

On the Service Differentiation Capabilities of EY-NPMA and 802.11 DCF ^{*}

Orestis Tsigkas, Fotini-Niovi Pavlidou and Gerasimos Dimitriadis

Department of Electrical and Computer Engineering
Aristotle University of Thessaloniki
54124 Thessaloniki, Greece
{torestis,niovi,gedimitr}@auth.gr

Abstract. Wireless Local Area Networks have gained popularity at an unprecedented rate over the last few years. However, most existing access mechanisms cannot provide Quality-of-Service assurances. Even those that are QoS aware can only provide relative service differentiation. Based on EY-NPMA, the HIPERLAN Medium Access Control algorithm, we propose a dynamic priority medium access scheme to provide time-bounded services. By approximating an ideal Earliest Deadline First (EDF) scheduler, the proposed scheme can offer delay and delay jitter assurances while achieving high medium utilization. Furthermore, we compare our scheme with the IEEE 80.11 MAC protocol. Simulation studies document and confirm the positive characteristics of the proposed mechanism.

1 Introduction

Holding the promise of making ubiquitous mobile access to IP-based applications and services a reality, wireless networks have gained popularity at an unprecedented rate over the last few years. Concurrent with the expansion of wireless networks is a high demand for real-time applications with very stringent and diverse QoS requirements. Providing QoS requires the network to guarantee hard bounds on a set of measurable prespecified attributes, such as delay, bandwidth, probability of packet loss, and delay variance (jitter). However, the unstable nature of WLANS and their different characteristics compared to those of their wired counterparts, have a direct impact on their ability to guarantee bounds on these QoS metrics.

Soft guarantees can be provided instead, to increase the satisfaction level of users. Recent advances in encoding techniques allow real-time applications to adapt to network conditions and adjust their sending rate in the presence of time-varying channel capacities. Moreover, they allow a low ratio of packet loss without the Quality-of-Service perceived by the end user being affected.

^{*} This work was supported in part by the Hellenic State Scholarships Foundation (I.K.Y) and by the General Secretariat of Research and Technology in the framework of the G.S.R.T. project: Irakleitos.

However, when delay bounds are concerned, average performance may be proven insufficient. Packets whose delay exceeds a given threshold will be dropped, resulting in wasteful consumption of the scarce wireless resources. Providing delay guarantees becomes even more important if wireless access is considered to be just another hop in the communication path. End-to-end delay guarantees can be provided only if delays are bounded at each node along the path. If a node fails to meet local delay requirements, the end-to-end delay experienced by a packet may exceed its delay budget, resulting in a waste of the resources assigned to it at each node along the path. It would be therefore, desirable to design a medium access scheme that can meet the delay requirements of delay-sensitive traffic.

When delay-sensitive traffic is to be supported by the network, the optimal choice is to use the Earliest Deadline First (EDF) service discipline [1]. The EDF scheduler is a dynamic priority scheduler where the priorities for each packet are assigned as it arrives. The priority of each packet is given by its arrival time plus the delay budget associated with the flow that the packet belongs to. The scheduler selects the packet with the smallest deadline for transmission on the link. The priority of the packet increases with the amount of time it spends in the system. It has been proven that for any packet arrival process where a deadline can be associated with each packet, the EDF policy is optimal in terms of minimizing the maximum difference between the deadline of a packet and the time it is actually transmitted on the link [2].

Using EDF to control channel access in a distributed wireless environment, comprises a number of challenges, since there is not a management entity that can obtain all the information (e.g. number of active sessions, link states, statuses of the session queues) needed to make a scheduling decision. Therefore, a station with a lower priority packet may not defer access. Implementing a dynamic priority scheme with multiple priority levels can ensure that the station with the highest priority packet will gain access to the common medium. However, the probability of correct scheduling will be less than one, since two or more packets with different deadlines may belong to the same priority class. The problem then is to design a medium access scheme that can approximate EDF to the largest extent possible while achieving high levels of efficiency.

EY-NPMA [15], the HIPERLAN MAC protocol, is a dynamic priority scheme, which provides hierarchical independence of performance by means of channel access priority. However, its ability to track an ideal EDF scheduler, and thus provide service differentiation degrades severely as traffic load increases and the number of contending nodes grows. This is mainly due to the fact that EY-NPMA supports only 5 priority levels.

Based on EY-NPMA, we propose a dynamic priority Medium Access Control protocol to support time-bounded services in wireless networks. We modify the channel access scheme of EY-NPMA to support a high number of priority levels. Simulation results show that our scheme can approximate an ideal EDF scheduler to the largest possible extent while achieving high medium utilization. Applications are assumed to generate traffic with QoS requirements, which are

expressed in terms of a maximum allowable delay. Upon its arrival, each packet is assigned a lifetime which is equal to the delay budget associated with the flow that it belongs to. The residual lifetime of a packet is then used to compute its priority level. Packets that cannot be delivered within the allocated lifetime are discarded. Moreover, we assume that stations are able to hear each other (i.e. the network is fully connected). Furthermore, we compare our scheme with a mechanism that enhances the IEEE 80.11 MAC protocol with QoS support and assess the ability of each scheme to provide service differentiation while achieving high throughput.

