IEEE 802.11b voice quality assessment using cross-layer information

Ian Marsh, CNA, SICS
Sweden
Email: ianm@sics.se
Juan Carlos Martín Severiano, and Victor Yuri Diogo Nunes, and Gerald Q. Maguire Jr.
School of Information and Communication Technology, KTH
Sweden
Email: juancarlosmartin@gmail.com, ynunes@gmail.com, maguire@kth.se

***Abstract*— This paper reports on the suitability of IEEE 802.11b networks for carrying real-time voice traffic, considering particularly the end terminals. More specifically we looked at such networks in different operating circumstances: an outdoor environment, an office environment, and the influence of competing traffic. Additionally we have investigated the link protocol in combination with the application layer. Based on over 2500 recorded sessions, it can be generally concluded that the 802.11b protocol can support real-time voice; particularly if the link transmission rate is immediately lowered after an unsuccessful initial transmission. However, we did find situations where the voice quality deteriorated below commonly accepted values, such as when competing with high-rate TCP traffic, when intervening obstacles blocked the transmission path, and with certain uses of the RTS/CTS mechanism.**

I. Introduction

IEEE 802.11b networks are being used in public hotspots, along with office and home networks. The resulting broadband wireless local area network (WLAN) has brought IP-based telephony into serious competition with the cellular telephony infrastructure. The goal of this paper is to assess the suitability of such networks based on extensive measurements. Our focus was the *application* perceived quality in different usage scenarios due to the environment's effect on voice quality. By environment we mean the physical context: the separation of the nodes and the intervening obstacles. We focus on the quality variations of a single voice over IP (VoIP) call in various circumstances: in ad-hoc mode with clear line of sight, inside and outside an office environment, in the presence of competing TCP and UDP traffic, and when using IEEE 802.11 infrastructure support (with an access point). Where possible, we also examine the contribution of the MAC layer, specifically retransmissions and the RTS/CTS mechanism to the application layer quality. The remainder of this paper is structured as follows: the next section explains our basic experimental setup, and how one can use information from two layers to obtain data about the overall VoIP quality. Section III presents some of the related work in this area, Section IV gives our findings in the four measurement settings. Finally we conclude with a summary of our findings and highlight issues for WLAN real-time voice users together with some suggestions for future work.

II. Basic experimental configuration and cross-layer measurement

The basic configuration used for our experiments comprises one node that sends a unidirectional flow of RTP VoIP packets to a second node that acts as a receiver. The basic configuration we used is shown in the illustration of figure II. In order to observe and capture over the air traffic, we used two monitoring devices placed close to the sender station. One monitor was in the same sending station and one was physically separate[1], "sniffer 2" in figure II. Capturing all the frames in the air is non-trivial, and our decision to use two capture devices and merge the traffic they observe was motivated by the experience and pitfalls reported by Yeo, et al. [YBA02].

[1]Experiments showed the probability of both monitors losing a frame was 0.04%.

