Capacity of an IEEE 802.11b Wireless LAN supporting VoIP

To appear in Proc. IEEE Int. Conference on Communications (ICC) 2004

David P. Hole and Fouad A. Tobagi
Dept. of Electrical Engineering, Stanford University, Stanford, California 94305
Email: {dhole, tobagi}@stanford.edu

***Abstract*— In this paper we evaluate the capacity of an IEEE 802.11b network carrying voice calls in a wide range of scenarios, including varying delay constraints, channel conditions and voice call quality requirements. We consider both G.711 and G.729 voice encoding schemes and a range of voice packet sizes.**

We first present an analytical upper bound and, using simulation, show it to be tight in scenarios where channel quality is good and delay constraints are weak or absent. We then use simulation to show that capacity is highly sensitive to the delay budget allocated to packetization and wireless network delays. We also show how channel conditions and voice quality requirements affect the capacity. Selecting the optimum amount of voice data per packet is shown to be a trade-off between throughput and delay constraints: by selecting the packet size appropriately given the delay budget and channel conditions, the capacity can be maximized.

Unless a very high voice quality requirement precludes its use, G.729 is shown to allow a capacity greater than or equal to that when G.711 is used, for a given quality requirement.

I. Introduction

Transmitting voice over wireless communication links has been in widespread use for many years - this is clearly shown by the huge takeup of mobile telephony around the world. Cellular networks provide coverage by locating antennae every few kilometres, either on purpose-built masts or on top of existing buildings. While these provide coverage over a large area, reception inside buildings is often very poor compared with the signal quality available outside.

In this paper we investigate whether IEEE 802.11 devices (the price of which has fallen significantly in recent years) could be used to create a low-cost wireless voice network that could be integrated with wired Voice over IP networks, or connected directly to cellular networks.

The scenario we are considering is shown in Fig. 1. The network comprises a single IEEE 802.11 basic service set (BSS) with one access point (AP), and a number of wireless users. The AP is connected to a wired network, to which other users are directly connected. Voice calls take place between a user in the BSS and a user connected to the wireless network (e.g. between users A and A').

We focus on the capacity of the wireless network as the principal metric of interest; this is important not only for deployment of these networks, but also as a means of comparing protocols and techniques in future work. We define capacity in this context to be the maximum number of simultaneous, bi-directional calls that can be supported, subject to a minimum voice quality requirement.

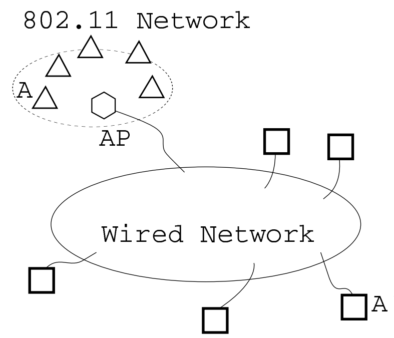

Fig. 1. Network Scenario

Voice traffic is generated by packetizing the output of a voice encoder (we consider both G.711 and G.729 schemes, without the use of silence suppression), creating packets each containing a fixed amount of voice data; we consider 10, 20, 30 or 50ms of voice data per packet. These packets are transmitted over the network using RTP over UDP/IP. As an example of parameters used in current VoIP implementations, Cisco 7960 VoIP phones can use either the G.729 or G.711 codecs, and the default amount of voice data per packet is 20ms.

The various IEEE 802.11 protocols [1] provide a 'best-effort' service: packets are carried without any delay guarantee, and may be dropped within the network. Here we consider that the wireless network operates using the Distributed Coordination Function (DCF) MAC protocol (without the RTS/CTS mechanism enabled); although the Point Coordination Function (PCF) protocol was designed to better handle stream-type traffic, this has not been widely implemented. The DCF protocol is based on CSMA/CA, whereby stations must detect that the medium is idle before transmitting. Because of the possibility of collisions and/or decoding errors (e.g. due to poor channel conditions), frames may require multiple transmissions before being successfully received. After an unsuccessful transmission (indicated by the lack of acknowledgment), the random backoff duration is selected from an exponentially increasing range. After a certain number of retransmissions, the sending station will drop the frame. Hence, the MAC protocol introduces significant randomness in the packet delay as well as the possibility of packet loss.