The rest of the paper is structured as follows. In section 2 we provide an overview of distributed QoS capable medium access algorithms. Section 3 reviews the IEEE 802.11 MAC protocol and a proposed mechanism that enhances it with QoS support. Section 4 reviews EY-NPMA, the MAC protocol for HIPERLAN. Section 5 presents the objectives and the design of the proposed protocol. Section 6 deals with the simulation results of the proposed schemes, while section 7 concludes the paper.

2 Related Work

In this section we review some of the existing approaches to provide service differentiation at the distributed wireless MAC layer. The common feature of these distributed medium access algorithms is their attempt to provide QoS support by implementing a priority scheduler, thus allowing faster access to the channel to traffic classes with higher priority.

[6][11][10] propose modifications to the IEEE 802.11 Distributed Coordinated Function (DCF) [3] to incorporate differentiated service by supporting two or more priority levels. [6] proposes to modify the backoff scheme of DCF and set different values of CW_{min} and CW_{max} for different priority classes such that higher priority packets are more likely to be transmitted first. In [11] priorities are introduced by having each priority class use a different exponential increase factor during the backoff stage. Since this mechanism cannot ensure that contending nodes will defer to other contending nodes with higher priority packets, the authors propose a second mechanism that uses different DIFS values for different priority classes. The DIFS of a priority class j is defined as the sum of DIFS and maximum contention window of the higher priority class $j + 1$. As each station waits for the entire DIFS duration before it starts to count down the backoff interval, this scheme provides nodes with higher priority packets absolute priority to channel access. In [10] the backoff scheme of the IEEE 802.11 is modified and different interframe space (IFS) intervals are used, to support multiple priority levels. Higher priority stations use a shorter IFS and smaller "Contention Windows"; however, it is not assured that higher priority stations will have shorter backoff time than lower priority ones.

In [8] the authors propose a MAC protocol that provides multiple priority levels and adopts the black-burst mechanism [9] to guarantee that higher-priority packets will be always transmitted earlier than lower-priority packets. Packets

with the same priority are then transmitted in a round robin manner. The Busy Tone Priority Scheduling (BTPS) [7] protocol makes use of two busy tone signals to create two priority levels and adds an additional interframe space to the 802.11 DCF to ensure channel access of high priority packets

All of the above schemes attempt to provide distributed service differentiation by assigning traffic to fixed priority classes. However, even if hierarchical independence of performance is achieved, ensuring channel access to high priority traffic, performance guarantees cannot be provided on a per-flow basis. Rather, it can only be assured that a higher priority class of traffic will receive better service than a lower one. When delay sensitive traffic is concerned, delay bounds can be guaranteed only if traffic belonging to a class is bounded. Even then, the delay guarantee provided to a single flow cannot be better than the delay bound guaranteed to the aggregate of the traffic belonging to the class. In most cases, especially when the priority class becomes crowded, this delay bound is insufficient.

Better service differentiation can be provided if the priority level of packets contending for access to the wireless medium is updated in a dynamic manner. This allows packets with loose QoS requirements obtain better service than they would in a static priority scheduler without sacrificing the tight QoS guarantees that may be provided to other flows. In [14] a distributed dynamic priority scheme is proposed which is based on the 802.11 DCF. It is shown that this scheme can achieve a closer approximation to an ideal deadline based schedule than IEEE 802.11. EY-NPMA, the MAC protocol for the HIPERLAN, is another QoS aware medium access scheme that follows the dynamic priority approach for providing service differentiation. Both of these approaches, to implement a dynamic priority scheduler in the wireless environment, are reviewed in more detail in sections 3 and 4.

3 IEEE 802.11 DCF

In this section we review the IEEE 802.11 MAC scheme and a mechanism developed in [14] that allows IEEE 802.11 to provide QoS support.

3.1 DCF - Base Scheme

The fundamental access method of the IEEE 802.11 MAC protocol is a Distributed Coordination Function (DCF) known as carrier sense multiple access with collision avoidance (CSMA/CA). DCF allows for automatic medium sharing through the use of CSMA/CA and a random backoff time following a busy medium condition. Carrier sense is performed through both physical and virtual mechanisms.

A node which intends to transmit a packet invokes the carrier-sense mechanism to determine if the medium is idle for a minimum specified duration before attempting to transmit. If the medium is determined to be idle, the transmission

may proceed. Otherwise, the node defers access until the end of the current transmission. After deferral, or prior to attempting to transmit again immediately after a successful transmission, the node defers until the medium is determined to be idle without interruption for a period of time equal to Distributed Inter-Frame Spacing (DIFS). After this DIFS medium idle time, the node generates a random backoff period for an additional deferral time before transmitting, unless the backoff timer already contains a nonzero value, in which case the selection of a random number is not needed and not performed.

The backoff interval is uniformly distributed over the interval $[0, CW]$, where CW is an integer within the range of values CW_{min} and CW_{max} . The Contention Window (CW) parameter takes an initial value of CW_{min} , and is doubled at every unsuccessful attempt to transmit a packet, until it reaches the value of CW_{max} . Once it reaches CW_{max} , the CW remains at the value of CW_{max} until it is reset. This improves the stability of the access protocol under high-load conditions. The CW is reset to CW_{min} after every successful transmission.

The backoff interval counter is decremented while the medium is idle, and when it is zero the node is allowed to transmit. This process minimizes collisions during contention between multiple nodes that have been deferring to the same event.

