

## RESEARCH ARTICLE

**E-MAC: an elastic MAC layer for IEEE 802.11 networks**Qing Wei<sup>1\*</sup>, Imad Aad<sup>2</sup>, Luca Scalia<sup>1</sup>, Jörg Widmer<sup>3</sup>, Philipp Hofmann<sup>4</sup> and Luis Loyola<sup>5</sup><sup>1</sup> DOCOMO Euro-Labs, Munich, Germany<sup>2</sup> Nokia Research Center, Lausanne, Switzerland<sup>3</sup> Insitute IMDEA Networks, Madrid, Spain<sup>4</sup> IABG mbH, Germany<sup>5</sup> SkillupJapan Corporation, Shinjuku-ku, Tokyo, Japan**ABSTRACT**

We present a system for real-time traffic support in infrastructure and *ad hoc* IEEE 802.11 networks. The proposed elastic MAC (E-MAC) protocol provides a distributed transmission schedule for stations with real-time traffic requirements, while allowing a seamless coexistence with standard IEEE 802.11 clients, protecting best-effort 802.11 traffic from starvation by means of admission control policies. Our scheduling decisions are based on an ‘elastic’ transmission opportunity (TXOP) assignment which allows for efficient wireless resource usage: whenever a real-time station does not use the assigned TXOP, the other real-time stations

57  
58  
59  
60  
61  
62  
63  
64  
65  
66  
67  
68  
69  
70  
71  
72  
73  
74  
75  
76  
77  
78  
79  
80  
81  
82  
83  
84  
85  
86  
87  
88  
89  
90  
91  
92  
93  
94  
95  
96  
97  
98  
99  
100  
101  
102  
103  
104  
105  
106  
107  
108  
109  
110  
111  
112

to implement and manage, which so far has hindered their deployment.

Besides the drawback of implementation complexity, strict guarantees rely on ‘medium reservations’ that are often done on a periodic basis: a given station reserves the channel for transmitting ( $P$  bytes) every  $T$  seconds. However, this periodicity may not match the real-time traffic pattern generated by the data sources. Consider for instance voice traffic, where the communication channel is typically idle 1/3 of the time. Several voice codecs, e.g., the ITU-T G.711  $\mu$ -Law codec [3], optimize bandwidth usage by applying silence suppression, leaving several reserved time-slots empty. One may try to adapt the period of the slot reservation to the voice codec pattern to optimize bandwidth efficiency or to reduce delays, but never both at a time. To reduce the average packet delay, short slot intervals must be used, but for such an over-provisioned reservation many of the reserved slots may be empty, therefore drastically reducing the efficiency, i.e., the ratio of the used slots to the total number of time slots. Conversely, increasing efficiency is likely to increase packet delays which is undesirable for real-time applications.

With these considerations in mind, we design an *elastic MAC* (E-MAC) protocol to provide strict QoS guarantees for real-time traffic, with ‘elastic’ reservations to allow for empty-slot reuse. A possible option is to adopt the centralized scheduling proposal HCCA of the 802.11e amendment. With HCCA, after the transmission of the beacon the access point (AP) reserves the channel for a specific amount of time, during which it polls the real-time stations. The polling is performed on the basis of Traffic Specifications (TSPECS), given by the stations through radio resource requests. This mechanism prevents the AP from scheduling clients that do not have real-time packets to transmit, while allowing it to order the polling list according to the specific traffic deadlines of the users. The remaining part of the superframe—the interval between two consecutive beacons—is left for legacy channel access contention. However, due to its implementation complexity, no HCCA-based APs can be found in the market. We design a protocol that meets the requirements for strict QoS guarantees and empty-slot reuse, for both infrastructure and *ad hoc* mode. It is compatible with existing 802.11 devices and deployed 802.11 networks and hotspots. To show the viability and efficiency of our approach, we implement it in a real testbed.

The paper is organized as follows. In Section 2, we give an overview of related work. In Section 3, we describe the design goals and the protocol in detail. The mathematical model is presented in Section 4. The protocol is evaluated both through simulations in Section 5 and in a testbed in Section 6. Finally, we discuss various issues and future work in Section 7 and conclude the paper in Section 8.

