An Experimental Study of Throughput for UDP and VoIP Traffic in IEEE 802.11b Networks

Sachin Garg
sgarg@avaya.com
Avaya Labs Research
Basking Ridge, NJ USA

Martin Kappes
mkappes@avaya.com
Avaya Labs Research
Basking Ridge, NJ USA

Abstract—**In this paper, we present experimental studies on the throughput of IEEE 802.11b wireless networks for UDP and VoIP traffic. Our experiments show that the maximum data throughput of a single station sending out UDP traffic is 6.1 Mbps. The maximum number of VoIP calls in a single cell of an IEEE 802.11b network is six if the ITU G711a-Law codec is used with 10 milliseconds of audio data per RTP packet. The experiments also show that the effective available bandwidth in the wireless network is reduced by ongoing VoIP connections. Specifically, for the above codec settings, each VoIP connection reduces the bandwidth available for data traffic by 900 Kbps.**

I. Introduction

In the last few years, wireless networks based on the IEEE 802.11b standard have gained popularity and have been widely deployed in enterprises mostly to provide wireless data access from Laptops, PDAs, etc. to the wired infrastructure of the enterprise. They have also been deployed in public hot-spots such as airports, hotels, conference facilities etc., mainly for internetwork connectivity.

The maximal data rate 802.11b currently supports is 11 Mbps. Although task groups, in particular 802.11a and 802.11g [10], are working on allowing higher maximum data dates, it is very likely that 802.11b deployments will continue to operate in both enterprises as well as the residential market for the next few years. As opposed to IEEE 802.3 (also known as Ethernet), where the maximal data rate of the network is indeed close to the throughput observed, the maximal achievable throughput for 802.11b networks is far lower than the data rate due to the nature of 802.11's CSMA/CA medium access protocol. Moreover, the bandwidth is shared among all participants in the network whereas most of today's Ethernet deployments are "switched".

As converged networking in the wired world is becoming more and more popular, it is very likely that in the near future wireless networks will also be increasingly used for voice traffic. In this paper, we present our experimental studies on the throughput for VoIP and UDP traffic in a single cell of an 802.11b network.

We chose UDP as transport layer protocol due to its connection-less nature. A UDP sender "flooding" the network gives an accurate estimate of the actual bandwidth that is available in the network. Results obtained with UDP constitute an upper bound for the throughput possible with TCP as transport layer protocol in any of the scenarios we studied. Another reason why we chose UDP is the 802.11 MAC. It was constructed in such a way that higher frame loss and collision rates in the wireless network as compared to a wired network are mitigated by the MAC-layer itself. In particular, the MAC protocol mandates acknowledgments of received frames by the receiver and retransmissions of frames by the sender that are not acknowledged within a time interval. The timer values are such that the higher layer transport protocols, in particular TCP, do not get affected by the loss of a frame once in a while. Hence, studying phenomena related to the connection-oriented nature of TCP does not lead to new insights into network problems that are related to the specific use of an 802.11 network.

This paper is organized as follows. First, in Section II, we outline the laboratory setup for all conducted experiments. Each section from III through VI consists of the description of a single experiment, the results obtained from the experiment and a detailed explanation. The first experiment, in Section III, studies the payload throughput of an 802.11b network as a function of the frame payload size. The second experiment, in Section IV, studies the throughput with multiple senders. The third experiment, in Section V, is aimed to find the maximum number of simultaneous VoIP connections in a single cell. In Section VI, we study the throughput of such a network in the case of converged networking, i.e., simultaneous voice and data traffic. Finally, Section VII discusses the ramifications of the experiments and Section VIII concludes the paper.

II. Experimental Setup

Eight clients, all PCs (some Laptops, some Desktops) running Windows 2000, are associated with a single access point. The access point is connected to an IEEE 802.3 (Ethernet) LAN. On the wired side of the network, we are also using PCs running Windows 2000 which serve as endpoints for VoIP or UDP data connections. All endpoints in the wireless network are in one subnet whereas the nodes in the wired network are on another. Both subnets are connected through a single router and are thus one hop apart from each other.