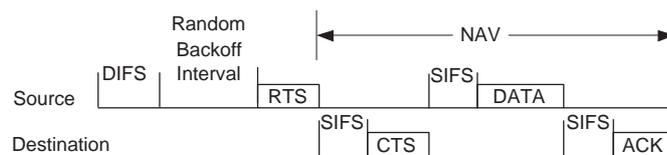


Fig. 1. IEEE 802.11 access method

A refinement of the method may be used under various circumstances to further minimize collisions, by providing a virtual carrier-sense mechanism. This mechanism is referred to as the network allocation vector (NAV). The NAV maintains a prediction of future traffic on the medium based on duration information that is announced in RTS/CTS frames prior to the actual exchange of data.

The RTS and CTS frames contain a Duration/ID field that defines the period of time that the medium is to be reserved to transmit the actual data frame and the returning ACK frame. All nodes within the reception range of either the originating node or the destination node learn of the medium reservation. The NAV may be thought of as a counter, which counts down to zero at a uniform rate. When the counter is zero, the virtual carrier-sense indication is that the medium is idle; when nonzero, the indication is busy. The medium shall be determined to be busy whenever the node is transmitting. In Fig. 1, an example of an IEEE 802.11 DCF access cycle is presented.

3.2 Priority Broadcast

In [14] a distributed priority scheme is proposed which takes advantage of the broadcast nature of the wireless medium to approximate an ideal deadline based schedule.

The authors propose to piggyback the priority index of a Head-of-Line packet onto existing handshake messages of the 802.11 DCF. If the RTS suffers no collisions, then all nodes in the broadcast region hear the RTS and add an entry in their local scheduling table. When the receiving node grants a CTS, it also appends the priority in the CTS frame. This allows the hidden nodes which are unable to hear the RTS, to add an entry in their scheduling tables upon hearing the CTS. Upon the successful completion of the packet transmission, each node removes the current packet from its scheduling table. Moreover, when transmitting a packet, each node also piggybacks its HOL packet priority. The priority is also copied in the ACK frame to allow hidden terminals to hear the HOL priority index. Neighbors monitor these transmissions and add another entry in their scheduling table.

Nodes keep a table of these times in order to assess the relative priority of their own Head-of-Line packet. Specifically, given a node's j local scheduling table \mathbf{S}_j and its rank r_j in its local scheduling table, the following equation is used in [14] to calculate the backoff interval,

$$f_l(\mathbf{S}_j) = \begin{cases} \text{Uniform}[0, CW_{min} - 1], & r_j = 1, \\ CW_{min} + \text{Uniform}[0, CW_{min} - 1], & r_j > 1. \end{cases} \quad (1)$$

The proposed backoff policy prevents nodes which are not ranked one in their scheduling table from contending in the first CW_{min} slots, thereby reducing contention for the top ranked nodes. The performance of the above policy improves when the scheduling table contains a higher fraction of the backlogged nodes' HOL indexes.

4 EY-NPMA

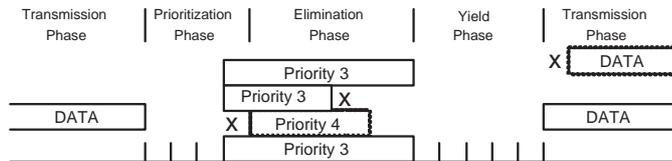


Fig. 2. EY-NPMA's synchronized channel access cycle

EY-NPMA, stands for Elimination-Yield Non-Pre-emptive Priority Multiple Access. Elimination-Yield describes the contention resolution scheme, while

NPMA refers to the principle of the HIPERLAN medium access mechanism that provides hierarchical independence of performance by means of channel access priority.

The channel access priority of a packet is defined by two parameters: the user priority (low and high) and the residual lifetime. When a new packet arrives, its lifetime is set to a value that cannot exceed 500 ms. At the beginning of the next channel access cycle the residual lifetime of the packet is updated. Depending on its residual lifetime, the packet is assigned one of the five priorities from 0 to 4, with 0 being the highest priority. Packets that cannot be delivered within the allocated lifetime are discarded. User priority determines whether the channel access priority of a packet can be upgraded to the highest priority level.

According to the HIPERLAN MAC protocol, every node that has data to transmit senses the channel for a period of 1700 bits. If no transmission takes place, the channel is considered free and the node starts transmitting immediately. Otherwise, the node synchronizes itself at the end of the current transmission interval and contends for the channel at the next channel access cycle according to the EY-NPMA scheme.

The NPMA channel access cycle is non-pre-emptive, so that only data transmission attempts ready at the start of a channel access cycle may contend for channel access in that channel access cycle. The synchronized channel access cycle comprises three phases: the prioritization, contention and transmission phase.

During the prioritization phase, priority resolution is performed, ensuring that only those data transmission attempts with the highest channel access priority will survive this phase. Each node having data to transmit senses the channel for as many slots as the priority of the packet in its buffer. If the channel is sensed idle for the whole interval, the contending node asserts the channel access priority by transmitting immediately a channel access burst. Otherwise, the node stops its data transmission attempt in the current channel access cycle.

The prioritization phase is immediately followed by the contention phase, during which only the data transmission attempts that have survived the prioritization phase contend for the right of transmission. The contention phase consists of two-subphases: elimination phase and yield phase. The objective of the elimination scheme is to eliminate as many as possible, but not all, contending nodes from competing for the right of transmission. A contending node transmits a channel access burst, whose length in slots is random between 1 and a predefined maximum, according to a truncated geometric distribution and then listens to the channel. If the channel is sensed as idle the node proceeds to the yield phase; otherwise, the node is eliminated and withdraws from the right of transmission in the current channel access cycle.

