Experimental Measurement of the Capacity for VoIP Traffic in IEEE 802.11 WLANs

Sangho Shin, Henning Schulzrinne
Columbia University, Department of Computer Science
Email:{sangho,hgs}@cs.columbia.edu

***Abstract*— we measured the capacity for VoIP traffic in an 802.11b test-bed and compared it with the theoretical capacity and our simulation results. We identified factors that have been commonly overlooked in past studies but affect experiments and simulations. We found that in many papers the capacity for VoIP traffic has been measured via simulations or experiments without considering those factors, showing different capacity in each paper. After these corrections, simulations and experiments yielded a capcaity estimate of 15 calls for 64 kb/s CBR VoIP traffic with 20 ms packetization interval and 34 calls to 36 calls for VBR VoIP traffic with 0.39 activity ratio.**

I. Introduction

As many public 802.11 wireless networks have been deployed not only in buildings but also in parks and streets, the importance of usage of VoIP over wireless networks has been increasing, and a lot of research has been conducted to improve QoS and increase the capacity for VoIP traffic.

However, most papers have used simulations to measure the performance of new or current MAC protocols in IEEE 802.11 wireless networks because of the difficulty of implementing the algorithms, which are contained in the firmware of wireless interface cards. In particular, most of the papers regarding the capacity for VoIP traffic, including [1], [2] and [3], have used simulation tools to measure the capacity due to the necessity of a large number of wireless clients and difficulty in the control of them and collection of data. As far as we know, very few studies ([4] and [5]) have experimentally measured the VoIP capacity in IEEE 802.11 wireless networks, however, without any comparison with corresponding. Also, many of them failed to take into account important parameters that affect the capacity, which resulted in different capacity in each paper.

For this paper, we measured the capacity for VoIP traffic using actual wireless clients in a test-bed and compared the results with the theoretical capacity and simulation results. Additionally, we identified factors that can affect the capacity but are commonly overlooked in simulations, and analyzed the effect of these factors on the capacity in detail.

The capacity for VoIP traffic in our experiments differed significantly from our initial simulation results, and we found that we can make the similar results **closely approximation the experiments** by adjusting some factors which we generally do not change in simulations. The following factors were identified to cause the difference between simulations and experiments:

- Preamble size: Many simulation tools use the long preamble size by default while current wireless cards use the short preamble size for better efficiency due to improved RF technology, improving the VoIP capacity by 25%.
- Transmission rate control: Usually, the transmission rate is fixed in simulators to avoid the effect of rate control algorithms. However, a rate control algorithm is used by default in wireless cards, which decreases the capacity in congested channels.
- Packet generation intervals among VoIP sources: The packet generation intervals among VoIP sources are arbitrarily decided according to the call arrival pattern and the start time of each sender process of nodes, in a real environment and the experiments, respectively, while a fixed value is commonly used in simulations.
- Scanning Access Points (APs): Probe requests and response frames are used to scan APs when performing handoffs, and no such frames are transmitted in simulations unless mobility of clients is involved. However, actual wireless stations transmit probe requests not only during handoffs.
- Retry limit: The 802.11 defines both a long and a short retry limit. The short retry limit is used in transmitting RTS and CTS frames and the data frames whose size are smaller than the RTS threshold value, and the long retry limit is used for the other packets. However, the retry limits of some wireless cards differ from the standards and simulations.
- Network buffer size at the AP: When an AP receives many packets from the Distribution System (DS) and the wireless channel is congested, the buffer of the AP becomes full and packets get lost. In this case, the buffer size affects the packet loss rate and also the end-to-end delay.

We will use our experimental and simulation results to explain the factors.

This paper is organized as follows: Section II describes related work; Section III presents the theoretical capacity of VoIP traffic; Section IV describes the VoIP capacity from simulations; Section V explains the test-bed for the experiments and shows the experimental results; Section VI describes in detail the factors that can affect the experiment and simulation results; Section VIII concludes the paper.

II. Related work

Hole et al. [1] provides an analytical upper bound on the capacity for VoIP applications in IEEE 802.11b networks, evaluating a wide range of scenarios including different delay constraints, channel conditions and voice encoding schemes using an analytical method, however, assuming the long preamble only. Veeraraghavan et al [2] analyzed the capacity of a system that uses PCF for Constant Bit Rate (CBR) and Variable Bit Rate (VBR) voice traffic, using Brady's model [6] for VBR voice traffic. In their analysis, they used a value of 75 ms and 90 ms as the Contention Free Period (CFP) interval, which causes a delay that is not acceptable for VoIP. The capacity for VoIP with a 90 ms CFP interval was 26 voice calls, but the maximum delay was 303 ms. Chen et al. [3] evaluated via simulations the capacity of VoIP with IEEE 802.11e Enhanced DCF (EDCF) and Enhanced PCF (EPCF), which are called EDCA and HCCA in the final standard. They used G.711, G.729 and G.723.1 as voice codecs and assumed CBR traffic. IEEE 802.11e provides low end-to-end delay for voice packets even if mixed with best effort traffic.

In [4] and [5], the capacity for VoIP traffic was measured experimentally. However, most of these factors we mentioned in the previous section were not taken into account, and no comparison with simulation results was provided. Sachin et. al. [4] experimentally measured the capacity for VoIP traffic with 10 ms packetization interval and the effect of VoIP traffic on UDP data traffic in 802.11b. They found that the capacity of such VoIP traffic is six and the effective available bandwidth is reduced by ongoing VoIP connections. Anjum et. al. [5] also measured the capacity and the performance of their new scheme, Backoff Control and Priority Queuing (BC-PQ) experimentally. However, in order to decide the capacity for VoIP traffic, they used the packet loss rate, which depends on the network buffer size of the AP in DCF, unless the wireless link is unreliable, as we will show in Section VI-D.4. They found that the capacity with 20 ms packetization interval is 10 calls, which differs from our results and we believe that this is because of the effect of the Auto Rate Fallback (ARF) and preamble size, but such parameters are not mentioned in the paper.

