contending for access to the medium increases, the AP must defer decrementing the Backoff Counter while another station is transmitting on the medium. Experimental results show that the performance degrades with increased contention despite the offered load in the network remaining the same. Furthermore we have shown that the complexity of the video content affects the degree of performance degradation. In our analysis we focused on the Frame Transmission Delay (FTD) which is the delay incurred transmitting the entire video frame from the server to the client. The FTD is important for video streaming applications since a video frame cannot be correctly decoded at the client until all of the packets relating to the video frame have been received within a given delay constraint. The delay constraint imposes an upper bound delay threshold for the video frames. Packets that exceed this delay constraint are effectively lost since they have not been received at the client in time for play out. We statistically analysed the results to determine the probability of the FTD being within a given delay constraint and have shown that this can be modeled as a Weibull distribution. This analysis can be used as part of a WLAN access control scheme or used in a cross-layer contention-aware video playout buffering algorithm. The QoS capabilities of the IEEE 802.11e QoS MAC Enhancement standard [15] facilitates new management mechanisms by allowing for traffic differentiation and prioritization. Work is ongoing [16] [17] with 802.11e standard. Further work is required to investigate the benefits for video streaming afforded by this standard.

ACKNOWLEDGEMENT

The support of the Science Foundation Ireland, grant 03/IN3/1396, under the National Development Plan is gratefully acknowledged.

REFERENCES

[1] J. Wexler, "2006 Wireless LAN State-of-the-Market Report", Webtorials, July 31, 2006, [Online]. Available: http://www.webtorials.com/abstracts/WLAN2006.htm

[2] Insight Research Corp., "Streaming Media, IP TV, and Broadband Transport: Telecommunications Carriers and Entertainment Services 2006-2011", Insight Research Corp., April 2006, [Online]. Available: http://www.insight-corp.com/reports/IPTV06.asp

[3] S. Moon, J. Kurose, P. Skelly, D. Towsley. "Correlation of packet delay and loss in the Internet". Technical report, University of Massachusetts, January 1998.

[4] Y. Wang, S. Wengers, J. Wen, A.K. Katsaggelos, "Error resilient video coding techniques", IEEE Signal Processing Mag., vol. 17, no. 4, pp. 61-82, July 2000

[5] N. Cranley, M. Davis, "The Effects of Background Traffic on the End-to-End Delay for Video Streaming Applications over IEEE 802.11b WLAN Networks", 17th Annual IEEE Personal, Indoor and Mobile Communications, PIMRC Helsinki, Finland, September 2006

[6] Darwin Streaming Server, http://developer.apple.com/darwin/projects/streaming/

[7] VideoLAN Client, http://www.videolan.org/

[8] WinDump, http://windump.polito.it/

[9] NetTime, http://nettime.sourceforge.net/

[10] S. B. Moon, P. Skelly, D. Towsley, "Estimation and Removal of Clock Skew from Network Delay Measurements", in Proc. of IEEE InfoComm'99, March 1999

[11] Distributed Internet Traffic Generator (D-ITG), http://www.grid.unina.it/software/ITG/download.php

[12] A. C. Begen, Y. Altunbasak, "Estimating packet arrival times in bursty video applications," in Proc. IEEE Int. Conf. Multimedia and Expo (ICME), Amsterdam, The Netherlands, July 2005

[13] N. Cranley, M. Davis, "Delay Analysis of Unicast Video Streaming over WLAN", 2nd IEEE International Conference on Wireless and Mobile Computing, Networking and Communications, WiMob 2006, Montreal, Canada, June 2006

[14] N. Feamster, H. Balakrishnan, "Packet loss recovery for Streaming Video", Proc. of 12th International Packet Video Workshop, April 2002

[15] IEEE STD 802.11e, September, 2005 Edition, IEEE Standards for Local and Metropolitan Area Networks: Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements

[16] Nicola Cranley, Mark Davis, "Video Frame Differentiation for Streamed Multimedia over Heavily Loaded IEEE 802.11e WLAN using TXOP", IEEE PIMRC 2007, Athens, Greece, September 2007

[17] Nicola Cranley, Tanmoy Debnath, Mark Davis, "An Experimental Investigation of Parallel Multimedia Streams over IEEE 802.11e WLAN Networks using TXOP", IEEE ICC 2007, Glasgow, Scotland, June 2007