The yield phase complements the elimination phase by further resolving contention between the contending nodes that survived from the elimination scheme. During the yield phase, the contending nodes sense the channel for a random number of slots, and if the channel is sensed idle, they immediately enter the transmission phase by transmitting the packet stored in their buffer. All other stations sense the beginning of the transmission and refrain from transmitting.

Each phase reduces the number of contending nodes, so that there is a high probability that at the start of the transmission phase, the transmitting node will be unique. Moreover, the contention resolution scheme ensures that each contending node has a statistically equal chance to gain the right of transmission. In Fig. 2, an example of an EY-NPMA access cycle is presented. Solid line boxes represent actual transmissions, while dashed line boxes represent projected transmissions that did not take place because the station left the contention process. The X marks show when stations left the cycle.

The parameters in the HIPERLAN standard were chosen so as to achieve a collision rate of 3.5% that is independent of the number of simultaneous contending nodes, for a predefined maximum population of 256 simultaneous transmitting nodes. The maximum number of slots for which a station may burst during elimination (m_{es}) was set to 12, while the maximum number of slots that a station may backoff during the yield phase (m_y) was set to 9. Finally, the probability that defines the truncated geometric distribution used for deciding for how many slots a station should burst was set to 0.5. The working parameters of EY-NPMA may be expressed as a triplet of values $\{m_{es}, m_y, p_e\}$, which obviously the HIPERLAN standard defines as (12, 9, 0.5).

A performance study of EY-NPMA can be found in [12] and [13], where extended analytical and simulation results are presented. Further, it has been compared with DCF and EDCF in [4] and [5] respectively.

5 Proposed Protocol

5.1 Objectives

Based on EY-NPMA, we propose a dynamic priority Medium Access Control protocol, DP-TB, to support time-bounded services in wireless networks. The proposed medium access scheme provides support to traffic with delay requirements by approximating an ideal EDF schedule to the largest extent possible. To better achieve this, it has been designed with the following objectives in mind. Firstly, it should provide hierarchical independence of performance, causing contending nodes to defer access to other nodes with higher priority packets. Secondly, contention resolution of packets with the same priority should be achieved in a fair manner. Finally, any increase in the overhead of our scheme should not have an adverse impact on its throughput.

5.2 DP-TB

The proposed scheme preserves all of the three phases of the synchronized access cycle of the EY-NPMA scheme; yet, it features a different structure for the prioritization phase. Instead of a maximum of 5 prioritization slots, we propose a scheme that uses at most N slots for the prioritization phase. The prioritization phase, in the proposed DP-TB scheme, is further sub-divided in n sub-phases, where sub-phase j consists of at most p_j slots, such that $\sum_{i=1}^n p_i = N$. We do

not fix N and n to constant values, but rather let them be parameters of the system. Depending on the choice of N and n there is a trade-off between the extent that the ideal EDF scheduler can be approximated to and the throughput that can be achieved.

EY-NPMA uses N prioritization slots to support N priority levels. By subdividing the prioritization phase in n sub-phases we can provide a maximum of $P = \prod_{i=1}^n p_i$ priority levels, with 0 denoting the highest priority and $P - 1$ the lowest one.

As it has been mentioned the lifetime of a packet that has just arrived, is set to a value that cannot exceed 500 ms. We divide the interval of 500 ms into P time intervals, each of which has a duration of $t_p = 0.5/P$ sec. Then the priority index PI of a packet with residual lifetime RL can be computed as:

$$PI = \{k : k \cdot t_p \leq RL < (k + 1) \cdot t_p\} = \left\lfloor \frac{RL}{t_p} \right\rfloor \quad (2)$$

Given the priority index of a packet, the algorithm presented below can be used to determine for how many slots a node should sense the channel in each sub-phase in order to determine if it has the currently higher priority packet for transmission.

Algorithm

for($i = 1; i < n; i + +$)

$$\{ps_i = \left\lfloor \frac{PI}{\prod_{j=i+1}^n p_j} \right\rfloor$$

$$PI = PI - ps_i \cdot \prod_{j=i+1}^n p_j$$

}

$$ps_n = PI$$

As soon as the set of parameters $\{ps_1, \dots, ps_n\}$ has been computed a packet can contend for channel access in the prioritization phase. The prioritization phase of DP-TB works as follows. At the beginning of the first sub-phase, a station that has a packet ready for transmission senses the channel for as many as ps_1 slots. If the channel is idle for the whole sensing interval, the station transmits a burst of one slot and proceeds to the second sub-phase. Otherwise, the station exits contention and will have another chance for accessing the channel at the next cycle. In the same manner, during the second sub-phase the station senses the channel for ps_2 slots, and if the channel is sensed idle it transmits a burst slot. The procedure is repeated until the last sub-phase, where the node transmits a burst of random length, instead of just one burst slot. The length of this burst is between 1 and a predefined maximum number of slots.