III. Theoretical capacity for VoIP traffic

First, we analyze the capacity for VoIP traffic theoretically to get an upper bound, and compare it with the capacity observed in simulations and experiments.

A. Capacity for CBR VoIP traffic

In this section, we analyze the capacity of Constant Bit Rate (CBR) VoIP numerically. We define the theoretical capacity for VoIP traffic as the maximum number of calls that are allowed simultaneously for a certain channel bit rate; we assume that all voice communications are full duplex.

The capacity of VoIP traffic (N_{max}) can be expressed as [7]:

$$N_{max} = \frac{T_p}{2(T_{DIFS} + T_{SIFS} + T_v + T_{ACK}) + T_s \cdot CW_{min}/2}, \quad (1)$$

where T_p is the packetization interval, T_v and T_{ACK} are the time for sending a voice packet including all headers and the 802.11 ACK frame, respectively, CW_{min} is the minimum contention window, T_{DIFS} and T_{SIFS} are the lengths of DIFS and SIFS, and T_s is a slot time. We used Eqn. 1 to compute the capacity for G.711, a 64 kb/s codec and 20 ms packetization interval, which generate 160B VoIP packets. Other parameters are taken directly from the IEEE 802.11b standard [8]. We used the short preamble for the comparison with the experimental results using actual wireless nodes, which also use the short preamble, and the effect of the preamble size will be discussed in Section VI-A. All the parameters used in our analysis are shown in Table I.

TABLE I
PARAMETERS IN IEEE 802.11B (11 MB/S)

Parameters	Time (μs)	Size (bytes)
PLCP[1] Preamble	72.00	18
PLCP Header	48.00	6
PLCP Header Service	192.00	24
MAC Header+CRC	24.73	34
IP+UDP+RTP headers	29.09	40
Voice	116.36	160
ACK	10.18	14
SIFS	10.00	
DIFS	50.00	
Slot	20.00	
CW_{MIN}	31 slots	

TABLE II
THEORETICAL CAPACITY FOR VOIP TRAFFIC

Packetization Interval	10 ms	20 ms	40 ms
Capacity (calls)	8	15	26

Using Eqn. 1 and the parameters in Table I, the theoretical capacity for VoIP traffic with 20 ms packetization interval and G.711 codec was computed as 15 calls. The VoIP capacity for other packetization intervals is shown in Table II.

B. Capacity for VBR VoIP traffic

Typically, VoIP traffic is half duplex rather than full duplex considering that when one side talks, the other side usually does not talk. Thus, in order to avoid wasting resources, silence suppression is used, which prevents sending background noise, generating VBR VoIP traffic. The VBR VoIP traffic is characterized by on (talking) and off (silence) periods, which determine the activity ratio and also the capacity for VBR VoIP traffic. The activity ratio is defined as the ratio of on-periods and the whole conversation time.

[1]Physical Layer Convergence Protocol

TABLE III
VOICE PATTERN PARAMETERS IN ITU-T P.59

Parameter	Average duration (s)	Fraction (%)
Talkspurt	1.004	38.53
Pause	1.587	61.47
Double Talk	0.228	6.59
Mutual Silence	0.508	22.48

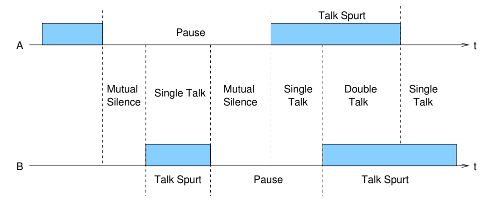

Fig. 1. Conversational speech model in ITU-T P.59

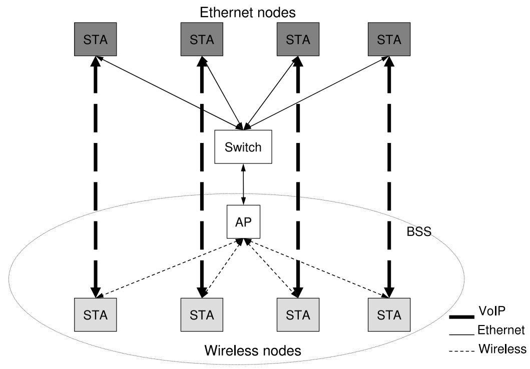

Fig. 2. Simulation topology

In this analysis, we use the conversational speech model with double talk described in ITU-T P.59 [9]. The parameters are shown in Table III, and the conversation model is shown in Fig. 1. The activity ratio in the conversational speech model is about 0.39 from the fraction of talkspurt in Table III. Therefore, the capacity for VBR VoIP traffic can be computed as: $\lfloor \text{CBR capacity}/\text{activity ratio} \rfloor = \lfloor 15/0.39 \rfloor = 38$ calls.

IV. Capacity for VoIP traffic via simulation

In this section, we measure the capacity for VoIP traffic via simulations. We used the QualNet simulator [10], which is a commercial network simulation tool and is known to have a more realistic physical model than other tools such as ns-2 [11].

We measured the 90th percentile value of the end-to-end delay of voice packets at each node with a varying number of wireless nodes. The one-way end-to-end delay of voice packets is supposed to be less than 150 ms [12]. We assumed the codec delay to be about 30-40 ms at both sender and receiver, and backbone network delay to be about 20 ms. Thus, the wireless networks should contribute less than about 60 ms delay [7]. Therefore, we defined the capacity of VoIP as the maximum number of wireless nodes so that the 90th percentile of both uplink and downlink delay does not exceed 60 ms.

A. Simulation parameters

As shown in Fig. 2, we used the Ethernet-to-wireless network topology to focus on the delay in a Basic Service Set (BSS). In the simulations, the Ethernet portion added 1 ms of transmission delay, which allows us to assume that the end-to-end delay is essentially the same as the wireless transmission delay. We used the parameters in Table I in simulations.

