

No Ack in IEEE 802.11e Single-Hop Ad-Hoc VoIP Networks

Jaume Barceló, Boris Bellalta, Anna Sfairopoulou, Cristina Cano, Miquel Oliver

Abstract This paper analyzes the impact of the No Ack policy in VoIP ad-hoc single-hop IEEE 802.11e networks. The No Ack policy consists on suppressing the MAC layer acknowledgement packets. This option dramatically reduces the duration of MAC frames and thus increments the maximum number of VoIP flows that can coexist in the network. The negative side is an increased packet loss rate and the consequent drop of voice quality as perceived by the users. The article presents a model to quantify the benefits of suppressing acks and assess the viability of the No Ack policy. It follows an argumentation to identify which is the best codec to combine with acknowledgement suppression and in which conditions the usage of No Ack can be beneficial. Finally, it is suggested to modify the stations to switch to No Ack policy as the network approaches congestion.

1 Introduction

Both IEEE 802.11 [1] networks and voice-over-IP (VoIP) are mature technologies and have been widely adopted for personal and enterprise use. Calls in a single-hop ad-hoc network are less frequent, since they require the proximity of the caller and callee. Nevertheless, when the proximity requirement is satisfied, the calls can be established free and in infrastructure-less scenarios.

Even in areas equipped with infrastructure, the use of ad-hoc communications is beneficial. The use of infrastructure mode implies that each packet has to be relayed by the access point, and therefore it doubles the number of transmissions in the network. This is a waste of the scarce radio resources. Moreover, changing to infrastructure mode also increments the end-to-end delay. And finally, the access point usually represents the bottleneck of the infrastructure network having to receive and transmit flows from and to all the stations. This is illustrated by Fig. 1.

Barceló, Bellalta, Sfairopoulou, Cano and Oliver
Universitat Pompeu Fabra, Passeig de Circumval.lació 8 e-mail: jaume.barcelo@upf.edu

The station-to-station communication is possible in infrastructure scenarios using Direct Link Setup (DLS) thanks to the quality of service (QoS) amendment to the standard [2].

However, the combination of VoIP and IEEE 802.11 is highly inefficient. The number of calls that can coexist in a IEEE 802.11 VoIP network is an order of magnitude lower than what could be expected when comparing the rates of the network to the rates of the codecs. The reasons for such impairment have already been analyzed [13] and can be summarized as:

- The time wasted in contention (i.e. the channel is idle and all the stations are backing off with packets ready to transmit).
- The physical preambles and the overhead placed by the different layers of the protocol stack.
- Each packet needs to be separately acknowledged.

The high overhead per transmitted packet seriously penalizes the transmission of small payload packets, which is in fact the case of VoIP packets. The problem can be alleviated in four different ways:

- Header compression, to reduce overhead. The 40-byte RTP/UDP/IP header can be compressed to 2-7 bytes [5, 4, 6, 11].
- Packet aggregation. Combine different packets to produce only one packet that contains all the data [13].
- Use only one ack packet to acknowledge a set of packets (Block-Ack)[12].
- Refrain from sending acknowledgement packets (No Ack).

While the first three solutions have been extensively studied, the fourth is still an open research issue. The possibility of not sending acknowledgements is available in the IEEE 802.11e [2] extension for quality of service. Intuitively, this option will decrease the packet delay and increase the number of simultaneous calls. However, it is expected that this solution will also increase the packet loss ratio which negatively affects the QoS.

This article presents a simple model to analyze a VoIP over IEEE 802.11 scenario and to assess which is the impact of using the IEEE 802.11e No Ack option.

After this introductory section, the rest of the paper is outlined as follows. Section 2 offers a brief description of the scenario and the protocols involved, details




Fig. 1 This figure compares ad-hoc (left) to infrastructure-based (right) VoIP calls in a IEEE 802.11 network.

the causes of the inefficiency, and presents a model that can be solved numerically. Section 3 compares the collision probabilities and the maximum number of flows to those obtained when acks are suppressed. It also takes some perspective on the results and analyzes which are the real benefits of suppressing the acks. It details in which conditions it would be desirable to apply ack suppression, which are the more appropriate codecs, and give some implementation clues for a real-world deployment. Finally, Section 4 concludes the paper.