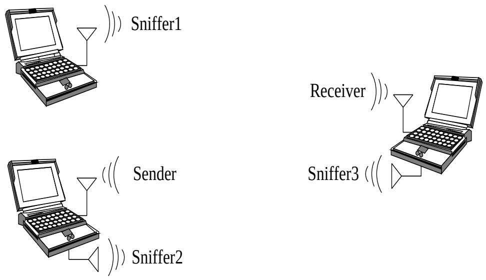

Fig. 1. VoIP measurement testbed

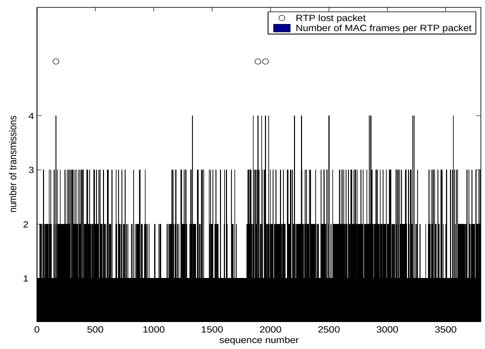

Fig. 2. Typical data link layer behavior

We also captured the traffic sent back from the receiver using a third monitor, sniffer3; however this should be mostly ACK frames. Ethereal and Sphone, an in-house VoIP tool, were used to capture the link layer frames and VoIP packets respectively. During the experiments, the RTP traffic corresponded to a simulated call of 80 seconds generating 160 byte packets 20ms apart, no silence suppression or signaling were used. As an example of combining cross-layer information is shown in Figure II and shows the retransmission pattern of $\sim$4000 packets (i.e., a call, 1/20ms * 80 seconds) when observed at both the data link layer and the application layer for a single VoIP flow. Vertical lines up to $y = 3$ show the number of transmissions that were received. In some cases the 4^{th} transmission (i.e., $y = 4$) were not received and are indicated in the plot as circles. The default maximum number of transmission attempts by the MAC layer was four with the hardware we used, after which the frame is silently discarded. Note that only by using the application layer information at the receiver was it possible to establish whether transmission was actually received. This is because the ACK may not always be received and failed link transmissions are not immediately indicated by the drivers to the higher layers.

III. Related work

Anastasi et al. measured the performance of IEEE 802.11b ad hoc networks [GAEB04], specifically the range of the end-terminals, the impact of different data rates and their variability. They observed that the transmission range was highly dependent on the data rate up to 100m, whilst the physical carrier sensing range was independent of rate up to 200m. Unlike their results in ad hoc mode, we didn't observe different ranges for different rates at up to 320 meters. Even at 400 meters there was no conclusive data rate dependency on range.

Hertrich looked at mixed traffic (including real-time voice) in IEEE 802.11 networks [Her03]. He used a MAC booster and by tailoring it could alter the number of retransmissions for different positions to achieve the required throughput. We did not try to change the number of transmissions. This work is similar to ours in that he considered the environment as important, however he used VoIP and MPEG4, while we used VoIP and TCP. Additionally, Hertrich focused on the home, whereas we focused on an office environment. Nevertheless we found that certain positions also did not permit any communication to take place at all.

Dimitrou et al. address issues that can make the deployment of multimedia communications difficult in 802.11 networks. [DS03]. They cite interference and users moving out of range as limiting factors for good VoIP quality in WLANs. They suggest the use of smart speech coding (including an enhanced version of the G.711 coding developed by their company) to make the speech more resilient to loss.

Hoene et al. examined the effect of motion on the performance of wireless links through a series of experiments with moving nodes [HGW03]. They conclude that other factors such as modulation type, quality of power supply, environmental setup, and number of retransmissions may have greater impact on 802.11b performance than the motion itself. In general the greater the speed of the terminal the lower the correlation of loss events. In our experiments the nodes were not moving, i.e., movement only occurred between measurements; thus movement should only decrease the observed losses.

IV. Results

We will now present four distinct series of experiments regarding VoIP quality: (A) outdoor measurements, (B) measurements in an office, (C) the effect