## 2. RELATED WORK

In the past decade, there has been considerable research in wireless networks with particular focus on QoS [4–6].

Existing approaches cover a wide range of applications, requirements, and assumptions; however, the lack of feasibility is a common drawback of many past works. With the availability of several open-source Linux-based 802.11 MAC drivers, developed thanks to the reverse-engineering work made by user communities [7–10], there is a significant increase of experimental work on QoS-oriented 802.11 solutions.

One common aspect of these works was the adoption of deterministic QoS provisioning mechanisms directly at the MAC layer. In particular, a lot of research targeted TDMA-based MAC protocols able to support PCF-like contention-free channel access without incurring the unpredictable delays and overhead of a centralized polling scheme [11–20]. The rationale behind this design choice lies in the deficiency statistical prioritization mechanisms supported in the 802.11e amendment. Apart from the inherent 802.11e limitations on achieving the expected QoS guarantees, it has been shown in References [21,22] that several commercially available WLAN devices do not exhibit standard compliant MAC behavior. Lower contention window sizes, absence of backoff mechanism, incorrect AIFS timing, NAV neglect, and so on have made the deployment of 802.11e QoS-oriented mechanisms infeasible in real-world scenarios, where fair bandwidth sharing between the devices is expected. Even for standard compliant network deployments, it has been shown in Reference [23] that IEEE 802.11e QoS-oriented enhancements can provide the desired delay performance only when the number of real-time traffic sources is very low. As shown in Reference [24], with a larger number of real-time traffic sources, the statistical prioritization offered by the 802.11e mechanism becomes insufficient to preserve the delay requirements. This has led to the common understanding that, in order to effectively support traffic prioritization over 802.11 legacy hardware, deterministic approaches have to be followed. Most of the proposed work explicitly designed non-802.11 compliant access schemes, whereas coexistence with legacy 802.11 traffic can be achieved at the cost of introducing a PCF-like reservation mechanism, together with its related inefficiencies [25]. Furthermore, as shown in Reference [26], the introduction of a contention-free period can also have the adverse effect of excessively delaying real-time traffic, thus requiring a more elastic traffic scheduling mechanism able to follow the traffic generation pattern. Zen *et al.* [27] propose a flexible scheduling scheme by adaptively segregating the real-time traffic and non-real-time traffic. However, that approach is still based on PCF and thus has some complexity issues.

The implementation of TDMA-like schemes on top of commodity 802.11 hardware requires modification of the driver source code: low-level 802.11 functions and parameters like the exponential backoff, the physical and virtual carrier sense, the slot-time duration and IFS size need to be modified, disabled, or reconfigured. In addition, the lack of central coordination and the distributed nature of these access schemes inevitably requires the introduction of

tailored synchronization mechanisms able to align the time slots of all clients to some time reference.

The authors of References [19,20] have built TDMA-like protocols and scheduling policies on the basis of the SoftMAC framework [16]. Peer synchronization is achieved by means of guard times between consecutive time slots and a wired connection to a central node is used to achieve a precision of 25  $\mu$ s.

In 2P [11], the authors provide a customized TDMA MAC protocol for interference mitigation in a multichannel environment. A node is in transmission mode for a specific time period that is globally known, and then explicitly notifies the end of its transmission period to each of its neighbors using marker packets. A receiving node waits for the marker packets from all its neighbors before switching over to transmission mode. In the event of a loss of a marker packet, a receiving node uses a timeout to switch into the transmission mode. The 2P protocol suffers from performance impairments in lossy environments, where marker packets can be easily lost. The same authors, in Reference [17], adopt a looser synchronization scheme tailored for lossy environments. The approach resembles the Network Time Protocol (NTP) principles, that implicitly correct the offset between the beginning of the transmission and the beginning of the time slot.

The authors in Reference [13] propose a TDMA MAC protocol able to exploit elastic TDMA transmission scheduling, thanks to an out-of-band synchronization mechanism able to achieve a precision of the order of a few microseconds.

In Reference [18], Rao and Stoica propose an overlay TDMA MAC layer on top of 802.11 hardware to overcome the typical 802.11 MAC performance impairments. They fix the slot size to 10 ms and use a leader node to generate the ‘clock’ to synchronize all the nodes in the network. Time stamps and latency estimation are appended to each packet header to compute the clock skew at each receiver.