Although the playout buffer at the receiver reduces the effect of variation in delay, and packet loss concealment (PLC) can mitigate the effects of packet loss, the combined effects of delay (i.e. the end-to-end delay as set at the playout buffer) and packet loss (either within the network, or at the playout buffer due to excessive delay) on the quality of the voice call must be taken into consideration. For this, we use the "Mean Opinion Score" (MOS) metric. Call quality is rated on a scale of 1 (worst) to 5 (best) - typically, a score of 3.6 or higher is considered satisfactory. Prior research [2] enables us to quantify the effects of packet loss and delay in terms of the reduction of a call's MOS. Even without packet loss or delay, voice encoding schemes introduce some quality degradation; we refer to this as the 'intrinsic' MOS of a scheme.

There have been few previous studies of the ability of the IEEE 802.11 DCF protocol to support voice traffic. One such example is [3], in which physical layers supporting only 1 and 2 Mbps were considered. In [4], a wider range of data rates was considered, as well as varying channel conditions; however, in this paper, the metric of interest was mean channel access delay, rather than capacity. In [5], the capacity of an IEEE 802.11a network was evaluated, taking into account loss at the playout buffer, for a single delay constraint scenario.

Our work is most closely related to that of Garg and Kappes [6], [7]. In these papers, the authors present a means of estimating the capacity of a voice-only 802.11b network, and validate the estimation by means of limited experimentation. Our study, although driven by the same motivation, uses simulation to evaluate the capacity for a much wider range of scenarios and constraints than considered in [6], [7].

In none of these papers has any metric for the quality of the voice call (such as MOS) been considered. However, such studies are common in the context of wired networks: [8] is one such example.

In the following section, we present an analysis of the MAC protocol which, by making simplifying assumptions, leads to an upper bound on the capacity of the network. The tightness of this bound for error-free scenarios is then evaluated by simulation. In section III, we use simulation to show how delay constraints and channel conditions affect the capacity of a network.

For the simulations, we used the ns-2 network simulator [9]. For the purposes of this study, various modifications were made. For example, the 802.11 code was rigorously checked and various corrections made (with specific attention to the use of the Extended Interframe Space), and a new module to model the playout buffer was created.

II. Mathematical Analysis of Network Capacity

In this section we present an upper bound on the network capacity, by making certain assumptions about the performance of the network. These assumptions are that: i) no collisions occur; ii) frames are always received without errors, and iii) all frames arrive at the playout buffer before their respective playout deadline. We refer to such a scenario as being *throughput constrained*.

The analysis is based on the following argument. At any point in time, one of the following is taking place on the wireless network:

- The frame sequence for the transmission of data from the AP to the stations is ongoing. We define a frame sequence as the transmission of the voice data frame (with transmission time T_{VOICE}), the Short Interframe Space (SIFS), the transmission of the acknowledgement (T_{ACK}) and the DCF Interframe Space (DIFS) following the acknowledgement.
- The frame sequence for transmission of data from a station to the AP is ongoing.
- The medium is idle, and the AP is counting slots as part of a backoff procedure. Recall that all stations must wait ('backoff') for a random number of idle slots following each transmission.
- The medium is idle, and the AP is not counting down idle slots.

We consider an arbitrarily long period of time T seconds during which N calls are in progress. Let R be the number of packets generated by each encoder per second.

Of the T seconds, the time required for the frame sequences for transmissions to and from the AP is given by $2NRT\left(T_{\text{VOICE}} + \text{SIFS} + T_{\text{ACK}} + \text{DIFS}\right)$. Similarly, the minimum amount of idle time required for the AP to complete its backoff procedures is $\left[\sum_{i=1}^{i=NRT} CW_i\right] \times T_{\text{SLOT}}$, where T_{SLOT} is the slot duration, and CW_i is the number of slots picked from a uniform distribution over (0, CW_{MIN}) for the ith transmission. For large T this expression converges to $NRT\left(T_{\text{SLOT}} \times CW_{\text{MIN}}/2\right)$.