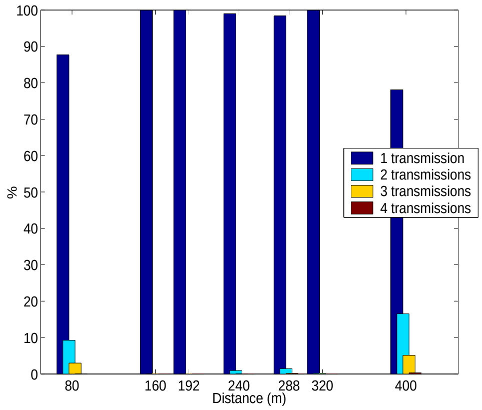

Fig. 3. Number of successful MAC transmissions vs distance

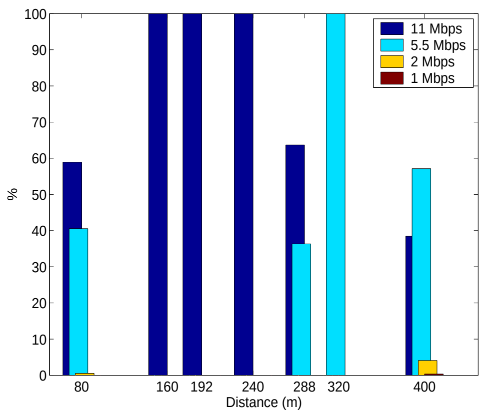

Fig. 4. IEEE 802.11b rates used vs. distance

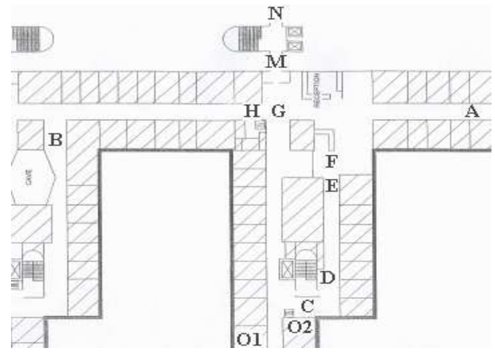

Fig. 5. The office floor-plan for the indoor experiments

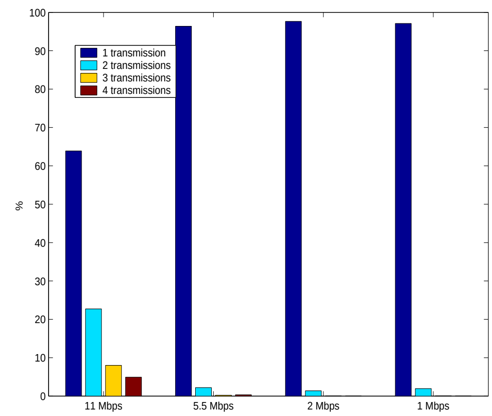

Fig. 6. The bitrate histograms for the indoor experiments

of competing traffic, and (D) usage of the RTS/CTS mechanism.

A. *Quality as a function of distance*

The first measurement series we conducted were designed to examine the effect of the distance between the sender and the receiver when using ad hoc mode. The terminals were within line of sight in an outdoor environment. Figure 3 shows the averages of transmitting a single VoIP flow eight times at each distance. It is a histogram of the percentage of MAC frames successfully delivered at the seven different distances.

Figure 4 shows the rates at which each of these frames were transmitted. The first observation is that the majority of transmissions were successful at the first attempt. This is particularly true for the middle distances in our measurements. Overall the loss percentages were either zero or very low ($< 0.1\%$), even considering the retransmissions at 80 and 400 meters were 0.025% at 80 meters and 0.05% at 400 meters. The loss and jitter figures can be found in Juan Carlos Martín Severiano's master thesis, along with further details and analyses [Sev04]. It is evident from these figures that interference from nearby WLANs can induce losses even at relatively short distances. One effect is the lower rates as easily seen in figure 4. No continual competing traffic was observed on the channel during these measurements, hence the delay and jitter values are low ($<$ 7ms for both, when added). Interference on the same channel was mainly the result of IEEE 802.11 beacon frames and probe requests from nearby various WLANs.

B. *Office environment measurements*

Next we measured the number of MAC layer transmissions inside a typical modern office environment. The purpose of this scenario was to measure the effect of the walls, windows, and intervening obstacles typically found in modern offices on the voice traffic. The physical environment is shown in figure 7 and the office floor-plan in figure 6. The letters on the floor-plan reflect the positions of the sender (O_1 and A), the receivers (the other

Fig. 7. Office environment with three different receiver orientations

letters) and orientations of the receiver (indicated as subscripts). The placements of the receiver were chosen arbitrarily to represent challenging locations for wireless communication, for example a computer room (containing servers) is located between O_1 and E, F. With the sender located at A and the receiver located at B, we exam at the impact of the office environment on quality.

Figure 6 shows the mean of four separate series of measurements where the transmission rate was fixed at each of the defined rates for an 802.11b interface. Unlike the example presented in section IV-A where the rate could vary, we observed that using a fixed sending rate of 11Mbits/sec led to many more retransmissions and higher losses, approximately 2.75%. By reducing the rate, the probability of a successful first transmission increases, as can be seen from the 5.5Mbits/sec values. A clear conclusion is that for voice traffic the rate should be immediately reduced (for the retransmission), rather than attempting retransmissions at a higher rate. This may not only alleviate loss, but reduce delay and jitter for this frame[2]. The loss figures for the 5.5, 2, and 1 Mbits/second rates were approximately 1%.