The contention phase in DP-TB works as in EY-NPMA. However, during yield, a station, instead of randomly choosing an interval to backoff, will compute the duration of the backoff interval as:

$$Backoff_Interval = \left\lfloor \frac{RL - PI \cdot t_p}{t_p} \cdot (m_y + 1) \right\rfloor \quad (3)$$

where m_y is the maximum number of slots that a station may backoff during the yield phase. This ensures that if there is a successful transmission, the station that transmits is the one with the lowest residual lifetime among those who survived the elimination phase.

The medium access protocol proposed, allows us define a large number of priority levels by using a relatively small number of prioritization slots. When the length of the data payload is large, the added overhead of the prioritization phase in DP-TB can be alleviated by the lower collision rates. Indeed, as the number of the provided priority levels increases, the probability that two or more packets belong to the same priority class decreases. Consequently, not only a closer approximation to the ideal schedule can be achieved but, moreover, the collision rate is drastically reduced, as most of the time, very few (if more than one) packets will proceed to the elimination phase.

However, the above assumption does not hold true for short payloads. In [12] the authors examine the relationship between the payload size and efficiency of the EY-NPMA scheme and show that the increased overhead has an adverse impact on the throughput of the protocol for short payloads. In such scenarios, the time needed to transmit a control slot is a significant fraction of the time needed for the transmission of the actual payload. By optimally choosing the working parameters of EY-NPMA, the authors achieve a significant improvement in the utilization of the channel, despite an increase in the collision rate.

6 Simulation Experiments

6.1 Simulation Environment

The experiments conducted in this work aim at comparing the performance of the proposed medium access schemes and not the respective implementations as expressed in the standards. Towards this end, the capacity of the common medium was set to 23.5 Mbps and was considered to be ideal, that is the only reason behind erroneous reception was the simultaneous transmission of more than one stations (packet collision). Furthermore, all network stations were within one hop from each other, eliminating thus the appearance of hidden/exposed terminals. The working parameters of the proposed MAC schemes were set to the values defined in their standards. Nodes generated traffic according to a Poisson process with mean rate 100 kbps. The performance metrics of interest were average medium utilization and probability of correct scheduling. The ability of each proposed scheme to approximate an ideal EDF scheduler while preserving high medium utilization was evaluated. The tool that was used for these experiments was customly coded by the authors in C++.

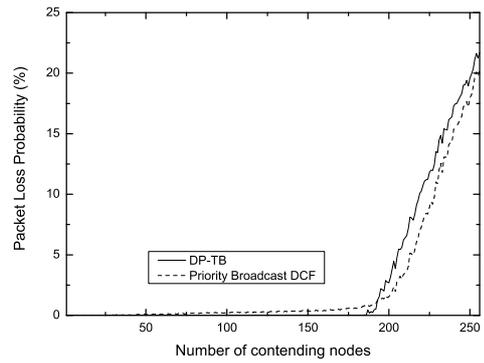
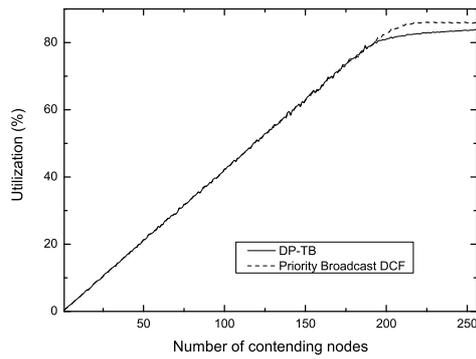
For DCF, a number of priority index assignment schemes is proposed in [14]. In our simulation experiments, we adopt the *Time To Live (TTL) allocation* scheme. In this scheme, a packet inserts its desired delay as its priority index. At the start of a contention cycle, each node updates the priority index of its HOL packet and the priority index of each entry in its local scheduling table. A packet is dropped when it experiences a queuing delay that is larger than its TTL index. This allocation scheme is the same as the one adopted in the HIPERLAN standard, where each generated packet is tagged with its residual lifetime. It should be noted that the authors in [14] propose the use of one byte to represent the priority tag. This would allow the definition of 256 priority levels; however, in their simulation experiments they make use of only 20 priority levels. We rather assume that each packet is tagged with its delay budget which can be represented with infinite precision. Although this assumption is not realistic, since the increased overhead would lead in underutilization of the available capacity, it let us define an infinite set of priority levels. Yet, it is shown that, even under this assumption, DP-TB which supports a finite set of priority levels can better approximate an ideal deadline based schedule.

DP-TB is designed to support 4840 priority levels. This is achieved by using at most 35 slots in the prioritization phase and subdividing it into 4 sub-phases, where the first two use at most 11 slots, the third 8 slots and the last one uses at most 5 slots. One can notice, that had we equally assigned the available prioritization slots in the 4 sub-phases, we could provide more priority levels without increasing the total number of prioritization slots. Indeed, this is the optimal way to assign the prioritization slots in order to support the maximum number of priority levels. However, simulation results indicated, that by using fewer slots for the last sub-phases, the average number of prioritization slots used decreases more rapidly as traffic load increases. This allows DP-TB to reduce the overhead introduced and achieve high levels of medium utilization. DP-TB uses at most 2 slots for the elimination and the yield phase.