2 Description and Model of IEEE 802.11 Voice Networks

A VoIP [8] application digitalizes the voice, then uses a codec to compress it and outputs a fixed-rate bitstream (See Fig. 2). Popular codecs include G.711 (64 Kbps) and G.729 (8 Kbps). Each codec has its own properties, such as the offered bitrate, the computational complexity of codification/decodification and the offered Mean Opinion Score (MOS). MOS is a measure of the voice quality as perceived by the user [7].

The packetizer collects the data from the encoder and periodically generates a voice packet. The packetization interval has a deep impact in the overall network performance [3]. Increasing the packetization interval introduces additional delay but drastically reduces the number of packets traversing the network. A choice of codec and packetization interval fixes the length of the payload of each packet. As an example, a G.711 codec combined with a 20ms packetization interval means a payload of 160 bytes per packet. The transport layer and network layer consist on the addition of RTP, UDP, and IP headers, which add up to 40 bytes of overhead.

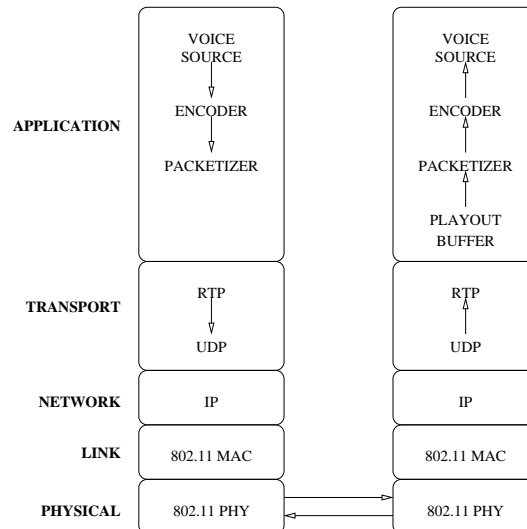


Fig. 2 The figure presents the complete protocol stack in implementing VoIP over IEEE 802.11.

Our proposal is to suppress acknowledgement packets at the link layer, which is an option contemplated in IEEE 802.11e amendment to the standard. In any case, the link layer also introduces its own headers. Finally the physical layer introduces preambles and headers that are transmitted at a fixed rate (physical rate), while the rest of the packet is transmitted at a speed that depends on channel conditions (data rate).

The time required to transmit a packet is:

$$T_{tx} = T_{plcp} + \frac{MAC + IP + UDP + RTP + VOICE}{DATA_{RATE}} \quad (1)$$

Where T_{plcp} is the duration of the transmission of the Physical Layer Convergence Protocol (PLCP) preamble and header at the physical bitrate. MAC, IP, UDP and RTP represent the length (in bits) of the respective headers. $VOICE$ is the length (in bits) of the payload. $DATA_{RATE}$ is the bitrate at which data is transmitted. The latter depends on channel conditions, but in this work it is assumed to be 11 Mbps for all the stations.

The duration of a successful transmission slot is computed as follows:

$$T_s = T_{tx} + SIFS + T_{plcp} + T_{ACK} + DIFS \quad (2)$$

where T_{ACK} is the duration of the acknowledgement packet. The duration of a collision slot is $T_c = T_{tx} + EIFS$, where $EIFS$ is the duration of the Extended Inter-Frame Space

Fig. 3 shows a complete successful slot. The shaded area correspond to the voice payload.

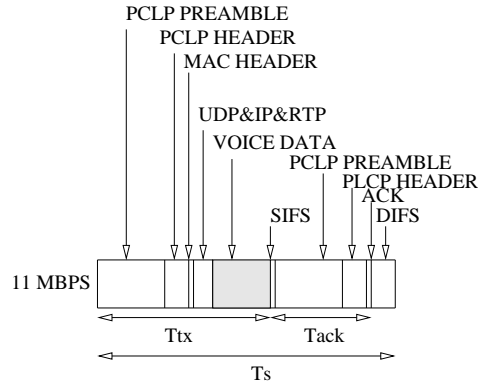


Fig. 3 The figure shows the a successful transmission slot. It is evident that the voice payload represents only a small fraction of the total. Suppressing the acknowledgement would reduce the duration of the slot significantly.