From table I we can see that the environment significantly affects the quality, particularly in the large loss values; in some situations communication was not even possible (F_2). Different locations and antenna orientations around the office were selected as shown in figures 6 and 7 respectively. From the table the quality differences were significantly affected by the antenna orientation. For each of the orientations there is 90 degree difference in the X, Y, and Z planes. In figure 7 the middle picture show when the antennas in parallel. More details of application layer measurements, particularly within the office environment can be found in [Nun04].

[2]This confirms a hypothesis by one of the authors (GQMjr) presented to Andrzej Duda following his talk at KTH on 2003.05.08.

Pos.	Loss (%)	RTT (ms)	Jitter (ms)
O_2	[0, 0.0, 0.1]	[1.9, 2.2, 2.4]	[0.2, 0.2, 0.3]
C	[0, 0, 0]	[1.9, 2.1, 4.0]	[0.1, 0.1, 0.2]
D	[0, 0, 0]	[2.1, 2.6, 3.1]	[0.1, 0.7, 1.0]
E_1	[0, 0.2, 2.6]	[2.8, 3.2, 5.4]	[1.0, 1.2, 2.0]
E_2	[9.4, 36.3, 89.1]	[5.3, 12.2, 24.3]	[4.2, 6.3, 16.1]
E_3	[0, 0.0, 0.2]	[2.8, 2.8, 4.0]	[0.9, 1.1, 1.5]
F_1	[20.1, 54.9, 88.7]	[5.4, 13.4, 24.6]	[0.9, 14.2, 57.9]
F_2	No signal	No signal	No signal
F_3	[1.8, 22.8, 84.9]	[4.6, 11.7, 13.7]	[2.9, 6.7, 37.4]
G	[0, 0.2, 1.9]	[1.9, 2.2, 3.8]	[0.1, 0.4, 1.0]
H_1	[0.4, 5.4, 30.2]	[3.5, 3.9, 8.1]	[1.5, 2.2, 4.2]
H_2	[3.4, 11.0, 28.3]	[5.9, 6.1, 11.7]	[3.4, 3.9, 4.6]
H_3	[0, 0.2, 2.7]	[3.4, 3.4, 5.4]	[1.1, 1.2, 1.8]

TABLE I

VoIP METRICS IN AN OFFICE ENVIRONMENT ([MIN,MEAN, MAX])

C. Competing UDP and TCP traffic in ad hoc mode

We now look at the effect of competing traffic on the quality of VoIP sessions. We first consider the case of ad hoc mode, i.e. without an access point. The two VoIP nodes (and monitors) were in the same room with up to four nodes generating UDP or TCP background traffic. The TCP and UDP packets were 1500 bytes in total, produced by the NTTCP traffic generator. Each node was responsible for generating a single stream. The goal was to observe the MAC protocol's behavior by measuring the delay, jitter, and loss caused by the failure of 802.11's collision avoidance mechanism under increasing load.

If stations select identical slot numbers and hence send simultaneously, so called "collisions" will occur resulting in lost frames. Usually the packet capture nodes cannot detect collisions so it is the responsibility of the application layer to infer lost packets due to heavy load on the medium. Usage of UDP was intended as a controlled traffic source, and TCP as a more representative, but more complex, traffic source. Figures 8 and 9 show the round trip time and jitter for the configuration described. This RTT is calculated from application layer RTCP information sent once per second, hence the high variance.

Zero nodes indicates the case without competing traffic, i.e.,solely the VoIP flow. The delay and jitter values are much higher compared to the values shown earlier. Since delay should be less than 150ms (an ITU G.114 recommendation) for good interactive communication, a significant proportion of the delay can be used up gaining channel access.