The IEEE 802.11b access point and all clients are situated in the same room with no physical obstacles between them. Hence, the probability of frame loss due to weak signal strength and/or presence of hidden stations is negligible. The experiments were conducted with access points from different vendors, with no significant differences in the results. The radio cards used in all clients were from the same vendor.

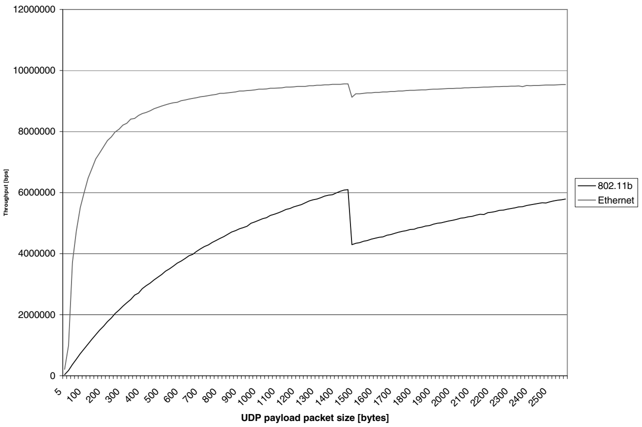

Fig. 1. Throughput as a function of the packet size for Ethernet and 802.11b. The drop in throughput for packet payloads larger than 1472 bytes is due to fragmentation.

III. Throughput for a Single UDP Sender

A. Experiment and Results

We measured the available throughput in terms of sent payload on Layer 4 with IP/UDP as bearer in the scenario of a single client.

We used a program constantly sending out UDP data frames to the wireless medium (and, hence, in some sense "flooding" the wireless network). The stream was sent to an endpoint in the wired network. Our results indicate that the maximal achievable payload data rate in this scenario is approximately 6.1 Mbps which is achieved when the payload of each UDP packet is chosen to be 1472 bytes. In this case, the maximal payload of an Ethernet frame (1500 bytes minus 20+8 bytes for the IP and the UDP header) is used. Payloads larger than 1472 bytes get fragmented into more than one frame which reduces the observed throughput significantly. When the 62 bytes of the IEEE 802.11 frame body and the IP/UDP headers (34+20+8) in each sent frame are accounted for, the overall throughput of the wireless network is approximately 6.36 Mbps.

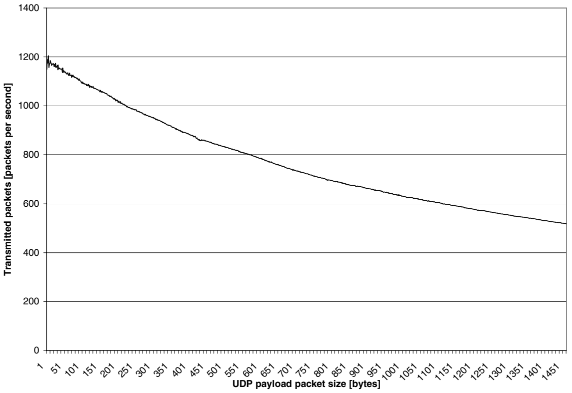

Fig. 2. Transmitted packets per second as a function of the packet size in IEEE 802.11b

Figure 1 shows our measurements of payload throughput as a function of the payload size in a single client scenario for 10 Mbps Ethernet and 802.11b. Figure 2 shows the number of transmitted frames per second as a function of the packet size.

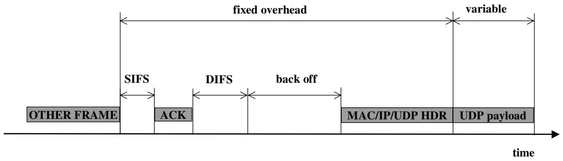

Fig. 3. IEEE 802.11 CSMA/CA medium access scheme.