2.1 Performance Impairment

The time that is used to transmit actual voice data ($VOICE/DATA_{RATE}$) compared to the time it takes a successful transmission slot (T_s) can be taken as a measure of efficiency. Depending on the choice of codec, packetization interval and data rate, this efficiency will vary. Table 1 presents a summary in which the efficiency is computed for various combinations. The last column of the table shows the efficiency improvement for the hypothetical case in which the acks are suppressed.

Table 1 Efficiency (T_{tx}/T_s) with and without acknowledgements.

Case	With acks	Without acks	Improvement
G711, 20ms, 11Mbps	0.18637	0.28231	51%
G711, 10ms, 11Mbps	0.10276	0.16436	59%
G729, 20ms, 11Mbps	0.027836	0.046866	68%
G711, 20ms, 2Mbps	0.46037	0.54329	18%

To obtain the efficiency improvement detailed in the last column of Table 1, the usage of No Ack is required. When the No Ack policy is used by a station, there is no MAC-level recovery, and the transmission reliability is reduced. Hence the standard [2] recommends to use this policy only with some other additional protective mechanisms. However, since VoIP applications can accept a certain level of packet loss, in the following section we will study the consequences of using No Ack without such protective mechanisms.

2.2 Modelling 802.11 with acks

Assume that n mobile stations participate in an ad-hoc single-hop IEEE 802.11e network. Each station generates a VoIP flow, which is characterized by the periodical transmission of a short fixed-size packet. ρ is the load that a station offers to the network.

Assume also that the network is uncongested and the MAC queues do not fill up. Nevertheless, some packets may be discarded after reaching the maximum retransmission limit (R). Thus the actual load successfully transmitted by the network is:

$$r = \rho(1 - p_{cc}^{R+1}). \quad (3)$$

where p_{cc} is the probability that a collision occurs when the station attempts a submission.

The successful transmitted bitrate can also be calculated as the amount of data in a packet (l bits) multiplied by the successful transmission probability of station i ($p_i(s)$), and divided by the average duration of a slot.

$$r = \frac{p_i(s) \cdot l}{p(s)T_s + p(c)T_c + p(e)T_e}. \quad (4)$$

Where $p(s)$, $p(c)$ and $p(e)$ are the probabilities that any given slot is successful, collision and empty, respectively. Manipulating Eq. 4 and plugging in Eq. 3 we obtain:

$$p_i(s) = \frac{r \cdot (p(s)T_s + p(c)T_c + p(e)T_e)}{l}. \quad (5)$$

$$p_i(s) = \frac{\rho(1 - \rho^{R+1}) \cdot (p(s)T_s + p(c)T_c + p(e)T_e)}{l}. \quad (6)$$

The probability that station i successfully transmits in a given slot is the probability that i transmits while the other $n - 1$ stations remain silent. It can be expressed as a function of τ (the probability of a transmission attempt) and n the number of flows.

$$p_i(s) = \tau(1 - \tau)^{n-1}. \quad (7)$$

The probability that one and only one station transmits in a given slot is

$$p(s) = n \cdot \tau(1 - \tau)^{n-1}. \quad (8)$$

The probability that no station transmits is

$$p(e) = (1 - \tau)^n. \quad (9)$$

And the probability of a collision is

$$p(c) = 1 - p(s) - p(e). \quad (10)$$

The probability that a packet collides, conditioned to the probability that station i is attempting a transmission is:

$$p_{cc} = 1 - (1 - \tau)^{n-1}. \quad (11)$$

Eqs. 7 - 11 can be substituted into 6. The resultant equation has only one unknown variable (τ) and can be solved using numerical methods.

$$\begin{aligned} l \cdot \tau(1 - \tau)^{n-1} = \\ \rho((1 - (1 - (1 - \tau)^{n-1})^{R+1}) \cdot (n \cdot \tau(1 - \tau)^{n-1}) \cdot T_s + \\ + (1 - n \cdot \tau(1 - \tau)^{n-1} - (1 - \tau)^n)T_c + (1 - \tau)^n \cdot T_e). \end{aligned} \quad (12)$$

Once the transmission probability τ is computed, it can be used to obtain the rest of performance metrics that depend on it. Eq. 12 is also useful to figure out the theoretical maximum number of active flows. If Eq. 12 converges, it means that the computed scenario is feasible. Otherwise, the number of flows is excessive and must be reduced.