Considering only the first three possible uses of the time then, we require that, in order to support the offered load,

$$\begin{aligned} T \geq & \left[2NRT\left(T_{\text{VOICE}} + \text{SIFS} + T_{\text{ACK}} + \text{DIFS}\right)\right] + \\ & \left[NRT\left(T_{\text{SLOT}} \times CW_{\text{MIN}}/2\right)\right]. \end{aligned} \tag{1}$$

This expression becomes an equality when the amount of idle time which is not counted towards the AP's backoff, T_{IDLE}, is zero.

It is very hard to find an expression for T_{IDLE} since it depends both on the load on the network and the way in which stations carry out their backoff procedures. However, we argue that as the load approaches capacity, this value will become very small. Since devices may count down backoff slots concurrently with each other, the most efficient use is made of the network when idle slots are counted by many devices simultaneously. In this network, one single device (the AP) is transmitting half of all the traffic, and so has the greatest single requirement for non-overlapping idle time. We therefore argue that as the load (i.e. the number of calls N) increases, stations will take advantage of the idle time required by the AP to fulfil all of their backoff requirements, thereby minimizing the amount of time during which the medium is idle, and not being counted towards the AP's backoff requirements. By making the assumption that, at capacity, $T_{\text{IDLE}} = 0$, we obtain the upper bound on the value of N given in (2).

$$N = \left\lfloor \frac{1}{R\left[2(T_{\text{VOICE}} + \text{SIFS} + T_{\text{ACK}} + \text{DIFS}) + (T_{\text{SLOT}} \times CW_{\text{MIN}}/2)\right]} \right\rfloor \quad (2)$$

TABLE I

COMPONENT TIMES OF T_{VOICE} & T_{ACK} AT 11MBPS

T_{VOICE}	PLCP Preamble & Header	192.0us
	MAC Header + FCS	20.4us
	IP/UDP/RTP header	29.1us
	Voice Data	(Voice octets $\times 8/11$) us
T_{ACK}	PLCP Preamble & Header	192.0us
	ACK Frame	10.2us

TABLE II

CAPACITY RESULTS FROM SIMULATION (ANALYSIS)

	Voice Data per frame			
	10ms	20 ms	30ms	50ms
G.711	6 (6)	12 (12)	17 (18)	25 (26)
G.729	7 (7)	14 (14)	21 (22)	34 (35)

For 802.11b, CW_{MIN}, SIFS, T_{SLOT} and DIFS are respectively 31, 10us, 20us and 50us. Assuming a data rate of 11Mbps, T_{VOICE} and T_{ACK} are comprised of the component times shown in Table I.

In order to assess the tightness of this upper bound, we compared the value of the upper bound with the capacity obtained by simulation (maintaining the assumptions ii) and iii) above). These results are shown in Table II (the calculated upper bound is in parentheses). Note that the capacity values shown are independent of any quality requirement: no packet loss occurs at or below capacity, while packet loss due to queue overflow at the Access Point for higher loads is extremely high (typically 10% and higher).

These results support our assertion that T_{IDLE} tends to 0 as the load increases, and also show that the effect of collisions on capacity is small. Indeed, we observed that the percentage of transmissions involved in a collision ranged from around 1.5% to 4% for the AP's transmissions and from 2% to 9% for the other stations' when the network is operating at maximum capacity. Because of the low rate of collisions, no frames were dropped by the MAC due to an excessive number of retransmissions (using the default retry limit of 7).

These results also highlight the effect of the large overhead at the MAC and physical layers. The majority of this overhead comprises the transmission time of the PLCP header and preamble (accounting for over 50% of the total time) and idle time, when no station is transmitting (typically 20-30% of the time). As a result, the capacity is affected more by the rate of packet generation R (and hence the amount of voice data per packet) than by the bit rate of the encoder. For example, although the output bit rate of a G.711 encoder is eight times that of a G.729 encoder, the reduction in capacity when G.711 is used is less than 50%.

