

2006-09-01

The Effects of Background Traffic on the End-to-End Delay for Video Streaming Applications over IEEE 802.11b WLAN Networks

Nicola Cranley

Technological University Dublin, nicola.cranley@tudublin.ie

Mark Davis

Technological University Dublin, mark.davis@tudublin.ie

Follow this and additional works at: <https://arrow.tudublin.ie/commcon>



Part of the [Systems and Communications Commons](#)

Recommended Citation

Cranley, N. & Davis, M. (2006) The Effects of Background Traffic on the End-to-End Delay for Video Streaming Applications over IEEE 802.11b WLAN . *The 17th Annual IEEE International Symposium on Personal, Indoor and Mobile Radio Communications (PIMRC'06)*, Helsinki, Finland. 11-14 September, 2006.

This Conference Paper is brought to you for free and open access by the Communications Network Research Institute at ARROW@TU Dublin. It has been accepted for inclusion in Conference papers by an authorized administrator of ARROW@TU Dublin. For more information, please contact arrow.admin@tudublin.ie, aisling.coyne@tudublin.ie, vera.kilshaw@tudublin.ie.

Funder: Science Foundation Ireland (Grant 03/IN3/1396)

THE EFFECTS OF BACKGROUND TRAFFIC ON THE END-TO-END DELAY FOR VIDEO STREAMING APPLICATIONS OVER IEEE 802.11B WLAN NETWORKS

Nicola Cranley
Communications Network Research Institute,
School of Electronic and Communications Engineering,
Dublin Institute of Technology,
FOCAS Institute,
Dublin 8, Ireland

Mark Davis
Communications Network Research Institute,
School of Electronic and Communications Engineering,
Dublin Institute of Technology,
FOCAS Institute,
Dublin 8, Ireland

ABSTRACT

The bursty nature of video streaming applications is due to the frame-based structure of video and this has an important impact on the resource requirements of the WLAN, affecting its ability to provide Quality of Service (QoS) particularly under heavily loaded conditions. In this paper we analyse this bursty behaviour in depth. We show how each video frame is queued at the AP causing the packet delay to vary in a sawtooth manner that is related to the frame rate, the number of packets per video frame, and the packet size. We infer the maximum background traffic load that can be supported so that it does not negatively impact on the video streaming application. We demonstrate that there is a critical threshold load value above which the AP can no longer reliably support the video stream and compare it to the threshold load values calculated through analysis. Using this knowledge, the AP can employ resource allocation mechanisms to regulate the incoming traffic to the AP transmission queue so that QoS can be provided for streaming applications.

Introduction

Streaming multimedia over wireless networks is becoming an increasingly important service [1]. 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 standards. Among the most significant are low delivery rates, high error rates due to media characteristics, 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 [2,3]. Under these conditions it is difficult to provide any Quality of Service (QoS) guarantees. In particular the delay constraints associated with streaming multimedia pose the greatest challenge since real-time multimedia is particularly sensitive to delay as multimedia packets require a strict bounded end-to-end delay. That is, every multimedia packet must arrive at the client before its playout time, with enough time to decode and display the packet. If the multimedia packet does not arrive on time, the packet is effectively lost. In a WLAN environment, lost or corrupted packets are repeatedly retransmitted until either the retransmitted packet is

successfully ACKed by the receiving station or until the retransmission limit has been reached. If a packet has expired,

there is no need for it to continue along its path to the client since its contents will be worthless when it arrives.

There are a large and diverse number of variables that must be taken into consideration when analysing multimedia streaming applications each of which has an effect on the performance of the system. Such variables include the complexity of the content, the compression scheme, the encoding configuration including the bitrate, the frame rate, the size of the video frames, the packetisation scheme used to transmit the video, and the streaming server used. In this work, the end-to-end delay for unicast video streaming over a WLAN is analysed in relation to frame rate, frame size, packet rate, and packet size. In particular we shall analyse how the performance is affected when there is a background traffic load. We show that there is a critical threshold load value that is related to the packet size, and offered load. Once this threshold background load value has been exceeded, the video streaming application experiences excessive delays.