The model described above can also be used in the case in which acks are suppressed. It is a special case in which $T_s = T_{tx} + DIFS$ and $R = 1$.

3 Performance Analysis

Using Eq. 11 from the previous section, the conditioned collision probability can be plotted (Fig. 4). In this example the codec is G.711, the packetization interval 20ms and the $DATA_{RATE}$ is 11Mbps. NS2 [10] is used to validate the results.

The price to pay for increasing the capacity of the network in terms of feasible number of flows is an increased packet loss probability. A packet is lost if it suffers a collision each time that a transmission is attempted. The packet loss ratio is:

$$P_{loss} = P_{cc}^R. \quad (13)$$




Fig. 4 It can be observed in the figure that the ack suppression leads to lower collision probability and increased number of concurrent flows.

If the acknowledgements are omitted, each packet is transmitted only once and therefore $P_{loss} = P_{cc}$. Figure 5 presents P_{loss} in a logarithmic plot.

The ultimate goal of the performance tuning of the IEEE 802.11e network is to admit the maximum number of calls with acceptable QoS. Fig. 4 shows that suppressing the acknowledgement packets significantly increases the number of VoIP flows that can coexist in the network, before reaching the congestion condition. The congestion is characterized by MAC queues building up and the packet loss due to queue overflow reaching values of 10% and higher [9]. Under this condition, the quality of all calls is unacceptable.

Given that the network is uncongested and no packets are lost due to MAC queue overflow, there are still two aspects that can threaten the quality of the calls: delay (and jitter) and packets loss due to collisions.

Delay and jitter negatively affect the MOS of the calls. Generally is considered that delays under 150ms are well suited for all user applications, and that up to 400ms are acceptable for international calls. The main contributors of the end-to-end delay are the encoding delay (about 10ms for G.729 and even less for G.711), the packetization delay (20ms in all the examples used throughout this article) and the jitter buffer (typically 60ms). The mean delay introduced by the (single-hop) network in non-congested conditions is well below 5ms [3]. Hence the total delay does not pose the MOS at risk.

The other threat for the MOS is the packet loss due to collisions. As can be observed in Fig. 5, the acknowledgement&retransmission mechanism drastically




Fig. 5 This logarithmic plot shows that the packet loss probability is significantly higher when acknowledgements are suppressed.

reduces the packet loss due to collisions. Therefore, the packet loss caused by collisions is negligible when acks are used and the MOS remains unaffected. This is not the case when the acks are suppressed.

Without acks, there is no chance of retransmission and the probability of losing a packet is P_{cc} . G.711 and G.729 react very differently to the loss of packets. While G.711 is quite robust and can accept up to 5% of losses and still deliver acceptable quality, G.729 behaves badly for packet losses as low as 3%. This means that G.711 could accept up to 26 flows (13 calls), G.729 allows only 20 flows (10 calls). The alleged advantage of G.729, its lower rate, is overcome by the excessive overhead mentioned in section 2.

3.1 Implementation Issues

Apparently, the new feature of No Ack introduced by IEEE 802.11e is of limited use, since it does not increment the number of acceptable quality calls. In our opinion, there is only one case in which it would make sense to suppress acks: when the network becomes congested. At this point, in which calls are dropped due to excessive delay and packet loss, switching to No Ack would allow the users to maintain low-quality calls.

To implement this switch in a distributed fashion, each station should constantly monitor its MAC queue. As soon as the queue builds up (a symptom of congestion), the stations should set the QoS control subfield of the packets to No Ack, until the queue returns to its previous empty state.

The applicability of this mechanism is reduced to collision-limited scenarios. If packet losses are due to low SNR and the stations react by suppressing acknowledgements, the underlying problem will remain. Even worse, the suppression of acknowledgements would prevent the data-rate fallback to a more robust modulation.