Based on the number of frames sent per second and the fact that in our setup the payload of each frame is transmitted at 11 Mbps, we can compute the fixed overhead for the transmission of each frame. For our measurements, the overhead evaluates to an average of around 840μs per frame. It should be noted that the transmission of a payload of 1472 bytes at 11 Mbps takes about 1070μs, hence even in this case the fixed overhead time is quite comparable to the actual transmission time for the payload.

B. Explanation

Clearly, the question arises why the fixed overhead per frame is so significant that it takes almost half of the total transmission time even when using the largest possible payload size of 1472 bytes. The overhead stems from the Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA) medium access scheme according to the Distributed Coordination Function (DCF) of the 802.11 standard [9] which we will briefly describe here. A more detailed description can be found in [8].

As collisions in the wireless medium cannot be detected, the MAC protocol is designed to prevent collisions from occurring. Furthermore, due to the possibility of undetected collisions and the higher number of transmission errors as compared to wired networks (e.g., by interference), unicast frames are acknowledged by the receiving station. The acknowledgment is sent out after the transmission has finished and a certain duration of time called short inter frame spacing (SIFS) has elapsed. If a node wants to transmit a frame and senses the medium idle for a certain duration of time called distributed coordination function inter frame spacing (DIFS), it may start transmitting. As DIFS is longer than SIFS, it is made sure that a well received frame can always be acknowledged for before the next frame is transmitted.

If a node wants to start transmitting while the medium is busy or if it wants to transmit another frame after just finishing a transmission, it also waits for the medium to be idle for the DIFS period. Then, the node does not begin to transmit immediately but enters a contention phase for the medium. Contention is done by choosing an integer random back-off between 0 and a parameter CW (CW stands for contention window size) which is initially set to a value CWmin. The probability distribution among these values is uniform. The random back-off determines the number of time slots the client defers its transmission in addition to the DIFS time.

If the medium is sensed idle in such a "slot", the back-off timer is decreased by one. If the random back-off has decreased to 0, the node starts transmitting. If another node starts transmitting before this happens, the node continues to count down the back-off timer after the medium has been sensed idle for the DIFS period. Thus, if multiple clients want to transmit a

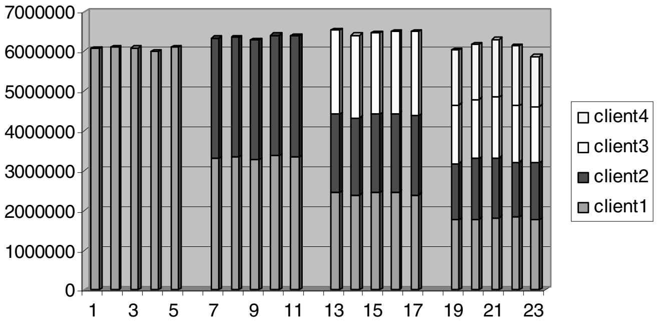

Fig. 4. Throughput for multiple UDP senders sending packets of 1472 bytes. Y-Axis shows throughput in bps.

frame, the one with the lowest random back-off time will win the contention for the medium. However, if more than one node happens to choose the same back-off time, they will start to transmit at the same time and a collision will occur. The clients assume that the frame was lost if an acknowledgment is not received within SIFS. In case of an unsuccessful transmission, the CW value doubles until a CWmax value is reached. The CW parameter is reset to the CWmin after each successful (i.e. acknowledged) transmission.

The IEEE 802.11/802.11b standard defines SIFS to be 10μs. A slot time is 20μs and the value of DIFS is defined to be the value of SIFS plus two slot times which is 50μs. The size of an acknowledgment frame is 14 bytes which take about $10\mu s$ to transmit at 11 Mbps. However, each transmitted frame also needs some physical layer overhead (PLCP header of 48μs and a preamble of 144μs) which is about 192μs. Thus, the total time to transmit an acknowledgment is 203μs. The IEEE 802.11b standard [9] defines CWmin to be 31. Therefore, in the scenario of a single client constantly transmitting, the average random back-off time is 15.5 slots which equals 310μs. For the actual data frame we have an overhead of 34 bytes for the 802.11 MAC header, 20 bytes of IP header and 8 bytes of UDP header totaling 62 bytes which take about 45μs to transmit at 11 Mbps. Together with the 192μs physical layer overhead this amounts to 237μs. Summing up these values, the fixed overhead per frame as illustrated in Figure 3 can be calculated as $10 + 203 + 50 + 310 + 237 = 810\mu$s.