For VBR VoIP traffic, we implemented the conversational speech model in Fig. 1 in the simulator. Each simulation ran

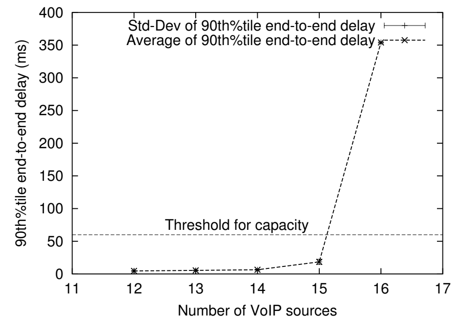

Fig. 3. End-to-end delay (90th percentile) of CBR VoIP traffic in simulations

for 200 seconds and was repeated 50 times using different seeds and VoIP traffic start time (The effect of the traffic start time will be explained in Section VI-C).

B. Capacity for CBR VoIP traffic

In order to decide the capacity for VoIP traffic, we collected the 90th percentile end-to-end delay of each VoIP flow, and calculated the average of them in each simulation, and computed the average and standard deviation of all simulation results, again. Fig. 3 shows the average and its standard deviation of the 90 percentile end-to-end delay of CBR VoIP traffic across simulations. The standard deviation represents how the 90th percentile delay fluctuates across simulations. There was no difference in delay among nodes because the downlink delay dominates the end-to-end delay and it is almost the same among nodes, as we have shown in our previous study [13]. According to the figure, we can see the capacity for the VoIP traffic is 15 calls, and the small standard deviation indicates that there is no fluctuation among simulations.

C. Capacity for VBR VoIP traffic

In the QualNet simulator, we implemented the VBR VoIP traffic with 0.39 activity ratio and exponentially distributed on-off periods, following the speech model in Fig. 1. Fig. 4 presents the end-to-end delay of VBR VoIP traffic. The end-to-end delay increases slowly compared with that of CBR VoIP traffic, and this is because only 50 kb/s (64 kb/s $\times$ 2 $\times$ 0.39) VoIP traffic is added to network as one VBR call is added. As we can see, the capacity of VBR VoIP traffic is 34 calls, which is four calls smaller than the theoretical capacity.

The standard deviation of 90%tile value of the end-to-end delay in VBR VoIP traffic is large, while the one in CBR VoIP traffic is small, and this is because of the fluctuation of the instant activity ratio, i.e., the number of nodes talking simultaneously.

V. Capacity for VoIP traffic via Experiments

We have performed experiments to measure the capacity for VoIP traffic in the ORBIT (Open Access Research Testbed for Next-Generation Wireless Networks) test-bed, which is a laboratory-based wireless network emulator located at WINLAB, Rutgers University, NJ.

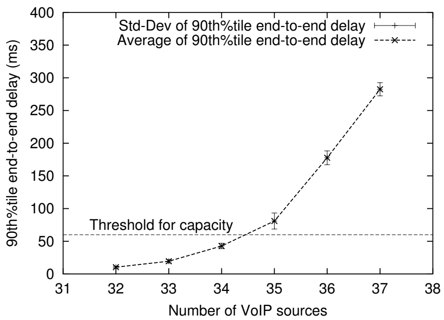

Fig. 4. End-to-end delay (90%tile) of VBR VoIP traffic in simulations

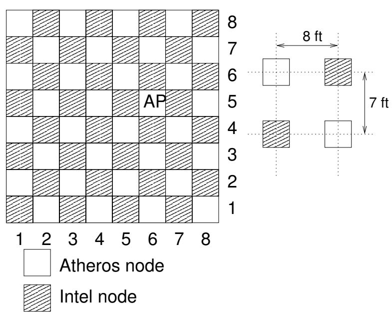

Fig. 5. Node layout in the sb9 ORBIT test-bed

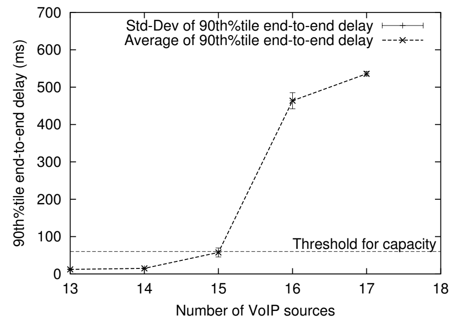

Fig. 6. End-to-end delay (90th percentile) of CBR VoIP traffic in the experiments

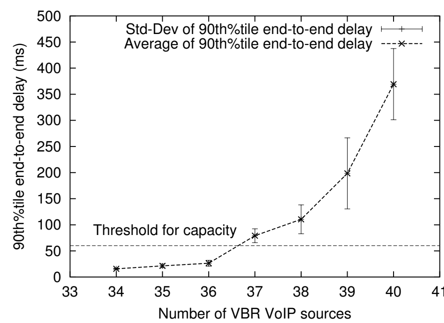

Fig. 7. End-to-end delay (90th percentile) of VBR VoIP traffic in the experiments

A. The ORBIT test-bed

ORBIT is a two-tier laboratory emulator and field trial network testbed designed to evaluate protocols and applications in real-world settings [14]. The ORBIT test-bed is composed of a main grid called 'grid' with 20 $\times$ 20 nodes and multiple smaller test-beds.

We used a small test-bed called sb9 with 8 $\times$ 8 nodes for CBR, and the main grid for VBR, which requires more nodes for experiments. The sb9 test-bed is composed of 32 nodes with the Intel chip-set wireless cards (Intel nodes), and 32 nodes with the Atheros chip-set wireless cards (Atheros nodes). Fig. 5 shows the layout of the sb9 test-bed, in which Intel and Atheros nodes are placed alternatively like a chess board. In the experiments, only Atheros nodes were used because the Atheros chip-set, which supports the Madwifi driver, is more flexible in configuring wireless cards. "AP" in the map indicates the position of the AP node. Thus, the distances between the AP and all nodes were within 50 feet, which is close enough to avoid the effect of Received Signal Strength Indication (RSSI) on packet loss. We also analyzed the RSSI of each node and will explain it later.

The main grid consists of 380 Atheros nodes and 20 Intel nodes, and forms a 20 $\times$ 20 grid with the same inter-node distance as the sb9 test-bed.