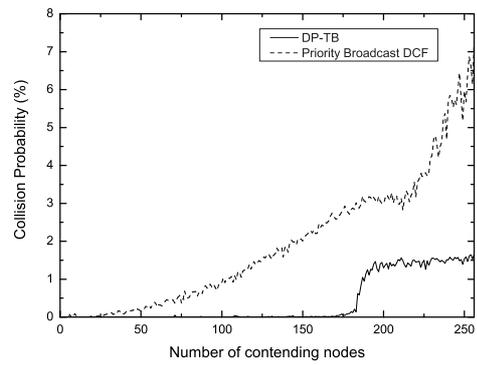
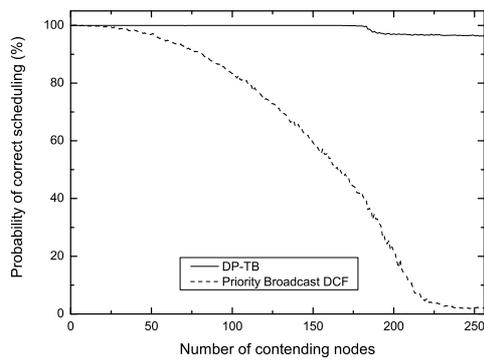
6.2 Simulation Results

Two sets of experiments were performed. Each node is assumed to generate one flow. The performance metrics were examined for different node populations (1-256 nodes) and packet sizes (2048, 1024, 512, 256, 128 bytes). To conserve space only the simulation results for 2048 and 128 bytes packet sizes are presented. As it has been already mentioned, these sizes correspond to two extreme cases. For 2048 bytes packet size the factor that mainly affects the throughput of the protocol is the collision rate, while for 128 bytes packet size the dominant factor becomes the protocol's overhead.

In the first set of simulations, the delay requirements of flows are looser as they are distributed over a wider range. Each newly generated flow has a delay budget, which is uniformly distributed in the interval [0, 500ms]. Upon an arrival of a new packet the residual lifetime assigned to it is equal to the delay requirement of the flow that it belongs to.



(a) Utilization vs. number of contending nodes (b) Packet loss probability vs. number of contending nodes



(c) Probability of correct scheduling vs. number of contending nodes (d) Collision probability vs. number of contending nodes

Fig. 3. Performance metrics for 2048 bytes packet size

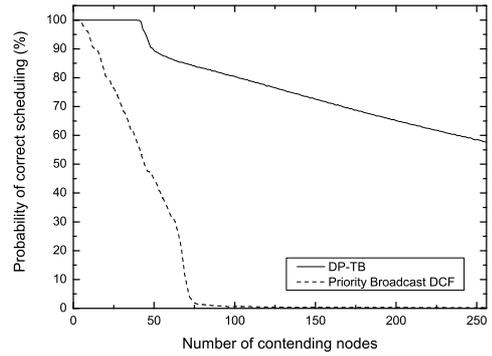
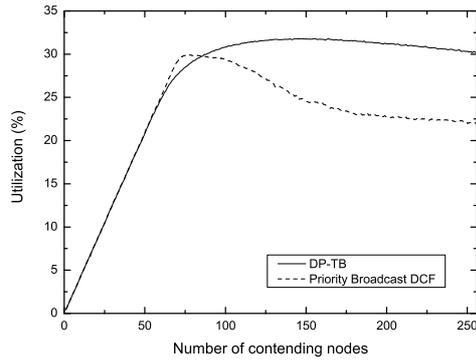
In Fig. 3(a) the mean medium utilization for 2048 bytes packet size is presented. For low and medium traffic load, DP-TB achieves slightly better medium utilization than Priority Broadcast DCF. Moreover, as the traffic load remains below 79%, DP-TB experiences no packet losses, while a small fraction of the contending packets are lost under Priority Broadcast DCF, as depicted in Fig. 3(b). Fig. 3(c) provides an explanation of the superiority of DP-TB under these traffic conditions. As it can be seen, DP-TB can approximate an ideal EDF scheduler to the largest extent possible. While traffic load is below 75%, DP-TB always makes the correct scheduling decision and schedules first the packets that their deadline is about to expire. The efficiency of DP-TB at approximating the ideal schedule is due to its ability to resolve the priorities of the contending packets. Fig. 3(d) shows that, under these traffic conditions, there are hardly any collisions, as the packet with the highest priority always captures the channel.

On the other hand, under Priority Broadcast DCF, the probability of correct scheduling decreases as the number of contending nodes increases. This results in a waste of resources (the delay budget assigned to each packet) as packets whose deadline is about to expire are pre-empted by packets that have enough delay budget to contend for channel access in the forthcoming access cycles. Moreover, even though Priority Broadcast DCF has lower overhead, it cannot achieve better medium utilization since it exhibits higher collision rate than DP-TB.

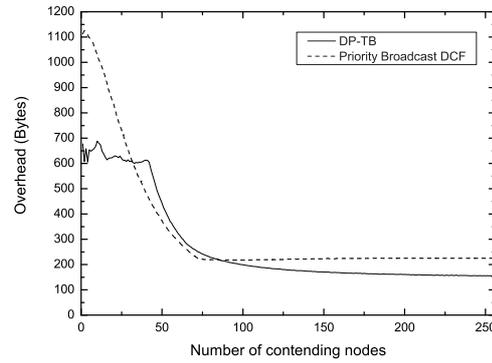
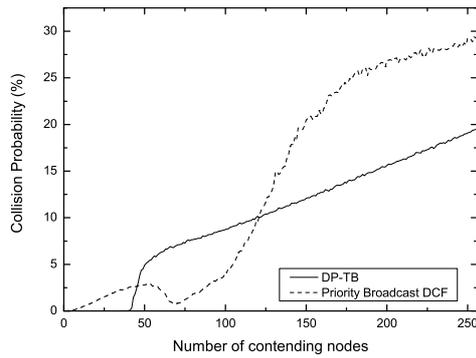
Under high traffic load, DP-TB makes the correct scheduling decision 96% of the time. However, the finite number of priority levels means that the probability of two packets having the same priority is nonzero, resulting in a slight increase in the collision rate. Priority Broadcast DCF suffers from a much higher collision rate; yet, the collisions are limited in the RTS/CTS exchange. Packet transmissions take place without collisions, allowing Priority Broadcast DCF to achieve slightly better medium utilization than DP-TB under high traffic load.