Indeed, this value matches well with the overhead measured in the experiment. The experimental values are a little bit higher due to the fact that the overhead of periodic beacons sent out by the access point is not included in the calculated value. Such a beacon contains management information about the network and is sent out about every 100 ms. When the effect of the beacons is discounted, the measured value matches with the calculated value.

IV. Throughput for Multiple UDP Senders

A. Experiment and Results

In this experiment, the program constantly sending out UDP data frames is run simultaneously on multiple wireless stations connected to the same Access Point. All UDP streams were sent to endpoints in the wired network.

When running two datagram senders, the aggregated maximal achievable payload data rate increased from approximately 6.1 Mbps in a single sender scenario to approximately 6.4 Mbps for a payload of 1472 bytes each UDP packet in both streams. Moreover, both senders equally share the available bandwidth. Each client is able to send around 3.2 Mbps.

For three senders, the maximal achievable payload data rate further increases to 6.5 Mbps which is also fairly shared among all clients. For four senders, the maximal achievable payload data rate decreases to around 6.1 Mbps.

Assuming all senders transmit with 11Mbps, the fixed transmission overhead per frame can be calculated to be 750μs, 720μs, and 830μs for two, three and four senders, respectively.

B. Explanation

As mentioned in the explanation of the previous experiment, the average back-off idle time between successive frames on the channel is exactly equal to the average back-off between successive frames of the station.

If more than one station is transmitting, there are two main factors, which determine the maximum channel throughput. First is the average idle time on the channel due to inter-frame back-offs performed by each station. Second is the probability that a transmission on the channel results in a collision. Obviously, a lower average idle time increases the channel throughput while a higher probability of collision reduces the channel throughput.

Let us first study the case of two senders always contending for the medium. In previous work on throughput analysis of 802.11 networks, the average actual back-off has been evaluated both via analytical approximations as well as via simulations [1], [2], [3], [7]. From these, the average back-off window size is determined to be 8.5 slots. Therefore, when calculating the fixed overhead, we obtain a value of 670μs.

However, there is a chance that transmissions of the two senders collide and in turn the time of the transmission is wasted. The collision probability has been shown to be 0.03. These failed transmissions account for “wasted” air-time that has to be taken into account when comparing the experimentally observed value. It takes 1070μs to transmit a 1472 byte long frame plus a fixed overhead of 670μs. Thus, the total time for such a frame transmission is 1740μs. Crediting this time to all successfully transmitted frames, we obtain a correction factor of $0.03 \cdot 1740 = 52\mu$s.

Hence, we would expect to observe a fixed overhead of 722μs per frame. This value is very close to the experimentally obtained overhead of 750μs. Again, the deviation between calculated and observed value is due to the transmission of beacons. Furthermore, the time accounted for a collided frame is too optimistic since the MAC scheme requires a longer idle period (Extended Interframe Spacing, EIFS) after a collision.

For three senders, the average back-off window size is determined to be 6.215 slots and the collision probability is 0.056. Thus, the expected fixed overhead per transmission is 719μs. The actual value we measured is 741μs.

In the case of four senders, the average back-off window size is determined to be 5.061 slots and the collision probability is 0.0774. Thus, we would expect a fixed overhead per transmission of 730μs and hence an aggregated throughput of 6.54 Mbps. However, the measurements revealed a much lower throughput.

The reason was that the peak transmission rate of one or more wireless stations was not 11Mbps, but dropped down occasionally to 5.5Mbps. The stations misinterpret the unusually high collision rate with four active stations as interference. In turn, they decrease the transmission speed to enable better interference robustness. Since interference detection and adaptive transmission rate control is not part of the standard and is implemented in a proprietary, undisclosed manner, it is not possible to calculate the fixed overhead.