We wrote a simple UDP client which sends 172 byte (160 B VoIP payload + 12 B RTP header) UDP packets to a specified destination. The UDP client records the sending time and receiving time to separate files with the UDP sequence number, which is included as a UDP packet payload, and the data were used to calculate the end-to-end delay and the packet loss. In order to synchronize the system clock of the nodes, the Network Time Protocol (NTP) was used. The AP needed to be used as an NTP server because the ORBIT test-beds were isolated from the Internet. Client nodes updated the system clock every second using the *ntpdate* application.

The Madwifi driver was modified to print out the information of all transmitted and received frames such as RSSI, retries, and 802.11 flags, which are reported to the driver from the firmware. The information was used to calculate the retry rate and the transmission rate.

B. Experimental results

The experimental results showed much higher fluctuation than the simulations. Therefore, we performed each experiment more than 50 times with 200 s experiment time over a few months. Also, we found that 802.11e feature was enabled in the wireless nodes in the test-bed by default, which has a siginificant impact on performance, and thus we disabled it in all nodes in our experiments. Furthermore, to avoid the effect of physical conditions of nodes, we removed the nodes with unreasonably higher retransmission rate than others in the experiments. **Explain more.**

1) Capacity for CBR VoIP traffic: Fig. 6 shows the average 90th percentile end-to-end delay and standard deviation with 13 to 17 CBR VoIP sources. The 90th percentile end-to-end delay with 15 CBR VoIP sources is 57 ms, which means that the capacity of CBR VoIP traffic is 15 calls, which is exactly the same as the theoretical capacity and the simulation result.

2) Capacity for VBR VoIP traffic: Fig. 7 shows the average 90th percentile end-to-end delay and standard deviation with

34 to 40 VoIP sources. According to the result, the capacity for VBR VoIP traffic with 0.39 activity ratio is 36 calls, which is a little bit greater than the simulation capacity, 34 calls, and slightly smaller than the theoretical one, 38 calls.

VI. Analysis of the experiment and simulation results

Our initial experimental results, which are not included in this paper, showed a big difference from the capacity from the theoretical analysis and simulations, and we found some parameters that are commonly ignored but affect the experiment and simulation results. In this section, we discuss the factors in detail with some additional experimental and simulation results.

We will focus on the CBR traffic in the analysis because we want to avoid the effect of activity ratio, which is a main factor to the experimental results with VBR VoIP traffic, and because the effect of the following factors on MAC layer would be the same in both CBR and VBR VoIP traffic.

A. Preamble size

The preamble size affects the capacity for VoIP traffic in IEEE 802.11b networks, and some simulators and wireless cards use different ones. Preamble is a pattern of bits attached at the beginning of all frames to let the receiver get ready to receive the frame, and there are two kinds of preamble, long and short. The long preamble has 144 bits and the short one has 72 bits. The long preamble allows more time for the receiver to sync and be prepared to receive while the transmission time becomes longer, since preamble is transmitted with 1 Mb/s. The QualNet simulator as well as NS-2 uses the long preamble by default while recent wireless cards use the short preamble due to advancing RF technology, improving the utilization of channels. Typically, users cannot change the preamble type since it is hard-coded in the firmware of wireless cards, but we can check what type of preamble is used in the wireless card by reading the register value where the value is written.

Considering the small packet size of VoIP packets and its low transmission speed, the preamble takes up a big portion of a VoIP packet; 144 bits of 2064 bits, taking 4% in size, 144 μs of 362 μs, taking 40% in the transmission time. Thus, the theoretical capacity for the VoIP traffic using DCF decreases from 15 to 12 calls when the long preamble is used.

Fig. 8 presents the 90th percentile value of end-to-end delay of CBR VoIP traffic with the long preamble in simulations, and we can see that the capacity decreased to 12 calls, as in the theoretical capacity.

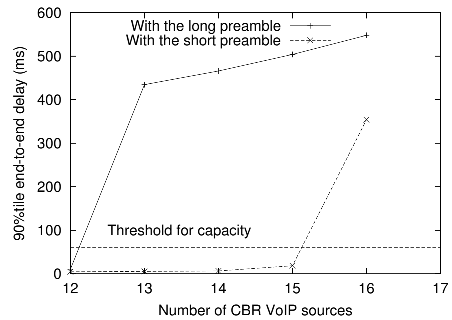

Fig. 8. End-to-end delay (90th percentile) of CBR VoIP traffic with long and short preamble via simulations

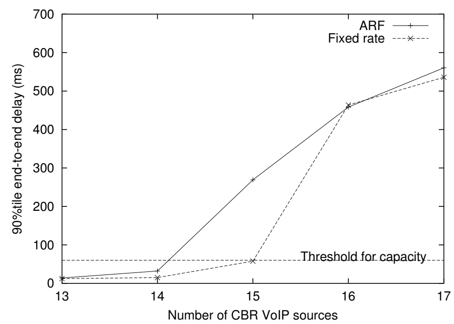

Fig. 9. End-to-end delay (90th percentile) of CBR VoIP traffic with and without the AMRR rate control algorithm in the experiments

B. Rate control

Most wireless cards support multi-rate data transmission, and wireless card drivers support Auto-Rate Fallback (ARF) to choose the optimal transmission rate according to the link status. Generally, the transmission rate decreases when the packet loss exceeds a certain threshold and increases after successful transmissions, but the specific behavior depends on the ARF algorithm.

A smart rate control algorithm improves the throughput and the channel utilization, however, only when the packet loss is caused by wireless link problems. ARF makes the channel utilization and throughput worse if the main reason for the packet loss is packet collisions [15]. In this case, the transmission with a low bit-rate makes the transmission time of frames longer and increases the delay without improving the packet loss, and the packet transmission with the highest available bit-rate achieves the best throughput.