Note that the use of the short PLCP preamble can significantly reduce this overhead, leading to an increase in capacity: our simulations showed an increase of around 25-50%. However, since this capability is optional, it cannot be relied upon, and for the remainder of the paper we assume the use of the long preamble only.

Garg and Kappes [6] performed an analysis using a simplifying approximation of the MAC protocol. The results in [6] correspond exactly to the values obtained here by simulation; however, Garg and Kappes validated their analysis, by carrying out a measurement study, only for the G.711, 10ms case.

III. THE EFFECTS OF DELAY CONSTRAINTS, NON-IDEAL CHANNEL CONDITIONS AND QUALITY REQUIREMENTS

In the previous section and throughout [6], it was assumed that the playout deadline was met by all packets and that no packets were received with error(s). The results showed that, to maximize capacity under such assumptions, G.729 is to be preferred over G.711 and that packets should contain as much voice data as possible. However, when considering the impact of poor channel conditions, call quality requirements and/or delay constraints, many further issues must be considered.

Because of the coding and modulation used in 802.11b, packet error rates are highly dependent on packet length: long packets are more susceptible to error than short packets. As the channel quality deteriorates, more retransmissions are required, resulting in a lower throughput-constrained capacity (more time is required per successful transmission), and higher per-packet delays.

Furthermore, although G.729 has been shown to permit greater capacity, the quality of such calls is limited by the encoding scheme's intrinsic MOS of 3.65, compared to 4.15 for G.711 [2]. The G.729 algorithm also requires a 5ms look-ahead, delaying packets by an additional 5ms compared to G.711.

Finally, for a given encoding scheme, the packetization delay is higher for larger packets. The loss of such packets is also harder to conceal (using PLC) than smaller packets [10].

In this section we first evaluate the effect of a delay constraint in the context of ideal channel conditions, and then present results showing the combined effects of delay constraints and non-ideal channel conditions.

A. The Delay Constraint

The delay constraint is set by the playout buffer at the receiver, which drops packets that have incurred an excessive end-to-end delay and arrive after their playout deadline. Throughout this section, we assume a fixed playout deadline for packets, allowing a maximum of 150ms end-to-end delay. This allows for the encoder's algorithmic and packetization delay, the propagation delay through the wired network, queuing delay at the wireless network interface, channel access delay and propagation time over the wireless medium.

TABLE III

MAXIMUM LOSS RATES (%) FOR MINIMUM MOS REQUIREMENTS

	Minimum MOS 4.0	3.6
G.711 (10ms)	1	4.9
G.729 (10ms)	N/A	0.33
G.711 (20ms)	1	3
G.729 (20ms)	N/A	0.19

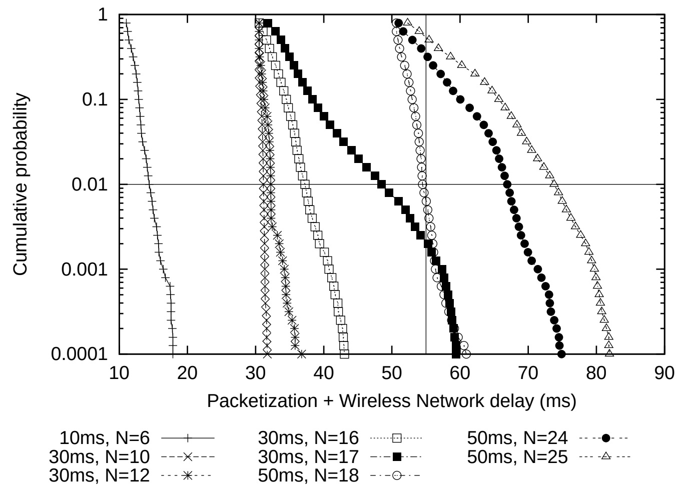

Fig. 2. CCDF for delay for G.711 with various packet sizes and number of calls

Since an end-to-end delay of 150ms causes a negligible decrease in quality [11], the MOS score is reduced from the intrinsic value only by the effect of loss. In [2], the effects of loss on MOS are given for G.729 and G.711, for 10ms and 20ms packets. The loss rates corresponding to MOS values of 3.6 and 4.0 are shown in Table III. We apply the values for 20ms packets given in [2] to 30ms and 50ms packet sizes considered here[1].