V. Maximal Number of VoIP Connections using ITU G711 a-Law Codec

A. *Experiment and Results*

The same set-up was used to explore the maximal number of VoIP connections possible in a single cell of an IEEE 802.11b network. One end-point of each VoIP call is a wireless client, while the other end-point is on the wired network. Given the short range of an 802.11 access point, "wireless-to-wireless" calls do not seem to be a typical scenario.

For each call, we used the ITU G711 a-Law codec where frames are sent out every 10ms. The codec output is sent over the IP, UDP and RTP protocols. It should be noted that it is possible to use this codec with varying audio data length per packet sent. Each millisecond of audio is encoded into eight bytes of data. In addition to the IP/UDP overhead each packet contains an additional fixed RTP layer overhead of twelve bytes. Hence, when sending audio data packets every 10ms, the payload of such a packet is 92 bytes. When making a VoIP phone call, two such data streams are established, one in each direction.

We tested the number of VoIP connections with acceptable voice quality by successively establishing new calls in addition to ongoing calls. The quality of the connections was monitored through measurements of loss, jitter and round-trip time by a commercially available tool. For the first five calls, the quality of all connections was fine. Loss (0%), round-trip time (around 5ms) and jitter (around 7ms) were all in ranges far below critical although higher than experienced in a single-hop wired Ethernet. When placing the sixth call, except for a sporadic increase in the round-trip time for some of the connections as measured by the commercial tool, the quality of the six simultaneous connections was still fine. Placing the seventh connection lead to unacceptable loss in all "wired to wireless" streams. Except for the round-trip time, all "wireless to wired" streams still were fine.

Thus, our experiments indicate that the maximal number of simultaneous VoIP connections in a single cell of an IEEE 802.11b network when using the G711 a-Law codec with 10ms audio data packets is six.

B. *Explanation*

The aggregated bandwidth of seven VoIP connections as specified above is approximately 1 Mbps. It can be argued that the bandwidth available when the network is exclusively used for VoIP traffic is equivalent to the bandwidth available in the case of two UDP senders flooding the network [5]. As the bandwidth in this scenario is lower than the bandwidth needed for seven VoIP connections, the unacceptable call quality is caused by overload.

When the seventh connection is started, two additional VoIP-streams are created, pushing the load offered by the nodes to the network beyond its throughput limit. Thus, some of the audio data traffic cannot be sent through the wireless network. Whereas each wireless node only sends a single VoIP stream, namely the one going from it to wired, the access point has to send out all audio data traffic coming from wired in the wireless network. Therefore, it is actually sending 50% of the overall load on the network before the network is saturated. As the CSMA/CD protocol is designed to be fair, in a situation where more load is offered than can be accommodated, the node using the largest number of transmission slots will be curtailed first. Thus, in our scenario the access point is the only node unable to put its load onto the network.

The paper [5] presents a detailed analysis on the number of VoIP connections possible in an 802.11b network for various codecs and audio data payload settings. Here, we would only like to recall the results with respect to the maximal number of connections possible when all stations transmit using a data rate of 11Mbps.

Audio (ms)	G711	G729	G723
10	6	7	
20	12	14	
30	17	21	21
40	21	28	
50	25	34	
60	28	41	42
70	31	47	
80	34	54	
90	36	60	61
100	39	66	

As this table shows, the analytical limit of six connection matches with our experimental findings.

VI. UDP Throughput in the Presence of VoIP Traffic

A. *Experiment and Results*

Now we consider the case of converged networking, i.e. the simultaneous presence of voice and data traffic in the wireless network. Again, our focus is on throughput. One station is running the UDP datagram sender, whereas other stations connected to the AP run VoIP connections using the same parameters as in the previous experiment. Again, all connections terminate in the wired network. The following table shows the throughput of the UDP sender as a function of the number of ongoing VoIP connections.