The Madwifi driver supports two rate control algorithms, ONOE and Adaptive Multi-Rate Retry (AMRR) [16]. AMRR is used by default, but we can disable it by modifying the driver source code. Fig. 9 shows the experimental results with the AMRR rate control algorithm. As we can see, the capacity decreased to 14, and this is because with 15 CBR VoIP sources, about 8% of the packets are transmitted with lower transmission rates due to the rate control algorithm (Table IV). Fig. 10 shows the retry rate of VoIP traffic with and without the AMRR rate control module. We can see that more collisions

TABLE IV
TRANSMISSION RATE OF CBR VOIP TRAFFIC WITH ARF ENABLED IN THE EXPERIMENTS

	11 Mb/s	5.5 Mb/s	2 Mb/s	1 Mb/s
12 calls	96.0%	3.6%	0.0%	0.0%
13 calls	95.8%	3.7%	0.0%	0.0%
14 calls	92.5%	5.3%	1.7%	0.0%
15 calls	92.0%	5.7%	0.0%	1.6%
16 calls	90.9%	8.2%	0.1%	0.0%
17 calls	90.5%	8.4%	0.1%	0.0%

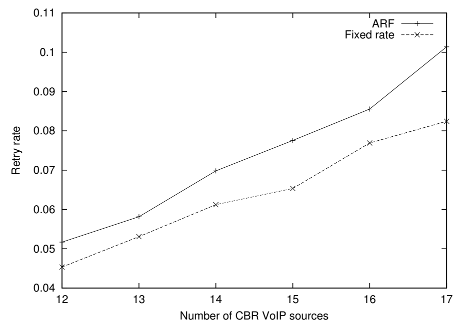

Fig. 10. Retry rate of CBR VoIP traffic with and without the AMRR rate control algorithm in the experiments

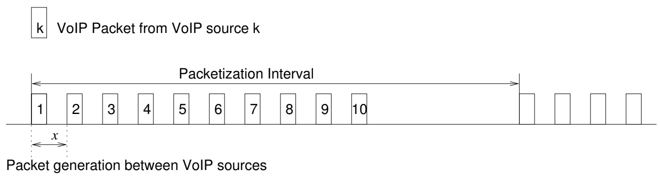

Fig. 11. An example of VoIP packet transmission in the application layer with 10 VoIP sources and the fixed transmission interval of x

were caused when the algorithm was used. The effect of the rate control on the capacity depends on the algorithm and the RF conditions, and the analysis of the algorithm is beyond the scope of this paper.

The QualNet simulator supports a few ARF algorithms, and NS-2 also has many external rate control modules. However, generally a fixed transmission rate is used in most simulations to avoid the effect of the rate control algorithm, while many wireless card drivers use a rate control algorithm by default. Therefore, when comparing the results from simulations and experiments, rate control should be disabled or exactly the same rate control algorithm should be used in both simulators and the drivers of all wireless nodes.

C. VoIP packet generation intervals among VoIP sources

In simulations, normally all wireless clients start to generate VoIP traffic at the same time, but the packet generation interval between clients affect the simulation results.

Wireless clients can transmit VoIP packets if the medium is idle without further backoff because backoff is already performed right after the transmission of the previous VoIP packet. Note that backoff is done right after a successful transmission, and wireless clients generate VoIP packets every packetization interval, which is typically 10 ms to 40 ms. We found in our previous study [13] that when the number of VoIP sources does not exceed the capacity, uplink delay is very small (less than a few milliseconds) due to the unfair resource distribution, which means that the outgoing queue of VoIP wireless clients is mostly empty.

Therefore, generally, when two VoIP sources generate VoIP packets at the same time, the collision probability of the two packets becomes high. Conversely, when the VoIP packet generation times of all VoIP sources are evenly distributed within a packetization interval, the collision probability between nodes

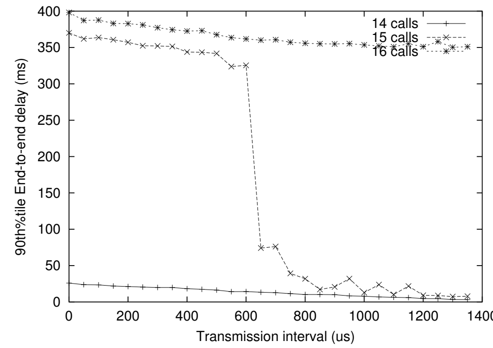

Fig. 12. The end-to-end delay (90th percentile) in each packet generation interval among VoIP sources

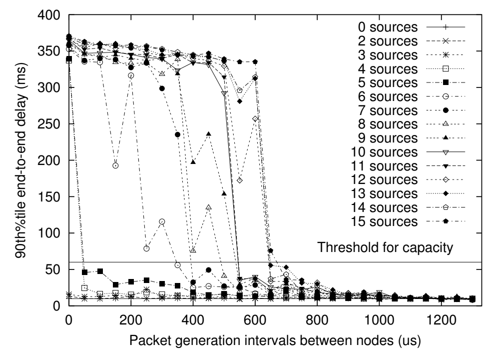

Fig. 13. The end-to-end delay (90th percentile) in each packet generation interval among VoIP sources and each number of VoIP sources with such intervals

becomes lowest. Fig. 11 shows an example of VoIP generation time of 10 VoIP sources. For simplicity, we assumed that all the packet generation intervals (x) between nodes are the same, although they are random in real environments.

The Fig. 12 shows the 90th percentile value of the end-to-end delay of the VoIP traffic in simulations with 14, 15 and 16 VoIP sources and 0 to 1350 μs packet generation intervals (x). We can see that the delay decreases as the interval value increases in all three cases. In particular, the delay of 15 VoIP calls, which is the capacity for VoIP traffic, decreases significantly at the interval of 650 μs. This is because CW_{min} is 31 and the initial backoff time is decided between 0 and 620 μs in 802.11b, meaning that when the interval (x) of two packets is larger than 620 μs, the collision probability of the two packets theoretically drops to zero.

In order to identify the effect of the packet generation intervals and the number of VoIP sources with such intervals, we performed additional simulations with 15 VoIP sources, changing the intervals and the number of VoIP sources with such intervals, assigning the largest interval (1330 μs in 15 VoIP sources) to the rest of the VoIP sources; for example, setting the packet generation intervals of 100 μs to two VoIP sources and 1330 μs to the remaining 13 VoIP sources.