Finally, Dechene and Shami [12] and Guo and Chiu [15] propose TDMA-based MACs explicitly tailored to achieve VoIP and video traffic improvements over the 2.4 GHz band. The TDMA scheme is similar to IEEE 802.16. A superframe structure, divided in uplink and downlink phases is used. A beacon packet informs the stations of the transmission schedule for the duration of the overall frame.

### 3. E-MAC PROTOCOL DESCRIPTION

This section describes the basic characteristics of the E-MAC protocol. For simplicity, the following description refers to a typical hotspot scenario, where all the traffic transits through an AP. However, the E-MAC protocol is independent of the 802.11 operation mode (*ad hoc* or infrastructure). We consider a network scenario in which  $n_{be}$  legacy IEEE 802.11 best-effort stations share the channel with  $n_{rt}$  E-MAC real-time stations. Of course, E-MAC real-time stations can also generate other traffic types at the same

time. The non-RT traffic from the E-MAC real-time station will be put into different queues and treated separately, just like another independent contending 802.11 station. Without losing generality, we focus on the real-time traffic from E-MAC real-time stations in this paper. The total number of active stations is thus  $n = n_{be} + n_{rt}$ . We start from a high-level overview of the E-MAC protocol procedures, and subsequently provide a more detailed discussion of the E-MAC protocol characteristics.

The E-MAC protocol is compatible with 802.11 standard compliant devices, ensuring inter-operability with legacy stations and APs.

#### 3.1. Overall E-MAC characteristics

The E-MAC protocol divides the channel access into two phases: a slotted TDMA-like access phase, available only to E-MAC enabled real-time (RT) stations, and a legacy 802.11 contention phase, available to all the contending stations and arbitrated according to the DCF access rules.

##### 3.1.1. Framing control and synchronization.

The length of the TDMA access phase and the legacy DCF access phase is regulated by a specific E-MAC station, which is referred to as the ‘Maestro station’. The Maestro station guarantees the loose synchronization of the E-MAC stations to the start of the contention-free phase and specifies the rules for admission control. This ensures a predictable level of fairness within the overall system: capping the traffic offered by RT stations during the contention-free period, and letting them contend fairly (using best-effort channel access parameters) with best-effort (BE) stations during the contention-based interval. The admission rules are used to divide the resource among the RT stations and guarantees a minimum length for the DCF phase which prevents low-priority best effort traffic from starvation.

##### 3.1.2. Scheduling and resource utilization.

One of the main features that differentiates the E-MAC protocol from similar channel access mechanisms like HCCA or PCF is the organization of the transmission schedule within the contention-free period. During the contention-free period, the transmission sequence is organized in a distributed manner by the RT stations.

Each E-MAC station overhears the admission rules as well as the highest sequence number  $S$  of the active E-MAC stations before joining the transmission. If this E-MAC station gets admitted according to the rule, it is assigned the sequence number  $S + 1$ . The sequence number decides the backoff time of each E-MAC station, which implicitly decides their transmission order.

The loose resource reservation *via* backoff and the distributed transmission schedule have an important impact on the resource utilization efficiency: different from the other reservation-based access schemes, which inevitably

waste the resources (slots) previously scheduled but subsequently not utilized, each MAC-enabled station can take over the transmission opportunity (TXOP) that her predecessor has skipped after short time interval (i.e., the difference of their respective backoff times), thus shortening the contention-free period and extending the duration of the contention-based period.