In Figs 2 and 3, we plot the complementary cumulative distribution function (CCDF) for the sum of the delay incurred within the wireless network and the packetization delay for several scenarios with ideal channels, using G.711 and G.729 respectively. On these plots, the y-axis gives the probability of a packet incurring a delay greater than the value on the x-axis. (We plot here the delay CCDF for packets transmitted from the wired nodes to the wireless stations, since the delay in this direction is almost always higher than in the opposite direction, due to queuing at the AP.)

These figures can be used to determine the fraction of packets that will be dropped at the playout buffer due to late arrival as a function of the delay budget allocated for packetization (including the 5ms look-ahead required by G.729) and queuing and transmission in the wireless network (assuming the delay in the wired network to be constant). For example, using G.711

[1]Because PLC is less effective for longer packets than short ones, the corresponding loss requirement would be stricter for 30ms and 50ms than for 20ms packets, for the same quality. Our approach will therefore lead to a somewhat optimistic value for capacity (from the point of view of PLC) when using these packet sizes.

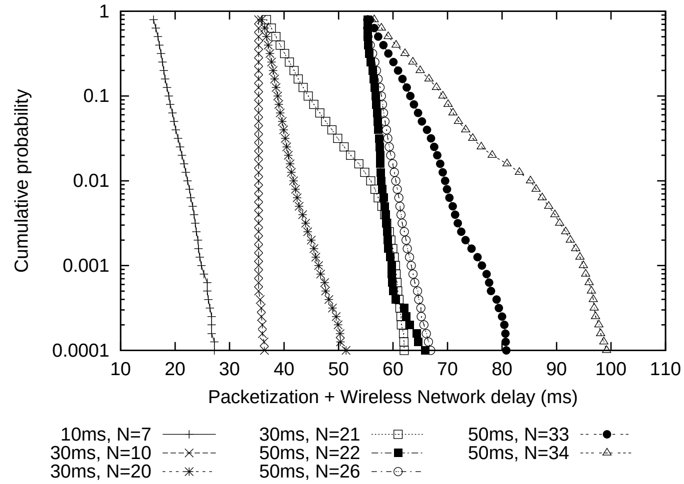

Fig. 3. CCDF for delay G.729 with various packet sizes and number of calls

TABLE IV

CAPACITY OF AN 802.11B NETWORK (ASSUMING IDEAL CHANNEL)

Delay budget (ms)	MOS = 3.6 G.711	MOS = 3.6 G.729	MOS = 4.0 G.711
≤ 10	0	0	0
20	6 (10)	6 (10)	6 (10)
30	11 (20)	11 (20)	11 (20)
40	16 (30)	16 (30)	16 (30)
50	17 (30)	20 (30)	17 (30)
60	23 (50)	23 (50)	22 (50)
70	25 (50)	31 (50)	24 (50)
80	25 (50)	33 (50)	25 (50)
≥90	25 (50)	34 (50)	25 (50)

with 50ms packets, and with a delay budget of 55ms, 18 concurrent calls can be supported with a loss rate of less than 1% (as indicated on Fig. 2).

Using the loss requirements specified in Table III together with the delay statistics, we can evaluate the capacity for various packet sizes and delay budgets. Clearly there is a trade-off between the throughput constraint (favoring larger packets) and the delay constraint (favoring smaller packets). In Table IV, we show the capacity that may be attained by selecting the optimum packet size appropriate to the wireless network and packetization delay budget. The optimum packet size (in ms) is given in parentheses following the capacity value. (Note that G.729 cannot be used for a minimum call quality MOS requirement of 4.0.)

Considering a MOS requirement of 3.6, it can be seen from the table that G.711 and G.729 provide similar capacity where the delay budget is less than 70ms: in general, G.711 is constrained by throughput, G.729 by delay. In particular, the 5ms look-ahead time required by a G.729 encoder makes the delay constraint particularly stringent in such scenarios. Where the delay budget is higher, throughput constraints dominate in both cases, and G.729 provides a much greater capacity.