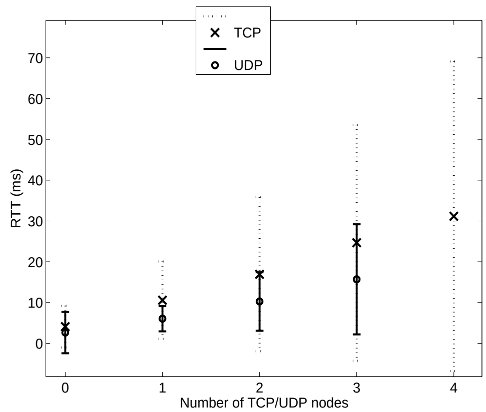

Fig. 8. The round trip time for 0-4 competing nodes in ad hoc mode

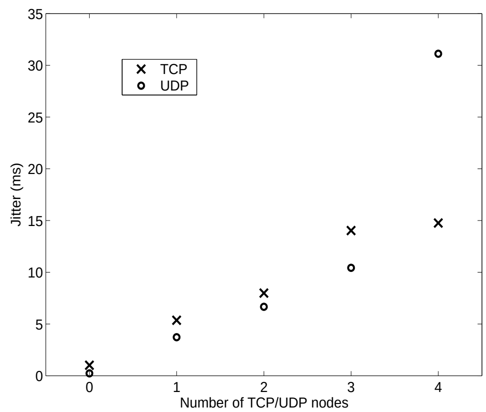

Fig. 9. The jitter for 0-4 competing nodes in ad hoc mode

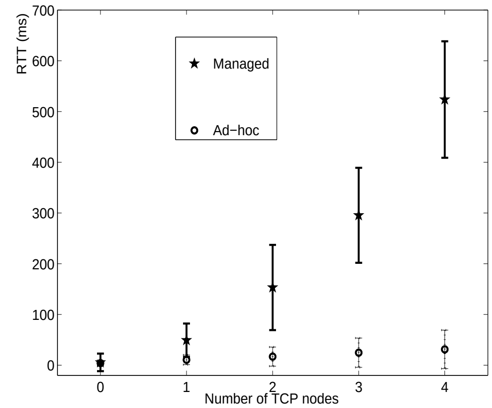

Fig. 10. The round trip time for 0-4 competing nodes

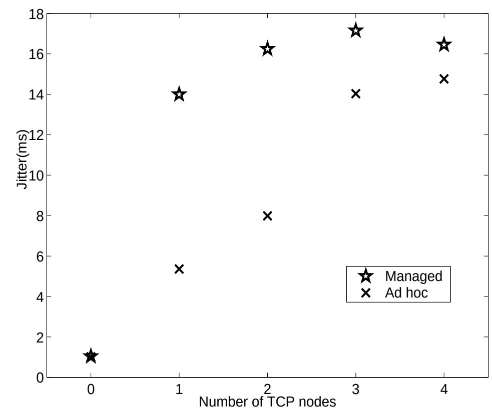

Fig. 11. The jitter for 0-4 competing nodes

Note also the standard deviation of the TCP traffic is much higher than the UDP, whereas the mean is only slightly higher. This variance is problematic for the jitter of VoIP sessions (leading to either loss or further delay due to the buffer playout algorithms).

D. Competing TCP traffic in infrastructure mode

Next we examined VoIP quality when using an access point between the communicating nodes, as this is likely to be a common scenario for deployed 802.11 wireless (voice) access to the Internet. We only considered TCP traffic as a competing traffic source for these experiments. The access point (AP) used was a D-link DI-614+, which is one of the most popular home APs. Due to the limited number of stations used, we decided to configure the nodes to send at full rate [Sev04], however, in reality a larger number of smaller TCP flows from several users would likely compete for the media on the same 802.11 channel.

In figures 10 and 11 we show the round trip time and jitter for zero to four competing nodes. One possible explanation of the higher delay in infrastructure mode is internal scheduling/queuing within the access point, the frame also must be transmitted on the medium twice, and the effect of the congestion avoidance mechanism [HRBSD03].

Alternatively the AP has the same probability of accessing the media as any other station, the rise in the delay could be simply due to the increased number of transmitting nodes. What is clear is that this situation can lead to delays that would be unacceptable for VoIP users, far exceeding the delay for good interactivity. We observed some loss, but it was low and within the acceptable values for VoIP quality for the speech coding scheme we used (G.711). A detailed examination of the performance of access points under heavy load can be found in [Pel04].