This paper is structured as follows. The experimental test bed is described in Section 2. In Section 3 we present results that demonstrate the behaviour of streamed video in a WLAN environment in the absence and presence of a background traffic load. We show that there is a threshold load value beyond which the video stream is negatively affected by excessive delays. Using experimental results we describe a method that can be used to infer the maximum background traffic load under ideal conditions and compare the experimentally observed and predicted intervals when the load exceeds this threshold load value. Section 4 presents some conclusions and directions for future work.

Experimental Test Bed

To evaluate unicast video streaming a video server was set up on the wired network and streamed to a wireless client via the AP. The AP used was a Cisco Aironet 1200 IOS version 12.3(8)JA. Both the client and server were configured with the packet monitoring tool, WinDump [4] and the clocks of both the client and server are synchronised before each test using NetTime [5] and any clock skew was removed using Paxson's algorithm as described in [6]. Given the large number of encoding parameters that can be varied whilst preparing the video content for streaming over the network, only the frame rate of the video and the size of the video frames are varied. As a result, the mean transmitted bit rate varies in an Additive Increase Proportional Decrease (AIPD) manner and reaches a maximum bit rate of 2.1Mbps after

1700sec as shown in Figure 1. The video was generated and streamed across the network using RTPTools [7]. The mean video frame sizes were varied from 3.1kB, 6.1kB and 9.2kB every 100sec and the frame rate was increased from 10fps to 30fps in steps of 5fps every 300sec. When streaming MPEG-4 files, each video and audio track must have its own associated hint track. The hint track setting can be used to indicate the Maximum Transmission Unit (MTU) of the packets to be sent. Throughout the remainder of this paper the MTU of the video stream is denoted as S_{VID} and the MTU of the background traffic is denoted as S_{BAK} . In these experiments the effects of using S_{VID} equal to 512B and 1024B are investigated. In this way when using S_{VID} of 512B,

the number of packets, n_{VID} , required to transmit the video frames varied from {6, 12, 18} packets and from {3, 6, 9} packets when using S_{VID} of 1024B.

Results

A. Delay with No Background Traffic

End-to-end delay is of critical importance to real-time streaming applications since if a packet is delayed past its playout time, the packet is effectively lost. 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. In our analysis, the delay is measured as the difference between the time at which the packet was received at link-layer of the client and the time it was transmitted at the link-layer of the sender. 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. 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 video frame cannot be decoded or played out at the client until all or most of the video packets for the particular video frame are received in time. Although, error resilient encoded video and systems that include error concealment techniques allow for a certain degree of loss tolerance [8], however the ability of these schemes to conceal bursty and high loss rates is limited.

In a WLAN environment, the bursty behaviour of video traffic has a sawtooth delay characteristic. Figure 2 shows how the delay varies for 4 consecutive video frames of size 9.2kB using S_{VID} of 512B and 1024B. It can be clearly seen that when using S_{VID} of 1024B, it takes more time to send each individual packet but since there are fewer packets required to transmit the video frame, it takes less time to send the complete video frame. When the server transmits a video frame, a burst of packets arrive at the AP. The arrival rate of the burst of packets is high and typically these packets are queued consecutively in the AP 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. The duration and height of the sawtooth delay characteristic depends on the number of packets in the burst and the packet size since when there are more packets required to transmit the video frame, it takes the AP longer to transmit all packets. Furthermore, it takes the AP less time to transmit a small packet than a large one. However, it does incur a larger overhead as more transmission opportunities are required.

The Inter-Packet Delay (IPD) is the difference in measured delay between consecutively queued packets of the same size within a burst of packets for a particular video frame at the receiver and gives an indication of the service rate of the AP. It also provides a means of analysing the backoff counter values. The IPD includes the time the AP spends accessing the medium, including DIFS, the backoff counter value, data transmission, SIFS, and the time to receive the MAC Acknowledgement [9]. Figure 3 shows the PDF of the IPD for video streamed with S_{VID} of 512B and 1024B. The IPD