### 3.2. Frame period, contention-free period, and contention-based period

One round of contention-free and contention-based access together form what we refer to as frame period ( $T$ ). One frame period starts with the transmission of a Reserved Access Marker (RAM) by the Maestro station: the role of the RAM is to define the duration of the frame period, the rules for admission control (e.g., the maximum amount of allowed RT traffic from each RT station and the minimum reserved contention period) and synchronizing all the RT stations to the beginning of the contention-free period ( $t_{rt}$  in Figure 1, including the RAM itself). Supposing there is no idle time between  $t_{rt}$  and the contention period  $t_{be}$ ,  $T = t_{rt} + t_{be}$  (Figure 1). During the contention-free period, packets from different RT stations access the channel sequentially according to the agreed schedule (e.g., A, B, C, D in Figure 1). For simplicity, we assume each RT station is only allowed to send one packet in  $t_{rt}$ . If an RT station (e.g., B in Figure 1) misses its chance to send its packet, for example because it does not have packet to send, the next RT station in the schedule (e.g., C in Figure 1) takes over after waiting for an additional timeslot. After all RT stations transmit their admitted packets in the contention-free period, BE stations compete for access to the channel during the contention period.

The frame period can be configured according to different design choices. One can opt for a fixed frame length structure or for a dynamic frame length structure. It is also possible to use a fixed ratio  $R$  between the contention-free period and the contention-based period. Figure 2(a) and (b) shows the fixed and the dynamic frame period structure. For simplicity, in the remainder of the paper we use a fixed frame length and a fixed minimum contention period.

### 3.3. Self-organization and 'Maestro' station

Before starting its transmissions, a real-time station overhears the channel for a given amount of time (RAM timeout) in order to join a group that is already using E-MAC. The RAM packet is broadcasted to all the nodes every frame period  $T$ . It has the role of partitioning the entire frame




Figure 1. Conceptual model.




Figure 2. Fixed frame length (a) vs. dynamic frame length (b).

period in the contention-free and contention-based periods, synchronizing all the RT stations to the start of each new contention-free period. It also has the purpose of 'pushing' the best-effort stations to the contention-based period. To this end, the duration field of the RAM is set to a SIFS +  $(n_{rt} + 1)t_{slot}$ . Real-time stations do ignore the NAV field. The RAM is sent with the highest priority, i.e., after a SIFS time plus one IEEE 802.11 time slot  $t_{slot}$  (SIFS = 10  $\mu$ s and  $t_{slot} = 20 \mu$ s for 802.11b). This allows the Maestro node to deterministically take the control of the channel, preventing the best-effort stations from accessing it, as further discussed in Section 3.

The Maestro node maintains a table of all active real-time stations, including their MAC address, sequence number  $i$  (explained in subsection 3.4), and transmission time required for their real-time packets. Additionally, the time of the last real-time packet transmission of the respective station is stored. Note that other stations should also maintain such a table in case they become the Maestro (see subsection 3.7). Compared to the polling process of PCF, the cost to maintain such a table at the Maestro station can be ignored. The Maestro node broadcasts information about the total number of real-time stations  $n_{rt}$  and the total time required for transmitting all real-time packets including the RAM, which is  $t_{rt}$ . If all real-time stations transmit their packets, the total required transmission time is (*cf.* Figure 3)

$$t_{rt} = \text{SIFS} + t_{slot} + t_{ram} + \sum_{i=1}^{n_{rt}} (2\text{SIFS} + 2t_{slot} + t_{data,i} + t_{ack}) \quad (1)$$

The transmission duration of RAM ( $t_{ram}$ ), data packets ( $t_{data,i}$ ), and ACK ( $t_{ack}$ ) depends on the current channel rate and thus may vary. How to cope with variable data transmission times, different data rate requirements, and mobility is discussed in Section 7.




Figure 3. Basic system model.

To maintain compatibility with conventional 802.11 devices, the RAM is a normal 802.11 data frame with specific information in the payload. It is transmitted using MAC layer broadcast and is not acknowledged by other stations. Conventional stations are not aware of the notion of frame periods and RAMs and may not finish their packet transmission before the end of a frame period, thus overlapping the next frame period by a time of  $\Delta T$ . In this case, the Maestro station sends the RAM with a delay  $\Delta T$  and schedules the next RAM after duration  $T - \Delta T$  (shortening the BE traffic period) to compensate for this delay, keeping the average frame period duration  $T$  constant.

### 3.4. Sequence establishment and admission control

The Maestro station assigns itself a sequence number of one. A new real-time station which wants to join the existing E-MAC group assigns itself a sequence number  $i$  ( $i = 2, 3, \dots$ ), without involving the Maestro, by simply adding one to the total number of real-time stations  $n_r$  previously advertised by the Maestro in the RAM.