Fig. 13 shows the simulation results giving the 90th percentile end-to-end delay for all possible combinations of the interval and the number of VoIP sources. For example, when packet generation intervals of seven VoIP sources among 15

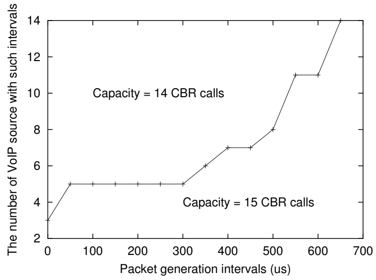

Fig. 14. Effect of the packet generation intervals and the number of VoIP sources with such intervals on the capacity. When the intervals and the number of VoIP sources with such intervals are in the left upper area in the graph, the capacity decreases to 14 calls.

VoIP sources are $350\,\mu s$, the 90th percentile end-to-end delay of 15 VoIP sources is around 250 ms, and if the packet generation interval increases to $400\,\mu s$, the 90th percentile end-to-end delay decreases below 60 ms and satisfies the condition for the capacity. The results of Fig. 13 is summarized in Fig. 14, which shows the intervals and the maximum number of VoIP sources with such intervals that satisfy the condition with 15 VoIP calls. For example, when the number of VoIP sources that generate VoIP traffic at the same time is more than three, the capacity decreases to 14 calls, and if the minimum packet generation interval is larger than $650\,\mu s$, the capacity is always 15 calls regardless of the number of nodes with such intervals.

We have shown that the capacity of VoIP traffic varies from 14 to 15 calls according to the interval. Therefore, in our other simulations, the starting time of each VoIP traffic was chosen randomly between 0 to 20 ms (packetization interval), as it is arbitrarily decided also in the experiments.

D. Other minor factors

In this section, we discuss the factors whose effect was minor and did not change the capacity for VoIP traffic in this study, but that may affect the experiment and simulation results.

1) Received Signal Strength Indication (RSSI): We analyze the effect of the RSSI on the experiments because the RSSI could be the main reason for the frame loss in experiments, but such effect is commonly avoided in simulations.

Fig. 15 shows the RSSI of uplink and downlink flow in each node. The vertical bar shows the range of the RSSIs of a node, and the height of the box and the center line represent the standard deviation and the average of RSSIs across the experiments, respectively. Moreover, in order to identify the correlation between the RSSI and the distance from the AP, the RSSI of each node is plotted according to the distance between the node and the AP.

We can see that the RSSIs of most of the nodes fall within the interval -55 to -70 dBm, and the fluctuation of RSSI of each node across the experiments was mostly within 5 dBm. Only Node 6 had a relatively weak signal, but it was still within effective ranges. Also, we also could not find any correlation between the RSSI and the distance from the AP: the weak signal was not because of the distance from the AP.

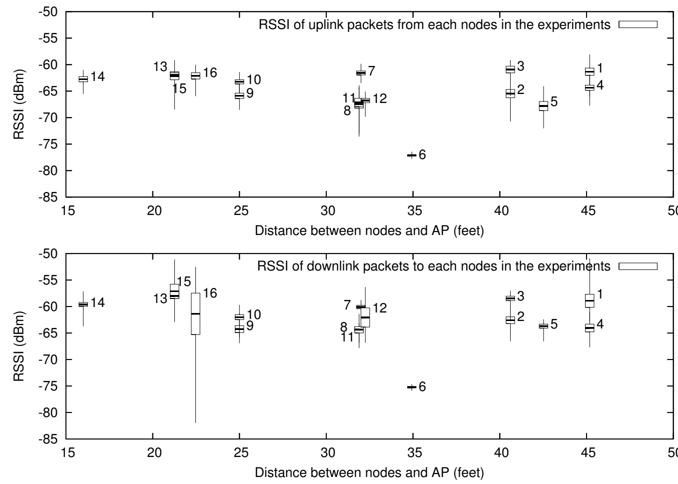

Fig. 15. RSSI values according to the distances between nodes and the AP. The vertical bar shows the range of the RSSIs of a node (min and max), and the height of the box and the center line represent the standard deviation and the average of RSSIs across the experiments, respectively. The numbers represent the node number.

Furthermore, in order to check the effect of RSSI on the experiments, we analyzed the correlation between the RSSI and the retry rate, but no correlation was found in both downlink and uplink. This means that the signal was strong enough and any bit error was not caused directly by the weak signal. However, the frames with the RSSI below -76 dBm (from Node 6 in Fig. 15) had a higher retry rate than the other nodes in uplink, in particular, with 17 VoIP sources. The behavior was caused by the capture effect where the frame with the stronger signal can be captured by the receiver through a collision [17].

2) Effect of scanning APs: We have observed probe request and response frames in the experiments, while no probe requests and responses were observed in simulations.

Probe request frames are transmitted from wireless clients to scan APs for handoff, and response frames are transmitted from the AP [8]. Typically, wireless clients transmit a probe request frame to a channel and scan all channels, for example, 11 channels in IEEE 802.11b. The handoff procedure is implemented in the firmware of wireless cards and each vendor uses different algorithms. Thus, it is hard to determine the exact reason for transmitting probe request frames. However, typically, wireless clients start to scan APs to find better APs when they experience a certain amount of packet loss [18]. We have also found in the experiments that as the number of VoIP sources increases, the retry rate and the number of probe request frames also increases, as shown in Fig. 16.

Probe request and response frames increase the delay of VoIP packet transmission due to traffic increase and the higher priority of management frames over data frames. As we have shown in Fig. 16, the effect on the capacity of VoIP traffic is negligible in the experiments. However, handoff algorithms depend on the firmware of the wireless cads; for example, some wireless cards regularly scan APs for better handoff [18]. In this case, the effect of scanning APs becomes bigger.