Figure 1: Offered Video Traffic Characteristics

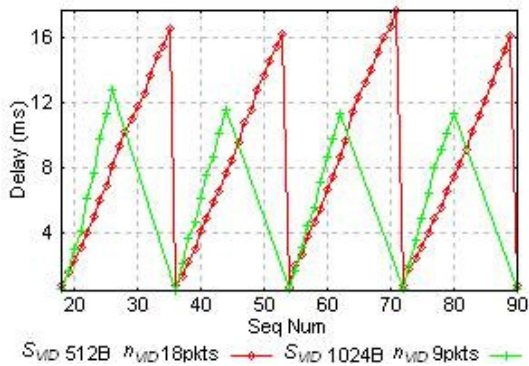


Figure 2: Sawtooth Delay Characteristic

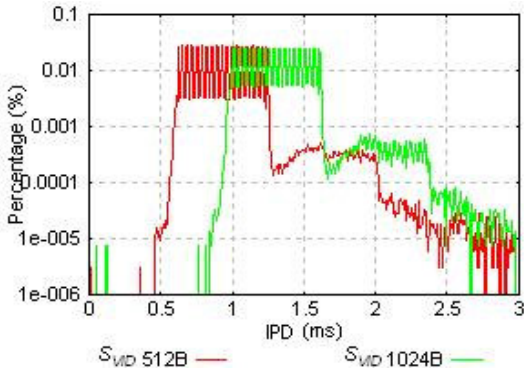


Figure 3: IPD for S_{VID} 1024B and S_{VID} 512B

Table 1: Mean Delay Analysis

Frame Size (kB)	$S_{VID}=512B$		$S_{VID}=1024B$	
	Frame Delay (ms)	Pkt Delay (ms)	Frame Delay (ms)	Pkt Delay (ms)
3.1	5.08	2.70	2.88	1.54
6.1	11.68	6.40	7.22	3.86
9.2	17.16	9.00	11.48	6.06

was experimentally found to have an upper plateau in the range of (0.64ms, 1.28ms) with a mean of 0.96ms for a single 512B packet and in the range of (1.02ms, 1.66ms) for a 1024B packet with a mean of 1.34ms. Comprising this plateau, there are 32 peaks which directly relate to the randomly chosen Backoff Counter values of the 802.11 MAC mechanism which is a random number between 0 and 31 timeslots, where each timeslot is 20μs. The tail of this distribution is related to retransmissions.

Table 1 presents a summary of the results averaged for the frame size and S_{VID} settings. It was found that the IPD remains constant for a particular S_{VID} setting. The frame delay increases with the size of the video frame since there are more packets to be sent and as a consequence, the mean packet delay is greater since there are more packets to be served in the queue at the AP ahead of it.

B. Delay with Background Traffic

A source of background traffic was introduced into the experimental setup. In these experiments, the traffic generator, MGEN, is used to generate a source of background traffic with a load of 1Mbps, 3Mbps and 5Mbps respectively using a packet size, S_{BAK} of 512B and 1024B. This background traffic is streamed via the wired network to the AP and received by a sink station in the WLAN.

It was observed that with a background traffic load of 1Mbps regardless of the packet size of both the video stream and the background traffic, the video stream was unaffected. However, as the background traffic load is increased to 3Mbps, depending on the packet sizes of both the background

traffic and the video traffic, the video stream experiences excessive delays when the total load reaches some threshold value. In Fig. 4(a) it can be seen that the bitrate of the video increases due to an increasing video frame rate and video frames size indicated by a thick line, once the bit rate of the video exceeds a certain threshold bitrate value indicated by a dashed line, the video stream experiences large delays indicated by the thin black line. Once the video bit rate falls below this threshold value, the delay returns to a low value. This threshold bitrate value varies with the packet size of both the video and background traffic. It can be seen that when using a large packet size for the background traffic, the threshold bitrate of the video is higher. Fig. 4(a) and Fig. 4(c) show the case for S_{VID} of 512B with a background load of 3Mbps, the bit rate threshold of the video stream is 1.1Mbps and 2.21Mbps when using S_{BAK} of 512B and 1024B respectively. With a background load of 5Mbps the video stream cannot be supported at all using S_{BAK} of 512B as shown in Fig. 4(c) and Fig. 4(g) however when using S_{BAK} of 1024B, the threshold value is 0.61Mbps as shown in Fig. 4(d). In contrast when using S_{VID} of 1024B and a background load of 3Mbps, the threshold bit rate is 1.84Mbps and 2.21Mbps using S_{BAK} of 512B and 1024B respectively as shown in Fig. 4(e) and Fig. 4(f).