Comparing the capacity for MOS requirements of 3.6 and 4.0, we are clearly constrained to the lower capacity of G.711 for the higher MOS value. However, comparing the capacity using G.711 with different quality requirements, we observe

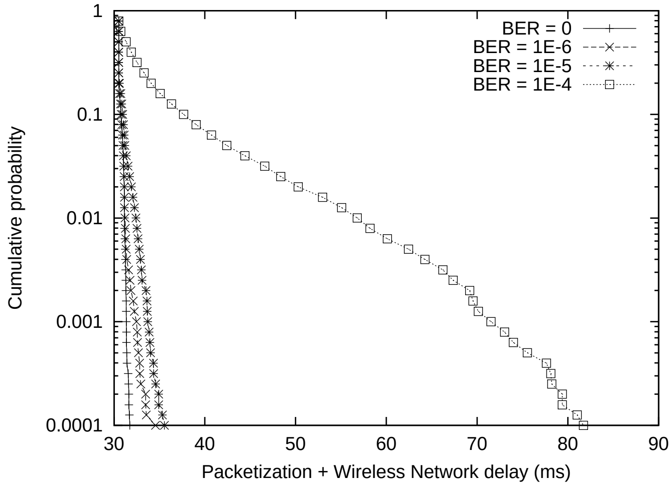

Fig. 4. CCDF for delay for 30ms G711 packets, 10 calls

TABLE V
THROUGHPUT BOUNDED CAPACITY FOR VARIOUS BER VALUES

	Voice Data per frame (ms)							
	G.711				G.729			
BER	10	20	30	50	10	20	30	50
0	6	12	17	25	7	14	21	34
10^{-6}	6	12	17	25	7	14	21	34
10^{-5}	6	12	16	24	7	14	20	33
10^{-4}	5	9	12	15	6	12	18	29
2×10^{-4}	4	7	8	7	5	11	16	25

that the difference is minimal even though the maximum loss rate is reduced from 4.9% to 1%.

B. Non-ideal Channel Conditions

To assess the effect of non-ideal channel conditions, we use a simplified channel model, represented by a constant Bit Error Rate (BER). We assume that the channels between all pairs of nodes are subject to this BER value, and that all bit errors occur independently. We maintain the abstraction of the channel at the BER level, rather than consider the Signal-to-Noise Ratio (SNR) at the receiver, since the SNR to BER mapping is implementation-dependent. Finally, we assume that, unless a collision occurs, the PLCP preamble and header are received and decoded correctly.

As has been described, poorer channel conditions, leading to higher BER values, cause an increase in per-packet delay. To illustrate this, we plot in Fig. 4 the delay CCDF for a network carrying 10 calls using 30ms G.711 packets for BER values ranging from 0 to 10^{-4}. The effect on the throughput constraint is shown in Table V, which lists the throughput-constrained capacity (i.e. without any delay constraint) for a minimum MOS requirement of 3.6. As in the ideal channel case, packet loss only occurs at the AP's interface queue.

We observe that for BER $\leq 10^{-6}$, the packet error rate for both G.711 and G.729 is so low that the difference in capacity between such a channel and an error-free channel is negligible. For $10^{-6} <$ BER $\leq 2 \times 10^{-4}$, the capacity decreases by varying degrees, depending on the amount of voice data contained in a packet (longer packets are more susceptible to errors, hence are more likely to require retransmission). For BER $\geq 10^{-3}$, not even one voice call can be supported by the network due to the very high packet error rate, which causes the MAC to retransmit frames so often that the probability of a frame being dropped due to excessive retransmissions becomes significant.