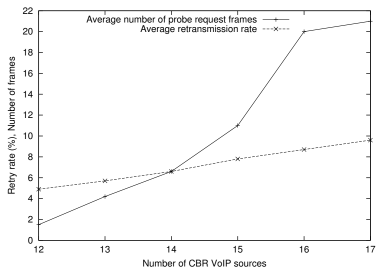

Fig. 16. Probe request frames and retry rate in the experiments

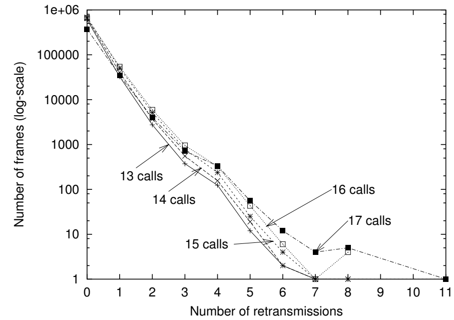

Fig. 17. Distribution of number of retransmissions of CBR VoIP traffic in the experiments

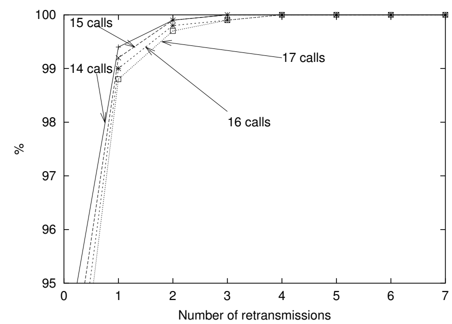

Fig. 18. Cumulative distribution function of the number of retransmissions of CBR VoIP traffic in the experiments

Therefore, the handoff behaviour of wireless cards needs to be investigated before performing experiments with the wireless cards.

3) Retry limit: IEEE 802.11 standard defines two kinds of retry limit, long and short. The short retry limit is typically used when the MAC frame size of packets is equal to or smaller than the RTS threshold value, and the long retry limit is used otherwise [8]. Although the specific values are not defined in the standard, seven and four are accepted as appropriate long and short retry limit values, in that order, and they are not configurable in wireless cards.

The wireless cards with the Atheros chip-set in the ORBIT test-bed used the long retry limit even if the RTS threshold value was set to off (infinite). Fig. 17 shows the distribution of the number of retransmissions in the experiments. According to the figure, when the number of VoIP sources exceeds the capacity, packets are retransmitted at most 11 times, which indicates that the long retry limit is 11 in the Atheros nodes. However, the QualNet simulator uses 7 and 4 as the short and long retry limit, respectively, and the short retry limit is used when the packet size is smaller than or equal to the RTS threshold value, as in the standard.

The retry limit affects the packet loss and also the delay. Fig. 17 shows that the retry limit did not cause packet loss as long as the number of VoIP sources remains below the capacity, 15 calls; there was no packet loss due to the retry limit before the number of VoIP sources reaches the capacity, and the packet loss due to the retry limit is also negligible, even with 17 VoIP calls. Fig. 18 shows the cumulative distribution function of the number of retransmitted packets. According to the figure, we can see that the packet loss would be the same even if 7 was used as the retry limit in the experiments, which shows that the difference in the retry limits did not affect the experiments.

4) Network buffer size and packet loss: The packet loss rate is also another metric to measure the capacity and QoS for VoIP traffic, and it is known that 1 to 3% packet loss is tolerable for VoIP traffic [12]. Table V shows the packet loss rate in the experiments. With 15 VoIP sources, which is the capacity of VoIP traffic, the packet loss rate was only 0.6% in the experiments, and this satisfies the condition.

We have found that the packet loss happened only in the downlink and we have already shown that the packet loss due to excess of the retry limit was almost zero in the previous section, which means that the packet loss was caused by buffer overflow at the AP.

The network buffer size at the AP directly affects the packet loss and also the end-to-end delay. For example, when the buffer size is 128 KB and 15 CBR VoIP sources send packets with 20 ms packetization interval, the AP receives 129 KB data (= 50×172B×15) every second on average, which means that the data for 992 ms (128 KB/129 KB) can be stored at the AP, causing a maximum delay of 992 ms. This also means that the packet loss due to the buffer overflow occurs when the delay exceeds 992 ms. The formula to get the maximum delay can be summarized as follows:

$$D_{max} = \frac{B}{\frac{1}{P} \cdot S \cdot N} \quad (2)$$

where, D_{max} is the maximum delay, B is the buffer size of the AP, P is the packetization interval, S is the VoIP packet size, and N is the number of VoIP sources. Fig. 19 shows the maximum queuing delay from Eqn. 2 and the packet loss rate assuming 500 ms queuing delay at the AP. We can see that as the network buffer size increases, the packet loss rate decreases, but the maximum queuing delay increases. Therefore, having a too long buffer at the AP does not improve QoS of VoIP traffic.

TABLE V

AVERAGE PACKET LOSS RATE IN THE EXPERIMENTS

	14 calls	15 calls	16 calls	17 calls
Packet loss rate	0.29 %	0.59%	4.86%	15.46 %

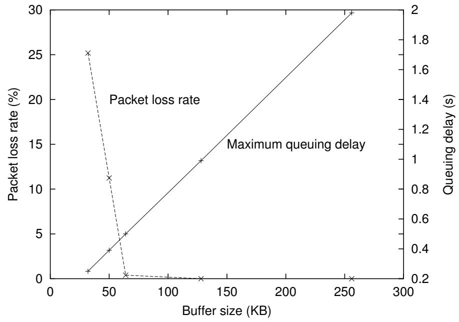

Fig. 19. Maximum queuing delay and packet loss rate according to the network buffer size of the AP ($S = 172B$, $P = 20\,ms$, $N = 15$)

VII. Limitations

We have identified many parameters that can affect the experiments and simulations, and confirmed that some did not affect the capacity directly in the experiments. However, after considering all those factors, there still exist minor differences between the experimental results and simulation ones. We believe that this is because there could be another factors or implementation issues that we could not identified clearly. One example is how to handle EIFS (Extended Inter-Frame Spacing). QualNet and NS-2 process EIFS in slightly different ways: when wireless clients receive another error frame during EIFS, QualNet restarts the timer for EIFS while NS-2 does not. The effect seems to be small, however, the effect on the capacity is not certain if the effect is combined with others in very congested channels, because such effect is easily amplified in congested channels. Here is such an example. The QualNet simulator had a minor bug in setting the backoff timer. The backoff timer was set within $[0, CW - 1]$ instead of $[0, CW]$ [2]. However, when this minor bug was combined with 300 μs packet generation interval, the 90th percentile end-to-end delay was changed to 26 ms from 365 ms with 15 VoIP sources, changing the capacity to 15 calls from 14 calls [3]. This confirms that very small differences in implementation can have a major impact on the experiments or simulations when the number of VoIP calls reaches the capacity.