4 Conclusion

This paper analyzes the applicability of a new feature introduced to wireless LANs by the standard amendment IEEE 802.11e that consists on suppressing MAC layer acknowledgements. The scenario under study is a VoIP ad-hoc network. This kind of networks suffers from an accentuated inefficiency problem, mostly due to the requirement of separately acknowledging each MAC frame. Additionally, VoIP applications can tolerate a certain number of packet losses. Hence we deemed this networks appropriate to benefit from the suppression of the acks.

A model that permits the quantification of the benefit of suppressing the acks is presented. Using this model, the collision probability is computed for an increasing number of flows. The results shows that by suppressing the acks, the collision prob-

ability drops and the number of VoIP flows that can be simultaneously allocated in the network grows. These results are validated by means of simulation.

The negative part of suppressing the acks is an increased packet loss rate. This is the metric that actually limits the maximum number of calls when the No Ack policy is used. The codec of choice to combine with the suppression of acks is G.711, since it admits a greater packet loss than G.729. The maximum number of calls that can be achieved using the G.711 codec and suppressing the acknowledgements is the same as with acknowledgements.

However, when increasing the number of calls over that maximum – that we have computed to be 28 flows, 14 calls – the behaviour of the network depends on the policy applied. If acknowledgements are sent as usual, the network congests and calls are dropped. If we opt for a No Ack policy, the perceived quality of the call drops below a MOS of 3.5 and some users will complain about unsatisfactory quality. However, in the latter case, network congestion is prevented.

Finally, it is suggested a distributed mechanism to switch to No Ack policy when the network approaches congestion.

References

1. 802.11, I.S.: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications. ANSI/IEEE Std 802.11 (1999 Edition (Revised 2003))
2. 802.11e, I.S.: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications; Amendment: Medium Access Control(MAC) Quality of Service Enhancements. IEEE Std 802.11e (2005)
3. Barceló, J., Bellalta, B., Cano, C., Sfairopoulou, A.: VoIP Packet Delay in Single-Hop Ad-Hoc IEEE 802.11 Networks. *Wireless On demand Network Systems and Services, 2008 IEEE/IFIP The Fifth Annual Conference on* **1** (2008)
4. Casner, S., Jacobson, V.: Compressing IP/UDP/RTP Headers for Low-Speed Serial Links. RFC 2508 (Proposed Standard) (1999). URL <http://www.ietf.org/rfc/rfc2508.txt>
5. Degermark, M., Nordgren, B., Pink, S.: IP Header Compression. RFC 2507 (Proposed Standard) (1999). URL <http://www.ietf.org/rfc/rfc2507.txt>
6. Engan, M., Casner, S., Bormann, C.: IP Header Compression over PPP. RFC 2509 (Proposed Standard) (1999). URL <http://www.ietf.org/rfc/rfc2509.txt>. Obsoleted by RFC 3544
7. ETSI, T.: 101 329-6 Actual measurements of network and terminal characteristics and performance parameters in TIPHON networks and their influence on voice quality (2001)
8. Goode, B.: Voice over Internet protocol(VoIP). *Proceedings of the IEEE* **90**(9), 1495–1517 (2002)
9. Hole, D., Tobagi, F.: Capacity of an IEEE 802.11 b wireless LAN supporting VoIP. *Communications, 2004 IEEE International Conference on* **1** (2004)
10. The network simulator ns-2, v2.30. <http://www.isi.edu/nsnam/ns>
11. Pries, R., Maeder, A., Staehle, D.: Do we need Header Compression for VoIP in Wireless LANs? *12th EUNICE Open European Summer School 2006* pp. 77–82 (2006)
12. Tinnirello, I., Choi, S.: Efficiency analysis of burst transmissions with block ACK in contention-based 802.11 e WLANs. *Communications, 2005. ICC 2005. 2005 IEEE International Conference on* **5** (2005)
13. Yun, S., Kim, H., Lee, H., Kang, I.: 100+ VoIP Calls on 802.11 b: The Power of Combining Voice Frame Aggregation and Uplink-Downlink Bandwidth Control in Wireless LANs. *Selected Areas in Communications, IEEE Journal on* **25**(4), 689–698 (2007)