VoIP conn.	UDP Throughput
0	6.06 Mbps
1	5.15 Mbps
2	4.26 Mbps
3	3.28 Mbps

As these values show, each VoIP connection diminishes the throughput of the UDP sender by approximately 900 Kbps even though the VoIP connection only consists of two UDP streams

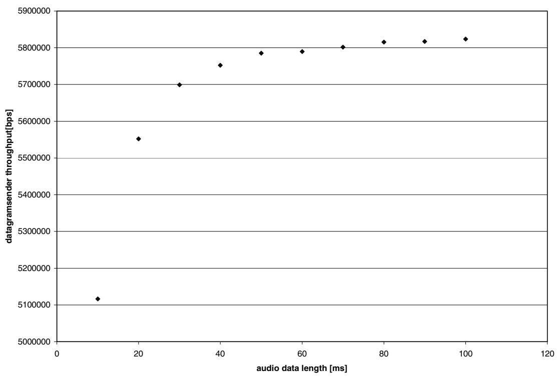

Fig. 5. Throughput of datagramsender when one VoIP connection is active.

each of which has a bandwidth of 74 Kbps (64 Kbps voice data plus 10 Kbps RTP overhead). Hence, the aggregated throughput of the network reduces by 760 Kbps for each VoIP connection.

B. Audio Data Size Variation

In a successive experiment, we ran a single VoIP connection using the G711 codec and varying audio data size as well as a datagram sender. We measured the payload throughput of the datagram sender. The results are displayed in Figure 5.

The throughput of the UDP sender increases steeply with an increase of the audio data payload size from 10ms (92 bytes per packet) to 20ms (172 bytes per packet). Then, the returns in terms of bandwidth for using larger audio data payload sizes diminish with increasing audio data payload size.

C. Explanations

When the wireless network is shared between voice and data, the remaining bandwidth for data when running a VoIP-Connection with a small payload size is lower than when running a data connection of the same bandwidth using full size frames.

Denote the maximal throughput possible for a packet payload size of k by R^k_{eff}. For a VoIP stream with payload size k and bandwidth b, the bandwidth equivalent b^{1472}_{eq} is defined by

$$b^{1472}_{eq} = \frac{b}{R^k_{eff}} \cdot R^{1472}_{eff}.$$

Intuitively, bandwidth equivalent is the bandwidth that would become available for data traffic if the VoIP stream was shut down. Figure 6 shows the bandwidth equivalent of a VoIP-Connection for varying audio data packet payloads. The values used for R^k_{eff} and R^{1472}_{eff} are taken from the results of the first experiment. As this figure shows, a single G711 a-Law VoIP-stream sending a packet every 10ms uses the network like a data connection sending UDP packets with 1472 bytes with 560 Kbps. Thus, when running one VoIP-Call with these parameters, the remaining bandwidth for data traffic in the wireless network is only 4.9 Mbps as each of the two VoIP data streams consumes bandwidth equivalent to 560 Kbps of data traffic.

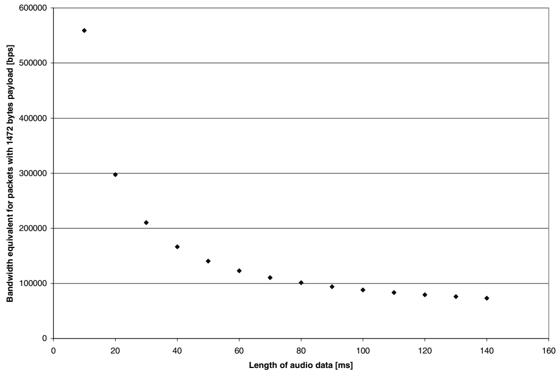

Fig. 6. Equivalent Throughput of a G711 a-Law stream when assuming that same transmission time is used to send UDP traffic with 1472 bytes payload instead.