If a real-time station wants to join and is not the Maestro, it first has to check whether sufficient transmission time in  $T$  is available to accommodate its real-time packets. If  $t_r + \delta_{\text{guard}} + \delta_{\text{min}}^{\text{be}} + 2\text{SIFS} + t_{\text{slot}} + t_{\text{ack}} + t_{\text{data}} \leq T$ , the real-time station may join. Otherwise, it has to refrain from transmitting real-time packets, contending instead using BE priority. Here  $\delta_{\text{min}}^{\text{be}}$  is the minimum reserved contention period for BE traffic (of course the degraded RT traffic can also use this period) and  $\delta_{\text{guard}}$  is used to accommodate occasional retransmissions of real-time packets due to channel errors and interference. The  $\delta_{\text{min}}^{\text{be}}$  guarantees the opportunity for BE stations to transmit their fair share of packets. Note that  $\delta_{\text{min}}^{\text{be}}$  and  $\delta_{\text{guard}}$  are only for admission control purposes and do not actually represent reservation periods.

Once the real-time station is admitted to the E-MAC group, the self-selected sequence number  $i$  will be used to configure its *fixed* backoff time to  $t_{\text{back},i} = (i - 1)t_{\text{slot}}$ . Choosing the backoff this way results in a fixed transmission sequence, avoiding collisions among real-time stations. However, there is a small probability that two stations join at the same time and hence select the same backoff. This results in a collision (detected by absence of an ACK). To resolve this conflict, the two colliding stations wait for a duration  $rT$  ( $r$  being a random integer number, e.g., between

1 and 10) before trying to join again. The use of this  $rT$  is triggered only upon occasional collisions between two (or more) RT stations that happen to join at the same time. It should not be confused with the *fixed* backoff mentioned earlier, which is always used in order to define the RT station's order in the RT transmission sequence.

### 3.5. Slot reuse

Real-time stations begin their contention-free access upon the reception of a RAM. If a real-time packet is already waiting in the buffer upon receiving the RAM, the real-time station starts decrementing its backoff after the channel gets idle for AIFS ( $= \text{SIFS} + t_{\text{slot}}$ ) time. If another real-time station is transmitting, the backoff is frozen until the channel becomes idle again. Hence, if all real-time stations have a packet in their buffer upon receiving the RAM, any two consecutive real-time packets are being transmitted with an idle time of AIFS between them. This case is illustrated in Figure 3.

If a real-time station does not have a packet to send, it skips its turn. The subsequent station in the sequence will then transmit next with an idle time of AIFS  $+ t_{\text{slot}}$  after the previous transmission. Generally, if  $k$  consecutive stations refrain from transmitting a RT packet after the RAM, the idle time between two packets becomes AIFS  $+ kt_{\text{slot}}$ . This idle time might be longer than the DIFS of legacy 802.11 stations for large  $k$  (or  $n_r$ ) resulting in possible collisions between RT and BE stations. This event is prevented by the real-time stations by setting the duration field of the packet to  $2\text{SIFS} + t_{\text{ack}} + (n_r + 1)t_{\text{slot}}$ , which in turn sets the Network Allocation Vector, NAV, of legacy 802.11 stations. The last RT station in the sequence announces a duration (for the NAV of legacy 802.11 stations) of SIFS  $+ t_{\text{ack}}$ . The duration field is ignored by real-time stations. Hence, best-effort stations will refrain from transmitting until all active real-time stations have transmitted.

### 3.6. BE traffic preservation and RT fairness considerations

Stations may generate real-time packets at the application layer at any time during a frame period. The first packet of a given station is assigned high priority, as described before. However, if a second packet of the same station arrives during the same frame period, it is either queued

until the next frame period or contends as BE. We refer to such ‘degraded’ packets as BE-RT packets in the rest of the paper. BE-RT packets are treated the same as the BE packets in the contention-based period. However, upon hearing a new RAM, a station with a waiting BE-RT packet can ‘promote’ it back to high priority, thus being able to transmit it during the RT period. Promoting BE-RT packets to RT avoids that packets of the same flow get reordered at the MAC layer. Without promotion, the older BE-RT packet could be transmitted after the current RT packet.