TABLE VI
CAPACITY OF AN 802.11B NETWORK (BER = 10^{-4})

Delay budget (ms)	MOS = 3.6 G.711	MOS = 3.6 G.729	MOS = 4.0 G.711
≤ 10	0	0	0
20	4 (10)	5 (10)	4 (10)
30	7 (20)	6 (10)	5 (10)
40	9 (30)	11 (20)	7 (20)
50	10 (30)	16 (30)	9 (30)
60	11 (30)	17 (30)	10 (30)
70	11 (30)	21 (50)	10 (30)
80	11 (30)	27 (50)	11 (30)
90	11 (50)	28 (50)	11 (30)
100	13 (50)	29 (50)	11 (30)
110	13 (50)	29 (50)	11 (50)

As we have already seen for the ideal channel case, adaptation to the delay constraint can be used to maximize capacity for a given quality (MOS) requirement. If stations can adapt jointly to the channel conditions and delay budget, the network capacity can be similarly maximized for non-ideal channel scenarios. In Table VI, we show the maximum capacity and corresponding packet size for the case of BER = 10^{-4} for delay budgets up to 110ms. Figs 5, 6 and 7 show the maximum capacities that can be achieved using G.711 and G.729 with optimum adaptation, for MOS requirements of 3.6 and 4.0 for the range of BER values considered.

From these figures we observe that, for a 3.6 MOS requirement, G.729 provides equal or greater capacity than G.711 in all cases. In particular, for a given BER, G.729 requires fewer retransmissions due to channel errors than G.711, since it generates smaller packets and has a correspondingly lower packet error rate. Although fragmentation can (in principle) be used to create smaller packets, and hence lower the effective packet error rate, its use in this application is limited by the 802.11 specification which requires a fragmentation threshold of no less than 256 bytes. In fact, we have observed through simulation that fragmentation cannot improve the capacity beyond that achievable with optimum packet size selection.

For most channel conditions (BER $< 10^{-4}$) the delay is dominated by the packetization delay alone; then, the optimum packet size can be determined without knowledge of the state of the channel by selecting the highest packet size smaller than the delay budget. However, as can be seen from Table VI, this is not the case for higher BER values.

C. Discussion

In presenting our results so far, we have assumed a fixed playout scheme, resulting in an end-to-end delay of 150ms. However, this value may not always be the most suitable choice. In [11], the additional degradation in MOS that would result if this delay were increased by less than 50ms is shown to be very small, e.g. increasing the allowed delay by 20ms to 170ms results in MOS degradation of less than 0.1.

Since we have shown that the capacity is highly sensitive to the delay constraint, relaxing the delay constraint by increasing the playout deadline would allow more calls to be supported with minimal degradation in MOS. In order to maximize call capacity and quality, an adaptive playout scheme would then be required, the details of which are outside the scope of this paper.