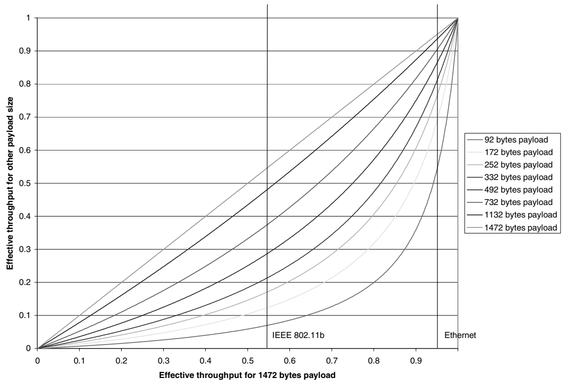

Fig. 7. Effective throughput of a network for 92, 172, 252, 332, 492, 732, 1132 and 1472 bytes (packet sizes used by a G711 a-Law codec with audio data length 10ms, 20ms, 30ms, 40ms, 60ms, 90ms, 140ms, respectively) of payload as a function of the effective throughput for 1472 bytes of payload. The vertical lines mark the values for IEEE 802.11b and Ethernet.

VII. Discussion

A. Use of 802.11b for VoIP

VoIP call quality in IEEE 802.11 networks, as determined by the loss, delay and jitter characteristics of the call, is fine as long as the network throughput limit is not exceeded. In other words, the throughput of the wireless network is the key to determining the number of VoIP connections which may be run simultaneously in a single cell. In contrast with Ethernet, the smaller packet size of VoIP traffic as opposed to data traffic has a dramatic impact on the effective throughput in a wireless network.

Figure 7 shows the effective throughput of a network for varying payload packet sizes as a function of the effective throughput for 1472 bytes of payload under the assumptions that the overhead for all different packet payload sizes is fixed and that the variable transmission time for the packet payload is linear.

Let the effective throughput of a network at payloads of 1472 bytes be denoted by R^{1472}_{eff}. The overhead of the transmission, denoted by T_{oh}, measured in time units needed to transmit one byte, is given by

$$T_{\text{oh}}(R^{1472}_{\text{eff}}) = 1472 \cdot \frac{1 - R^{1472}_{\text{eff}}}{R^{1472}_{\text{eff}}}.$$

Consequently, the effective throughput available for packets with a payload size of k bytes, R^k_{eff}, as function of the effective throughput for packets with payload size 1472 is given by

$$R^k_{\text{eff}}(R^{1472}_{\text{eff}}) = \frac{k}{k + T_{\text{oh}}} = \frac{k}{k + 1472 \cdot \frac{1-R^{1472}_{\text{eff}}}{R^{1472}_{\text{eff}}}}.$$

As shown in the figure, the efficiency of Ethernet which is around 95% for a payload packet size of 1472 bytes still allows for an efficiency of around 55% when the payload packet size is only 92 bytes.

There are two ways of increasing the efficiency of such a network for the VoIP-streams we consider here, namely to either to increase the overall efficiency of the network by reducing the fixed overhead per sent payload or to switch to transmitting larger payload packets.

Let us discuss the first of these alternatives. Whereas standard bodies and committees are in general trying to make their proposals as efficient as possible, fixed overhead per transmitted frame is especially crucial for smaller packet sizes which are common for real-time traffic. However, as the picture shows, the lower the efficiency of the network for payloads of 1472 bytes, the less impact has an increased efficiency on the throughput achievable for lower packet payload sizes. Adding to that, a small packet payload size also decreases the improvement in efficiency for such measures. In fact, if the fixed overhead of IEEE 802.11 could be cut from 840μs to 420 μsec, this would increase the network efficiency for packet payloads of 1472 bytes from 55% to 72% but the efficiency for packet payloads of 92 bytes would only increase from 7% to 14%. Hence, although decreasing the current overhead in IEEE 802.11b in general would significantly help to increase the available throughput for data traffic, it would not be as helpful for VoIP traffic.

So, increasing the audio data length per transmitted packet appears to be a better solution to the problem. Unfortunately, there is a tradeoff between the subjective audio quality of the call and the amount of audio data sent in a single packet. A higher audio data length per packet increases the time gap between the recording of the audio data on one side of the stream and its playback on the other side as the data has to be collected before it is sent. Moreover, loosing a single audio packet with a large amount of audio data is an event that will likely be heard by the receiving party. The right solution therefore might be one that exploits different optimal payload sizes for the wireless and the wired part of the network.