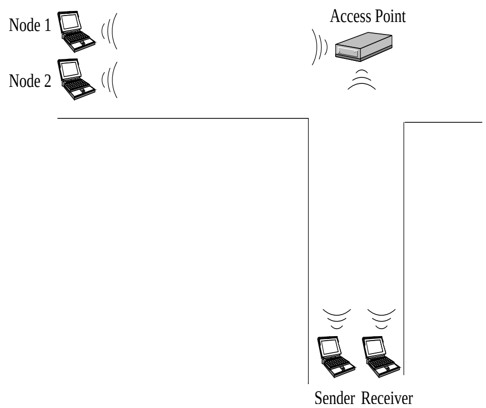

Fig. 12. A) RTS/CTS setup

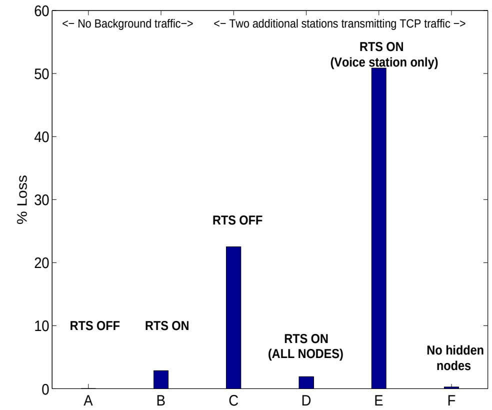

Fig. 13. B) RTS/CTS setup

E. RTS/CTS mechanisms for VoIP performance

In this set of experiments we wanted to learn if the RTS/CTS handshake mechanism is effective against the hidden node problem. For this purpose we placed the sender and receiver stations out of transmission range of two additional stations, so that each pair of stations were hidden from each other, as shown in Figure 12. Normal office walls were the obstacles preventing the radio signals from reaching the hidden stations. At the intersection of these hallways we placed an access point, which was in range of all the wireless stations. We examined 6 combinations of using the RTS/CTS mechanism:

- A No background traffic, RTS disabled in sender station
- B No background traffic, RTS enabled in sender station
- C Two stations sending TCP, RTS disabled in all the stations
- D Two stations sending TCP, RTS enabled in all the stations
- E Two stations sending TCP, RTS enabled in voice sender station only
- F Two stations sending TCP, RTS disabled, stations not hidden

Figure 13 shows the loss results. We will begin by considering the experiments without background traffic, i.e. A and B. The plot shows 3% of losses when the RTS mechanism is used.

By adding background traffic, via enabling the TCP flows between nodes hidden from the sender and receiver stations, we observe that the hidden nodes caused a 25% loss in the voice stream (the C case), whilst the loss percentage was only 0.3% when the stations were not hidden (F). In order to reduce the loss we enabled the RTS mechanism in all the stations, and the loss percentage reduced to approximately 2% (D). We also examined whether enabling RTS in the voice sender station alone would help to minimize loss. However, rather than reducing loss it increased it up to 50% (E). This shows that the RTS mechanism is effective against the collisions caused by hidden nodes, but only if all the stations enable it. In fact in our experiments enabling it in only one station causes greater degradation than leaving it disabled.

Experiment	A	B	C	D	E	F
Throughput (Mbps)	-	-	1.1	1.1	1.56	2.1

TABLE II

THROUGHPUT FOR DIFFERENT CONFIGURATIONS OF THE RTS/CTS MECHANISM

Clearly, there was no decrease in aggregate throughput after enabling RTS in all the nodes. However, throughput increased when the voice station was the only one that enabled RTS. Thus the RTS mechanism benefited the other stations whilst the enabling station experienced worse performance. The highest throughput was achieved when all the nodes were in range, because the probability of collision was lower. Figures 14 and 15 show the delay and jitter. Here we observe the overhead introduced by the RTS mechanism can be significant. The round-trip delays obtained in scenario F were very high, approximately 1.5 seconds for the RTT. Other than software errors we have no explanation for this high delay. We conclude from these measurements that the RTS/CTS mechanism is not beneficial to VoIP users unless *all* stations enable it. Since it