There is no guarantee of minimum throughput for each BE station, but there is a minimum  $t_{be}$  for all contending stations (i.e., BE packets and RT-BE packets) in each frame period. For simplicity, we assume that RT stations are allowed to transmit only one high priority RT packet per frame period. The general case of different stations with different RT traffic requirements is discussed in Section 7. Stations can have (unrestricted) BE traffic sources in addition to RT traffic sources, in which case the two traffic types should be managed separately by two different queues.

### 3.7. Releasing reservations

If a real-time station has finished its session, the previously reserved resources must be released. If a station has not transmitted a real-time packet for a duration of  $lT$  ( $l$  being a predefined integer valid for all stations, e.g., 100), the Maestro supposes that it has left the real-time session. The Maestro then informs the other real-time stations about this fact in the next RAM together with the sequence number of the station that left. Then, all real-time stations with a higher sequence number can decrement their fixed backoff by one.

If a station has not transmitted a real-time packet for more than  $lT$  although it has not finished its real-time session yet, it has to re-join as if it were a new station, before transmitting the next real-time packet.

In case the Maestro finishes its real-time session, it adds the number of remaining RAMs it will still broadcast to the last  $j$  (e.g., 10) RAMs. Thus, other real-time stations know when they have to decrement their sequence number. Furthermore, the real-time station with sequence number 2 then knows when it has to take over the role of the Maestro. In the event of sudden Maestro disconnection (e.g., mobility), there is an additional timeout, after which the next station in the schedule takes over the role of the Maestro.

### 3.8. Difference to TDMA and 802.11e

In comparison with TDMA, our mechanism has the advantage of slot reuse, making it more efficient since no time slots are wasted, for example in case of silence suppression by the corresponding (e.g., voice) application codecs.

The differentiation based on traffic categories defined by IEEE 802.11e does not give any guarantees for real-time traffic since at high load there is a high number of collisions

even for real-time flows. Hence, under high load traffic the delay performance of IEEE 802.11e deteriorates. Moreover, previous work [1] showed that in heavily loaded networks, low-priority traffic has extremely low transmission probability when using EDCA, an effect called starvation of low-priority applications. Conversely, the proposed E-MAC guarantees a minimum data rate and a very low delay for all real-time stations almost irrespective of the network load while avoiding the starvation of best-effort stations. In summary, E-MAC has the following advantages compared to IEEE 802.11e:

- Almost no collisions for real-time stations during the contention-free period, due to the order imposed by the sequence of backoff values. In 802.11e, high-priority stations still suffer from increasing collisions when the number of real-time stations increases.
- Strict throughput and delay guarantees for admitted real-time traffic, due to the ‘reservation’ of periodic slots. In contrast, 802.11e offers only statistical guarantees.
- A very low guaranteed delay even under heavy-load traffic conditions. In IEEE 802.11e, the delay performance deteriorates as the number of high-priority stations increases.
- Better protection for best-effort traffic, due to the limitation on RT transmission and frame period, and the specification of minimum contention period  $\delta_{min}^{be}$ . In contrast, real-time stations in 802.11e can consume the whole channel capacity depending on the data rates at the sources.

## 4. MATHEMATICAL ANALYSIS

As the RT stations are also allowed to contend during the contention-based period, the analysis focuses on (a) the contention-free period where  $n_{rt}$  RT stations transmit in a TDMA-like way and (b) a contention-based period where all the  $n_{rt} + n_{be}$  stations contend for channel access (including  $n_{rt}$  BE-RT stations). The analysis focuses on the behavior of the network under saturated conditions, i.e., at any time instant both RT and BE stations have at least one packet in their transmission buffer.

### 4.1. Throughput analysis

During the contention-free period, all RT stations transmit their packets in a TDMA-like way. Under the assumption of saturated conditions, the RT stations always have a packet to transmit so they also participate in the contention period. Hence, all  $(n_{rt} + n_{be})$  stations participate in the contention during the contention period. After the end of the current time frame, any BE-RT packet which could not be sent during the contention period is promoted to RT priority again and transmitted during its corresponding time slot in the contention-free period.