Two interesting events have been captured in Fig. 4. The first is when the arrival rate is just slightly greater than the service rate of the AP as seen in Fig. 4(a) during the time interval 800-900sec. It can be seen that the delay gradually climbs higher and higher as the AP queue is progressively filled. The second interesting event is when the arrival rate is just slightly less than the service rate as seen in Fig. 4(d) during the interval 1100-1200sec. The delay can be seen gradually reducing over time as the AP can service more packets than are arriving allowing the AP time to clear the backlog of queued packets in the transmission buffer resulting in the number of queued packets to slowly decrease over time.

It can be seen that in all cases, when the total combined bitrate of the video stream and background traffic exceeds the threshold bitrate value, the video stream experiences




Figure 4: Experimentally Measured Delay with Background Traffic Load

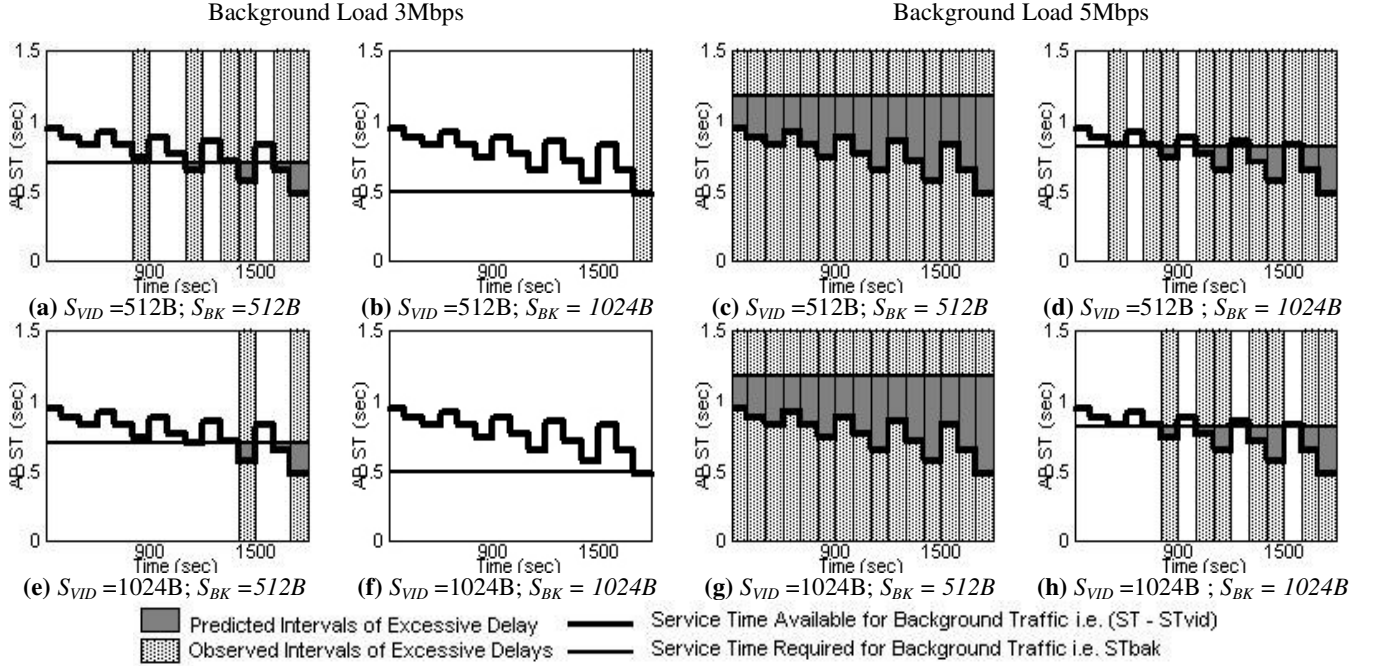


Figure 5: Comparison of Predicted and Observed Intervals of Excessive Delay

excessive delays. However, once the bitrate of the video falls below this threshold value, the delay of the video stream returns to a low value.