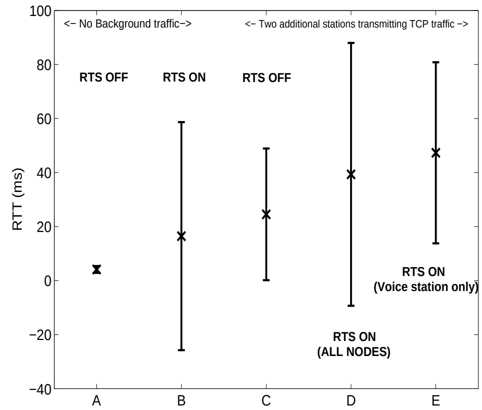

Fig. 14. RTS/CTS delays

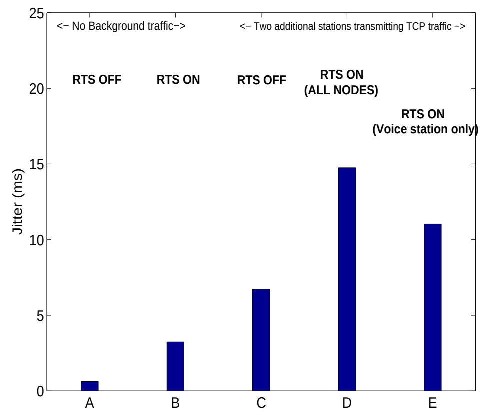

Fig. 15. RTS/CTS jitter values

is disabled by default it is safer not to enable it optimistically.

V. Conclusions

We have conducted many hundreds of experiments in order to assess the suitability of 802.11b networks for real-time voice communication. We have found that measuring the MAC layer behavior in *conjunction* with measuring the application layer performance is both useful and informative in estimating the quality of VoIP sessions. It is informative in that the occurrence of retransmissions, for example, can indicate that the terminal is entering/experiencing a period of poor quality and increase its use of forward error correction or another mechanism.

The contribution of the MAC layer itself is generally low, however delay, once introduced into a system, cannot be eliminated, unlike the perceptual effects of loss for example, so understanding the MAC layer's contribution, particularly for delay, is important. It is situation-specific as to whether 802.11b can deliver sufficient real-time voice quality. The major hurdles we encountered were attenuating objects between the end-terminals.

We have also seen that (at least) one popular access point can add delays that would seriously degrade of conversations under heavy load. Although the newer standards 802.11{a,g} allow greater capacity, they operate over shorter ranges. In the case of competing traffic, 802.11e will give priority to voice traffic when in competition with TCP bulk transfers, therefore we recommend its use.

References

[DS03] Eleftherios Dimitriou and Patrik Sörqvist. Internet Telephony over WLANS. Technical report, Global IP Sound, Sep 2003. http://www.globalipsound.com/solutions/wlan.usta_paper.pdf.

[GAEB04] Enrico Gregori Giuseppe Anastasi Eleonora Borgia, Marco Conti. IEEE 802.11b Ad Hoc Networks: Performance Measurements. In *Cluster Computing*, 2004.

[Her03] Daniel Hertrich. Experimental Performance Evaluation of an 802.11b WLAN supporting mixed Multimedia Traffic. Technical report, TU Berlin, Sep 2003.

[HGW03] Christian Hoene, André Günther, and Adam Wolisz. Measuring the Impact of Slow User Motion on Packet Loss and Delay over IEEE 802.11b Wireless Links - Row Measurement Traces. In *Proceedings of Workshop on Wireless Local Networks (WLN) 2003*, Bonn, Germany, October 2003.

[HRBSD03] Martin Heusse, Franck Rousseau, Gilles Berger-Sabbatel, and Andrzej Duda. Performance anomaly of 802.11b. In *In Proceedings of IEEE INFOCOM 2003*, San Francisco, USA, Mar 2003.

[Nun04] Victor Yuri Diogo Nunes. VoIP quality aspects in 802.11b networks. Master's thesis, IMIT, Royal Institute of Technology, Stockholm, Sweden, August 2004.

[Pel04] Enrico Pelletta. Maximum Throughput of IEEE 802.11 Access Points: Test Procedure and Measurements. Master's thesis, IMIT, Royal Institute of Technology, Stockholm, Sweden, May 2004.

[Sev04] Juan Carlos Martín Severiano. IEEE 802.11b MAC layer's influence on VoIP quality: Measurements and Analysis. Master's thesis, IMIT, Royal Institute of Technology, Stockholm, Sweden, October 2004.

[YBA02] Jihwang Yeo, Suman Banerjee, and Ashok Agrawala. Measuring Traffic on the Wireless Medium: Experience and Pitfalls. Technical report, Department of Computer Science University of Maryland, Dec 2002.