In VBR VoIP traffic, such effect appears bigger due to the fluctuation of the number of nodes talking simultaneously; it increased up to 27 nodes with 35 VBR calls in the experiments and simulations. For these reasons, after considering all those factors, still the capacity for VBR VoIP traffic shows slight difference between the simulations and experiments, 34 calls and 36 calls, respectively.

The source code of the firmware of wireless cards is not public, and we do not know how the IEEE 802.11 standard is implemented in the firmware of each vendor. Therefore, there could be some factors that have not been identified in this study, and they can affect the capacity depending on the experimental environment or simulation tools.

[2]This bug was reported by Hoon Chang (Columbia University) and fixed in QualNet 3.9, and all simulations were performed with the bug-fixed simulator.

[3]Originally, the capacity with 300 μs packet generation interval is 14 calls

VIII. Conclusion

We measured the capacity for VoIP traffic via experiments with actual wireless clients in the ORBIT test-bed, and compared it with the theoretical capacity and simulation results. We also identified some factors that are commonly ignored in simulations and experiments but affect the capacity significantly, and we analyzed the effect in detail with additional experiments and simulations.

We confirmed that after considering all those factors, the capacity for 64 kb/s CBR VoIP traffic with 20 ms packetization interval is 15 calls via both experiments and simulations, and this is the same with the theoretical one. VBR results still show minor difference between simulations and experiments due to the limitations mentioned in the previous section, and the capacity with 0.39 activity ratio is 34 to 36 calls.

Even though we analyzed the effect of those factors on the VoIP capacity in this study, those factors affect any experiment and simulation with 802.11 WLANs, and this study can be utilized in the analysis and comparison.

References

[1] D. P. Hole and F. A. Tobagi, "Capacity of an IEEE 802.11b wireless LAN supporting VoIP," in *IEEE ICC*, Paris, France, 2004, pp. 196–201.

[2] M. Veeraraghavan, N. Cocker, and T. Moors, "Support of voice services in IEEE 802.11 wireless LANs," in *INFOCOM*, vol. 1, Anchorage, Alaska, 2001, pp. 488–497.

[3] D. Chen, S. Garg, M. Kappes, and K. S. Trivedi, "Supporting VoIP traffic in IEEE 802.11 WLAN with enhanced medium access control (MAC) for quality of service," Avaya Labs Research, Tech. Rep., 2002.

[4] S. Garg and M. Kappes, "An experimental study of throughput for UDP and VoIP traffic in IEEE 802.11b networks," in *IEEE WCNC*, vol. 3, New Orleans, LA, 2003, pp. 1748–1753.

[5] T. Kodama and Y. Katsube, "Voice performance in WLAN networks - an experimental study," in *Globecom*, San Francisco, CA, 2003, pp. 3504–3508.

[6] P. Brady, "A model for generating on-off speech patterns in two-way conversation," *Bell Syst. Tech. Journal*, vol. 48, no. 7, pp. 2245–2272, Sept. 1969.

[7] T. Kawata, S. Shin, A. G. Forte, and H. Schulzrinne, "Using dynamic PCF to improve the capacity for VoIP traffic in IEEE 802.11 networks," in *IEEE WCNC 2005*, vol. 3, New Orleans, LA, Mar 2005, pp. 13–17.

[8] *Wireless LAN Medium Access Control (MAC) and Physical (PHY) specifications*, IEEE, 1999.

[9] *Artificial Conversational Speech*, ITU-T P.59, 1993.

[10] QualNet Network Simulator 3.7. [Online]. Available: http://www.scalable-networks.com

[11] *QualNet 3.7 User's Guide*, Scalable Network Technologies, Inc, 2005.

[12] *One-way Transmission Time*, ITU-T G.114, 2003.

[13] S. Shin and H. Schulzrinne, "Balancing uplink and downlink delay of VoIP traffic in IEEE 802.11 WLANs using APC," in *Qshine 2006*, Waterloo, Canada, Aug 2006.

[14] D. Raychaudhuri, I. Seskar, M. Ott, S. Ganu, K. Ramachandran, H. Kremo, R. Siracusa, H. Liu, and M. Singh, "Overview of the ORBIT radio grid testbed for evaluation of next-generation wireless network protocols," in *IEEE WCNC 2005*, New Orleans, LA, 2005, pp. 1664–1669.

[15] M. Rodrig, C. Reis, R. Mahajan, D. Wetherall, and J. Zahorjan, "Measurement-based characterization of 802.11 in a hotspot setting," in *ACM SIGCOMM 2005 W-WIND Workshop*, Philadelphia, PA, Aug 2005, pp. 5–10.

[16] M. Lacage, M. H. Manshaei, and T. Turletti, "IEEE 802.11 rate adaptation: a practical approach," in *MSWiM 2004*. Venice, Italy: ACM Press, 2004, pp. 126–134.

[17] H. Kremo, I. Seskar, and P. Spasojevic, "An ORBIT testbed study of 802.11b DCF: Throughput, latency, and the capture effect," in *IEEE Tridentcom 2006*, Barcelona, Spain, Mar 2006, pp. 308–309.

[18] H. Velayos and G. Karlsson, "Techniques to reduce IEEE 802.11b MAC layer handover time," Royal Institute of Technology, Stockholm, Sweden, Tech. Rep., April 2003.