C. Inferring the Maximum Ideal Load

In the best-case scenario, once the AP has serviced all the packets relating to a video frame, there remains unused or idle times when the AP can transmit other traffic before the next video frame arrives. In an ideal situation, the background traffic is perfectly interleaved with the video stream, that is, after the AP has serviced all packets in the queue relating to the video frame, the unused time between sending the video frame and the arrival of the next video frame is given to the background traffic.

This load represents the ideal maximum background traffic load that can be serviced. For example, given a video stream encoded with a video frame rate of X_{FPS} , where the number of packets required to transmit the entire video frame is n_{VID} and an IPD, IPD_{VID} , that is related to the packet size, S_{VID} . Then it takes the AP, $(n_{VID} \times IPD_{VID})$ ms to send the video frame. The idle time between sending the video frame and the arrival of the next video frame can then be used to service the background traffic. The time interval for background traffic in this case is therefore:

$$BkInterval = \left(1 \text{sec} / X_{FPS} \right) - (n_{VID} \times IPD_{VID}) \quad (1)$$

During this interval a number of background traffic packets n_{BAK} can be sent. However, this varies with the packet size S_{BAK} , of the background traffic which in turn affects the mean IPD per background packet, IPD_{BAK}

$$n_{BK} = \left\lfloor BkInterval / IPD_{BK} \right\rfloor \quad (2)$$

where $\lfloor \cdot \rfloor$ is the floor function. Given that for every video frame there is a corresponding interval during which the background packet can be sent, this results in a total ideal background traffic load.

$$TotalIdealLoad = n_{BK} \times IPD_{BK} \times 8 \times X_{fps} \quad (3)$$

The ideal background traffic load represents the maximum load that can be supported in such a way so as to not negatively impact on the video stream where the video is transmitted with a variety of frame rates and packet sizes.

Let us consider an interval ST of 1000ms at the AP, we can predict that excessive delays will occur when the total service time for the video stream ST_{VID} and the total service time for the background traffic load ST_{BAK} exceed the interval of interest. In this work, there are no other stations contending for access to the medium, which gives the AP full use of the service time during the interval of interest. However, when there are other stations contending for access to the WLAN medium, the service time at the AP is reduced. For example, the total service time at the AP ST_{VID} is 432ms to send video with X_{FPS} of 25fps, n_{VID} of 18 packets per video frame and a mean IPD_{VID} of 0.96ms. The total service time at the AP, ST_{BAK} , is 703ms to send a background traffic load of 3Mbps with S_{BAK} of 512B. Thus, the total service time for all offered traffic, ST_{TOTAL} (i.e. $ST_{VID} + ST_{BAK}$) exceeds ST , then it can be expected that excessive delays will be experienced by both the video stream and the background traffic.

Using this approach in Fig. 5, we compare the predicted intervals of excessive delay with experimental results presented in the previous section. The predicted intervals of excessive delays are defined as those intervals where the total service time for all offered traffic, ST_{TOTAL} exceeds ST . It can be seen that there is a good correlation between the predicted and observed intervals of excessive delay. Table 2 demonstrates how the experimentally observed maximum bit

Table 3: Summary of Threshold Load Values

	Bak Load 3Mbps $S_{VID}=512B$		Bak Load 5Mbps $S_{VID}=512B$		Bak Load 3Mbps $S_{VID}=1024B$		Bak Load 5Mbps $S_{VID}=1024B$	
	$S_{BAK}=512B$ Fig.4(a)	$S_{BAK}=1024B$ Fig.4(b)	$S_{BAK}=512B$ Fig.4(c)	$S_{BAK}=1024B$ Fig.4(d)	$S_{BAK}=512B$ Fig.4(e)	$S_{BAK}=1024B$ Fig.4(f)	$S_{BAK}=512B$ Fig.4(g)	$S_{BAK}=1024B$ Fig.4(h)
Experimentally Observed Video Threshold Load (Mbps)	1.1	2.21	0	0.61	1.84	2.21	0	0.98
X_{FPS} (fps)	15	30	0	25	25	30	0	20
n_{VID}	18	18	0	6	9	9	0	6
Mean ST_{VID} (ms)	316	631	0	175	348	417	0	185
n_{BK}	732	366	1220	610	732	366	1220	610
Mean ST_{BAK} (ms)	857	566	1428	943	857	566	1428	943
Mean ST_{TOT} (ms)	1172	1198	1428	1118	1205	983	1428	1128

rate thresholds shown in Fig. 4 can be predicted as shown in Fig. 5.