Fig. 4(a) shows the achieved medium utilization for 128 bytes packet length. DP-TB achieves higher throughput for a large number of contending nodes. However, it is outperformed by Priority Broadcast DCF when the number of contending nodes is in the range 60-85. In this range, Priority Broadcast DCF takes advantage of both its lower collision probability and its reduced overhead to achieve high medium utilization. The collision rate and the overhead of each protocol are depicted in Fig. 4(c) and Fig. 4(d) respectively.

The efficiency of each proposed scheme at approximating an ideal EDF scheduler is depicted in Fig. 4(b). DP-TB achieves a closer approximation to the ideal deadline based schedule than Priority Broadcast DCF does. It is this inefficiency of Priority Broadcast DCF at making the correct scheduling decision that has an adverse impact on its throughput for a large number of contending nodes. Moreover, DP-TB can be designed to achieve higher throughput under any traffic load in the expense of a reduction in the probability of correct scheduling. In our simulation experiments the parameters of DP-TB were chosen so as to achieve a good compromise between medium utilization and probability of correct scheduling for any number of contending users.

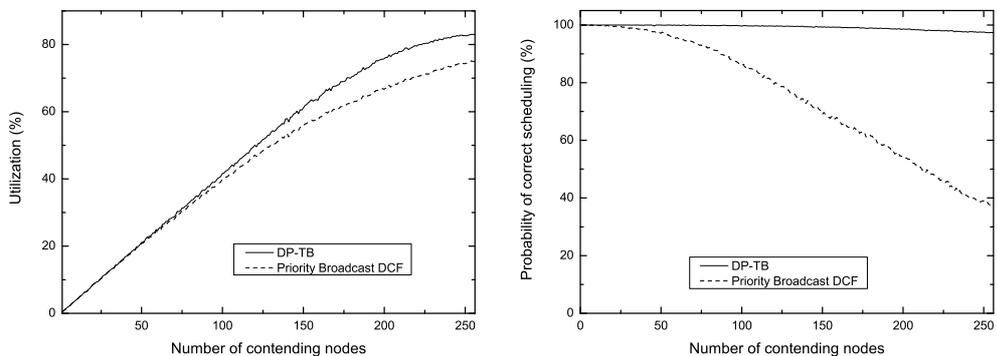


(a) Utilization vs. number of contending nodes (b) Probability of correct scheduling vs. number of contending nodes



(c) Collision probability vs. number of contending nodes (d) Overhead vs. number of contending nodes

Fig. 4. Performance metrics for 128 bytes packet size



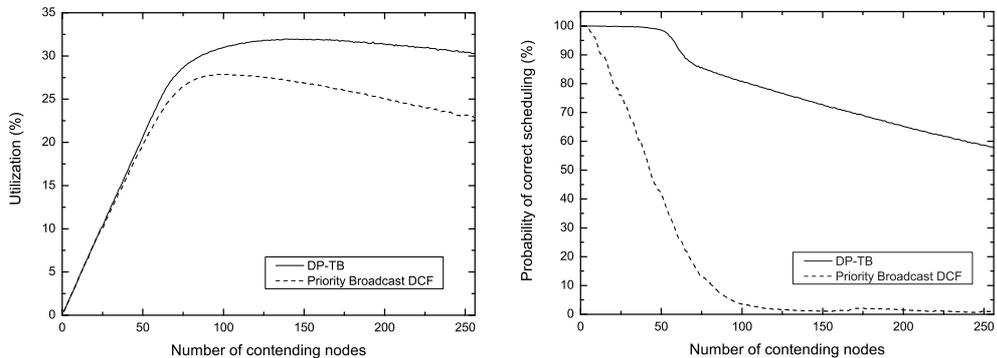
(a) Utilization vs. number of contending nodes (b) Probability of correct scheduling vs. number of contending nodes

Fig. 5. Performance metrics for 2048 bytes packet size

In the second set of experiments the QoS requirements are more stringent, since the delay budget of each contending flow is uniformly distributed in the interval $[0, 10\text{ms}]$. It should be noted that this is a more realistic scenario. Considering that wireless access is just another hop in a heterogeneous communication path that provides end-to-end delay guarantees, the delay budget of a flow at each node along the path will be small. The working parameters of each protocol were set to the values used in the first scenario.