B. Capacity Assessment for Wireless Network

As our experiments show, the maximal aggregated bandwidth available in an 802.11b network is highly dependent on the traffic in the network and can range anywhere from a few hundred Kbps up to 6.1 Mbps even if only a data rate of 11 Mbps is used. Adding to that, stations may also transmit at data rates of 1, 2 or 5.5 Mbps. Hence, the simple questions "can the network handle a flow with 2 Mbps bandwidth" or "can the network accommodate 800 packets per second" cannot be answered when only looking at bandwidth.

Thus, admission control for 802.11b networks cannot be based on bandwidth. However, taking into account the low number of VoIP connections possible, the need for VoIP admission control in wireless networks is apparent. Placing an additional call or an additional data connection that exceeds the capacity of the wireless network will result in unacceptable call quality for all ongoing VoIP calls. Further, if the load offered to the network is higher than its capacity, the DCF medium access scheme of 802.11 curtails the client with the highest load first. In most cases, the access point of the wireless cell puts more traffic on the air than the associated stations. Hence, it gets curtailed first which leads to unacceptable packet loss for all VoIP streams transmitted from the access point to a client resulting in bad call quality for all connections.

Therefore, in [6], an admission control metric was proposed that is based on the definition of bandwidth equivalence as outlined here. This metric allows to accurately assess the capacity of an 802.11b network by calculating the bandwidth equivalence for all flows which allows for an accurate picture of the current and future channel usage.

VIII. Conclusions

In this paper, we studied the behavior of UDP and VoIP over 802.11 networks, from the perspective of number of connections that a single access point can support. An important conclusion is that given the choice of payload sizes in IP phones, 802.11b base-stations prove to be inadequate in handling large number of VoIP calls. In fact, the inherent channel inefficiency of 802.11b, at smaller frames sizes, limits the maximum number of VoIP calls to a very low number, six in the case of G711a-law codec with 10ms of audio data per packet. Furthermore, our experiments revealed that the aggregated bandwidth of the wireless network is diminished by ongoing VoIP connections. Each such G711 VoIP connection reduces the bandwidth available for data traffic by approximately 900 Kbps. In turn, the maximal bandwidth available for other traffic when three of such connections are active is only 3.3 Mbps.

References

[1] G. Anastasi and L. Lenzini. QoS provided by the IEEE 802.11 wireless LAN to advanced data applications: a simulation study, In *Wireless Networks*, Vol. 6, pp 99–108, J.C. Baltzer AG, Science Publishers, 2000.

[2] G. Bianchi. Performance Evaluation of the IEEE 802.11 Distributed Coordination Function. In *IEEE Journal on Selected Areas in Communication*, Vol. 18, No. 3, March 2000, pp. 535–547.

[3] H. S. Chayya and S. Gupta. Performance modeling of asynchronous data transfer methods of IEEE 802.11 MAC protocol. In *Wireless Networks* Vol 3, 1997, pp. 217-234.

[4] S. Garg, M. Kappes and M. Mani, *Wireless Access Server for Quality of Service and Location Based Access Control in 802.11 Networks*, Proceedings of the Seventh IEEE Symposium on Computers and Communications (ISCC), Taormina, Italy, 2002.

[5] S. Garg and M. Kappes. *Can I add a VoIP call?*, submitted.

[6] S. Garg and M. Kappes. *A New Admission Control Metric for VoIP Traffic in 802.11 Networks*, submitted.

[7] A. Heindl and R. German. Performance modeling of IEEE 802.11 wireless LANs with stochastic Petri nets. *Performance Evaluation*, 44 (2001), 139-164.

[8] J. Schiller. *Mobile Communications*, Addison Wesley, New York, 2000.

[9] IEEE 802.11, 11a, 11b standard for wireless Local Area Networks. *http://standards.ieee.org/getieee802/802.11.html*

[10] *http://grouper.ieee.org/groups/802/11/Reports/tgg_update.htm*