This basic mechanism gives a convenient means to determine the offered loads that will affect the video stream causing excessive delays and as a consequence result in poor QoS. This method of determining the required service time at the AP can be applied to a resource allocation mechanism to regulate the incoming traffic to the AP transmission queue so that these excessive delays can be avoided. Equally, the QoS enhancement facilities of the emerging IEEE 802.11e standard [10] can be used to give priority to the video stream by varying the AIFS and contention window range.

Conclusions

In this paper, we have identified the primary challenge posed by streaming video over WLAN networks. Video is a frame-based media, whereby frames are generated at a particular rate. In general, several packets are required to transmit a video frame to the client and these are transmitted as a burst of packets which are queued for transmission in the AP buffer. Since each packet must wait for the packets in the queue ahead of it to be transmitted, the end-to-end delay varies with a sawtooth-characteristic. The rate at which the delay increases depends on the size of the packet to be transmitted since the AP can send a smaller packet faster than a large one. However, by using a smaller MTU for the video stream the packet bursts are much larger. This sawtooth delay characteristic has been shown by experimentation when there is no background traffic. The end-to-end delay was experimentally measured for video streamed both in the absence and presence of a background traffic load. The results demonstrate that there is a critical threshold load value that is related to the number and size of the packets arriving at the AP transmission buffer. When the load exceeds this threshold load value, the video streaming application is negatively affected incurring excessive delays which cannot be tolerated by real-time or near real-time applications. Using knowledge of the video frame size, frame rate and packetisation of the video, we can infer the maximum background traffic load that can be supported so that it does not negatively impact on the video streaming application. The results demonstrate a good

correlation between the predicted maximum load that observed through experimentation.

By using knowledge of the behaviour of video streaming applications in a WLAN, adaptation algorithms and resource management strategies can be developed to improve QoS for multimedia streaming applications. Work is being conducted in order to establish the effects of contention with other wireless stations with varying traffic loads and packet characteristics on the ability of the AP to serve video packets. With the imminent release of the IEEE 802.11e standard, our future work will measure the effects of varying the AIFS and contention window sizes on video streaming applications.

ACKNOWLEDGEMENT

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

REFERENCES

- [1] J. Wexler, S. Taylor, "2004 Wireless LAN State of the Market Report", Webtorials, Feb. 2004, [Online]. Available: <http://www.webtorials.com/main/resource/papers/taylor/paper4/2004-WLAN.pdf>
- [2] N. Cranley, M. Davis, "Performance Analysis of Network-level QoS with Encoding Configurations for Unicast Video Streaming over IEEE 802.11 WLAN Networks", WirelessCom 2005, Maui, Hawaii, June 2005
- [3] N. Cranley, M. Davis, "Performance Evaluation of Video Streaming with Background Traffic over IEEE 802.11 WLAN Networks", 1st ACM Workshop on Wireless Multimedia Networking and Performance Modelling (WMuNeP 2005), Montreal, Canada, October 2005
- [4] WinDump, <http://windump.polito.it/>
- [5] NetTime, <http://nettime.sourceforge.net/>
- [6] 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
- [7] RTPTools, <http://www.cs.columbia.edu/IRT/software/rtpools/>
- [8] 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
- [9] J. Jun, P. Peddabachagari, M. Sichertiu, "Theoretical Maximum Throughput of IEEE 802.11 and its Applications", in Proceedings of the Second IEEE International Symposium on Network Computing and Applications, Washington, DC, USA, 2003
- [10] Q. Ni, "Performance Analysis and Enhancements for IEEE 802.11e Wireless Networks", IEEE Network, Vol. 19, No. 4, July/August 2005, pp. 21-27.