Fig. 5-6 show that DP-TB outperforms Priority Broadcast DCF for any length of packet size. Not only does it achieve better medium utilization, but also it closely approximates the ideal deadline-based schedule. For 2048 bytes packet length, even when 256 nodes contend to gain channel access, the probability of correct scheduling for DP-TB is higher than 97%, while Priority Broadcast DCF starves to make the correct scheduling decision. The ability of DP-TB to track an ideal EDF scheduler, greatly depends on the number of the priority level it provides. The higher the number of the priority levels, the closer it can approximate the ideal schedule. By approximating an EDF scheduler to the largest possible extent, DP-TB minimizes the probability that a packet will be dropped due to lifetime expiration and, in this way the throughput of the protocol is increased. However, providing a very large number of priority levels is not beneficial. The efficiency of the protocol is adversely affected, since the introduced overhead increases as the number of priorities grows.

On the other hand, Priority Broadcast DCF attempts to approximate an ideal EDF scheduler by sharing information, which each node stores in its local scheduling table. It is evident that, the performance of the above policy improves when the scheduling table contains a higher fraction of the backlogged nodes' HOL indexes. However, it may be the case that when a node starts transmitting a packet, it does not have another packet stored in its buffer. This will be common for low rate applications that generate non-bursty traffic. Moreover, DCF allows



(a) Utilization vs. number of contending nodes (b) Probability of correct scheduling vs. number of contending nodes

Fig. 6. Performance metrics for 128 bytes packet size

a packet to gain channel access even when there is a packet with higher priority. In this scenario, where the QoS requirements are more stringent, the inefficiency of Priority Broadcast DCF at approximating an ideal EDF scheduler becomes evident.

7 Conclusions

In this work, we have proposed and evaluated the performance of a distributed dynamic priority medium access scheme to support time-bounded services in wireless networks. To better evaluate the performance of our scheme, we compared it with a mechanism proposed to enhance the service differentiation capabilities of IEEE 802.11 DCF. The mechanisms behind our proposed protocol that allow it to achieve better performance rely on the approximation of an ideal EDF scheduler. By closely approximating an ideal EDF scheduler we minimize the maximum difference between the deadline of a packet and the time it is actually transmitted on the wireless link. This allows us: 1) to save resources (i.e., some "delay budget" is not wasted) and in this way increase the medium utilization of the proposed scheme and 2) to minimize the delay variance and thus provide QoS in terms of both delay and delay jitter. The good characteristics of the proposed scheme were confirmed via simulations, where significant gains in performance were witnessed.

References

1. D. Ferrari and D. Verma, A scheme for real-time channel establishment in wide-area networks, *IEEE Journal on Selected Areas in Communications*, 8(3), pp. 368-379, Apr. 1990.

2. L. Georgiadis, R. Guerin, A. Parekh, Optimal Multiplexing on a single link: delay and buffer requirements, *IEEE Transactions on Information Theory*, 43(5), Sep. 1997
3. IEEE, 802.11: Wireless LAN medium access control (MAC) and physical layer specification, 1997.
4. G. Dimitriadis and F.-N. Pavlidou, Comparative performance evaluation of EDCA and EY-NPMA protocols, *IEEE Communications Letters*, vol. 8, no. 1, pp. 42-44, Jan. 2004.
5. J. Weinmiller et al., Performance study of access control in wireless LANs - IEEE 802.11 DCF and ETSI RES 10 HIPERLAN, *ACM/BALTZER Mobile Networks and Applications*, vol. 2, no. 1, pp. 55-67, 1997.
6. A. Veres, A. T. Campbell, M. Barry, and L.-H. Sun, Supporting service differentiation in wireless packet networks using distributed control, *IEEE Journal on Selected Areas in Communications*, 19(10), Oct. 2001.
7. X. Yang and N. H. Vaidya, Priority Scheduling in Wireless Ad Hoc Networks, *Proc. MOBIHOC*, pp. 71-79, June 2002.
8. J.-P. Sheu, C.-H. Liu, S.-L. Wu, and Y.-C. Tseng, A Priority MAC Protocol to Support Real-Time Traffic in AD Hoc Networks, *Kluwer Wireless Networks* 10, pp. 61-69, 2004.
9. J. L. Sobrinho, A. S. Krishnakumar, Quality-of-Service in Ad Hoc Carrier Sense Multiple Access Networks, *IEEE Journal on Selected Areas in Communications*, 17(8), Aug. 1999.
10. D. J. Dneg and R. S. Chang, A priority Scheme for IEEE 802.11 DCF access method, *IEICE Transactions on Communivation*, pp. 262-273, Jan. 1999.
11. I. Aad and C. Castellucia, Differentiation mechanisms for IEEE 802.11, *IEEE INFOCOM*, April 2001.
12. G. Dimitriadis and F.-N. Pavlidou, Two alternative schemes to EYNPMA for medium access in high bitrate wireless LANs, *International Journal of Wireless Personal Communications*, vol. 28, no. 2, pp. 121-142, Jan. 2004.
13. G. Anastasi et al., HIPERLAN/1 MAC protocol: Stability and performance analysis, *IEEE J. on Selected Areas in Communications*, vol. 18, no. 9, pp. 1787-1798, Sept. 2000.
14. V. Kanodia, C. LI, A. Sabharwal, B. Sadeghi, and E. Knightly, Distributed Priority Scheduling and Medium Access in AD Hoc Networks, *Kluwer Wireless Networks*, pp. 455-466, 2002.
15. ETSI, DTS BRAN-0020004-1 v0.m: Broadband radio access networks (BRAN): High PERFORMANCE Local Area Network (HIPERLAN) Type II: Data link control (DLC) layer: Part 1: Basic transport functions, 1999.