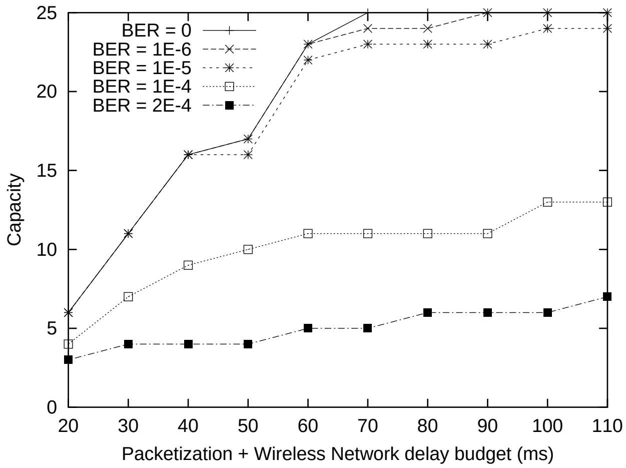

Fig. 5. Number of G.711 calls, MOS $\geq$ 3.6, that can be supported

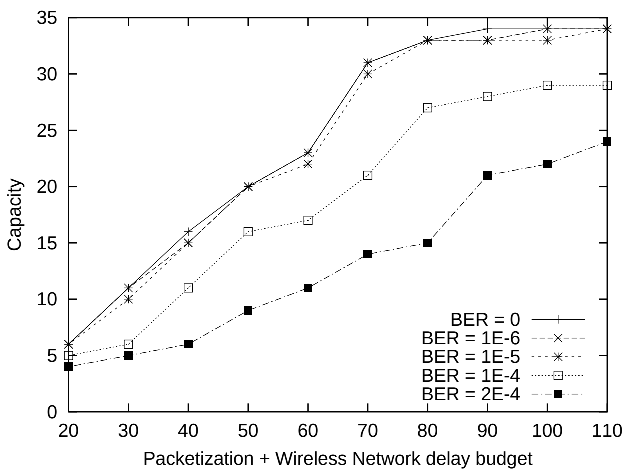

Fig. 6. Number of G.729 calls, MOS $\geq$ 3.6, that can be supported

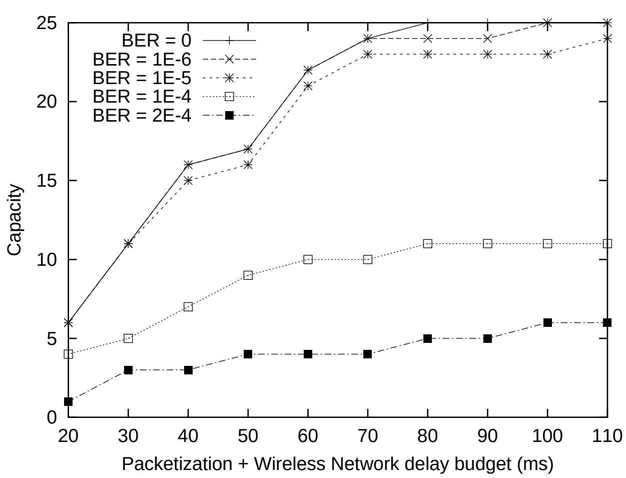

Fig. 7. Number of G.711 calls, MOS $\geq$ 4.0, that can be supported

IV. Conclusion

In this paper we have evaluated an upper bound on the capacity of an IEEE 802.11b network carrying voice calls, and found it to be tight in scenarios where channel quality is good and delay constraints are weak or absent.

We have then shown that capacity is highly sensitive to the delay budget allocated to packetization and wireless network delays. Concerning non-ideal channel conditions, we have shown that capacity is very close to that in an error-free channel for BER values of less than 10^{-5}; in more adverse conditions capacity is reduced considerably, and is zero for channels with BER $\geq 10^{-3}$. By selecting the packet size appropriately given the delay budget and channel conditions, the capacity can be maximized; furthermore, in the majority of cases optimum packet size selection can be made without knowledge of the channel conditions.

Throughout, the use of G.729 has been shown to allow greater capacity than the use of G.711, unless a voice quality corresponding to a MOS of greater than 3.65 is required, in which case G.729 cannot be used.

References

[1] ISO/IEC and IEEE Standard "Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications" (as supplemented), 1999

[2] ETSI Technical Report "Actual measurements of network and terminal characteristics and performance parameters in TIPHON networks and their influence on voice quality", ETSI TR 101 329-6 V2.1.1 (2002-02) 2002

[3] B. P. Crow, I. Widjaja, J.G. Kim and P.T. Sakai "IEEE 802.11 Wireless Local Area Networks" IEEE Communications Magazine, September 1997

[4] A. Köpsel and A. Wolisz "Voice transmission in an IEEE 802.11 WLAN based access network" Proc. of ACM Workshop on Wireless Mobile Multimedia, July 2001

[5] N. Smavatkul, Y. Chen and S. Emeott "Voice Capacity Evaluation of IEEE 802.11a with Automatic Rate Selection", Proc. of GLOBECOM 2003

[6] S. Garg and M. Kappes "Can I add a VoIP call?", IEEE Int. Conf. on Communications 2003 (ICC '03)

[7] S. Garg and M. Kappes "An experimental study of throughput for UDP and VoIP traffic in IEEE 802.11b networks", IEEE Wireless Communications and Networking 2003

[8] A. Markopoulou, F. Tobagi and M. Karam "Assessment of VoIP Quality over Internet Backbones", in Proc. IEEE INFOCOM 2002

[9] NS-2 Network Simulator (v2.1b9) http://www.isi.edu/nsnam/ns/

[10] J. Rosenberg "G.729 Error Recovery for Internet Telephony", Columbia University Computer Science Technical Report CUCS-016-01, 2001

[11] ITU-T Recommendation G.107 "The Emodel, a computational model for use in transmission planning" 1998