Temporal Fairness Provisioning in Multi-Rate Contention-Based 802.11e WLANs *

Ilenia Tinnirello*,+,1
*University of Palermo, ITALY
1ilenia@mwnl.snu.ac.kr

Sunghyun Choi+,2
+Seoul National University, KOREA
2schoi@snu.ac.kr

Abstract

The IEEE 802.11e extensions for QoS support in WLAN define the transmission opportunity (TXOP) concept, in order to limit the channel holding times of the contending stations in presence of delay-sensitive traffic. In this paper, we evaluate the use of TXOP for a different purpose: the "temporal fairness" provisioning among stations employing different data rates. We show that the equalization of the channel access times allows each station to obtain its throughput basically (1) proportional to its transmission rate, and (2) independent of the transmitted frame length. This also improves the aggregate throughput of the overall WLAN. For a given TXOP limit, i.e., a granted channel access time, a station is required to fragment its pending frame if the TXOP limit is too short, and is allowed to transmit multiple frames back-to-back in a burst if the TXOP limit is long enough, while different fragmentation and bursting rules are possible. Based on the analytical and simulation results, we demonstrate the advantages of the TXOP operations over the legacy 802.11 DCF, and compare different TXOP managing policies to find the optimal one.

1 Introduction

The distributed coordination function (DCF) of IEEE 802.11 wireless local area networks (WLANs) is based on carrier-sense multiple access with collision avoidance (CSMA/CA). In the long term, the DCF basically provides an equal opportunity to access the shared wireless channel to the contending stations. When contending stations transmit the frames of the same size, this equal channel access opportunity results in an equal share of the bandwidth or throughput. In this context, the 802.11 DCF is known to provide a "throughput-fair" channel access.

Today's 802.11 WLANs provide multiple data transmission rates by employing different sets of modulation and channel coding schemes. For example, the popular 802.11b physical (PHY) layer provides 1, 2, 5.5, and 11 Mbps [9]. These multiple transmission rates can be used for frame transmissions in an adaptive manner depending on the underlying channel conditions [11, 8, 5]. It should be noted that the above-mentioned throughput fairness is attained even if stations use different transmission rates, since a successful channel access results in a frame transmission irrespective of the frame transmission rate/duration. When different stations in a WLAN use different transmission rates, the performance of the stations using high rates is bounded by the performance of the stations using low rates. This phenomenon seems to be really weird and unfair in the viewpoint of the station employing high rates, and hence was referred to as a performance anomaly of the 802.11 [7]. However, as discussed above, this is nothing but an observation of the throughput fairness characteristic of the 802.11 DCF.

Unless we consider different quality-of-service (QoS) requirements, this throughput fairness characteristic does not seem desirable normally. It is more desirable to get the throughput proportional to the employed transmission rate. This can be basically achieved if "time fairness" can be provided. That is, stations should basically share the same amount of channel access time in average. For this purpose, a burst transmission mode can be introduced where a station is allowed to transmit multiple frames back-to-back in a burst provided that the entire frame exchange duration does not exceed a threshold value. This scheme is being standardized as part of the emerging IEEE 802.11e medium access control (MAC) [10, 3, 4, 6]. Interestingly, this idea is also introduced in [12] as part of the proposed Opportunistic Auto Rate (OAR) protocol independent from the 802.11e.

Under the "transmission opportunity" (TXOP) operation of the 802.11e contention-based channel access, a station is allowed to transmit multiple frames back-to-back with the total frame exchange duration bounded by a threshold, called *TXOP limit*. The TXOP operation

*The work reported in this paper was supported in part by the italian research program PRIN TWELVE and by the basic research program of Korea Science and Engineering Foundation (KOSEF) under Grant R08-2004-000-10384-0,

Proceedings of the Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM'05)

also requires that a frame should be fragmented into a set of smaller frames if the frame cannot be transmitted within the TXOP limit. To our best knowlege, analytical models for the 802.11e (e.g., [13, 14]) deal with the effects of the contention windows and inter frame spaces for service differentiation, but not the TXOP operation. The TXOP operation is briefly considered as part of the simulation-based evaluation of the 802.11e [4, 6]. In this paper, we analyze the performance of multi-rate contention-based 802.11e WLAN by considering different aspects of the TXOP operations. The contributions of the paper are mainly twofold. First, we provide a general analytical framework for the throughput and delay evaluation considering all the possible fragmentation and bursting rules of the TXOP operations. Second, we provide the guidelines for the optimal TXOP usage rules from both user and network perspectives via an extensive evaluation of different rules for many possible scenarios.

The rest of the paper is organized as follows: the TXOP operation and different managing rules are introduced in Section 2. Then, the throughput and delay performances of multi-rate 802.11e are derived in Sections 3 and 4. After presenting the performance evaluation based on both analysis and simulations in Section 5, we conclude the paper in Section 6 with the summary and future work.

2 TXOP Definition and Operations

2.1 Legacy DCF Operation

The IEEE 802.11 MAC receives a MAC Service Data Unit (MSDU) from the higher layer, i.e., 802.2 logical link control (LLC). Out of the received MSDU, one or more MAC Protocol Data Units (MPDUs) are generated and forwarded to the PHY for the transmission. When more than one MPDUs are generated out of a single MSDU, it is referred to as a *fragmentation*. The fragmentation is controlled via a parameter, called *Fragmentation Threshold*. When an MSDU is fragmented, all the fragmented MPDUs (or simply, fragments) except the very last one have the size equal to the threshold while the last fragment can be possibly smaller than the threshold.

In the legacy DCF, a station, which wins the contention, acquires the right to transmit an MSDU on the channel. When an MSDU is fragmented into multiple fragments, these fragments are transmitted back-to-back without a backoff procedure between two consecutive fragment transmissions. That is, after a successful fragment transmission, the subsequent fragment is transmitted after a Short Interframe Space (SIFS) interval from the end of the acknowledgment (ACK) corresponding to the previous fragment. This operation is referred to as the *fragment burst*. Note that the channel access unit is represented by a single MSDU transmission, irrespective of the time required to complete the transmission. The channel access is fair in terms of the access probability. This means that each station has the same probability to win the contention and, in a long term, obtains the same number of channel accesses (i.e., MSDU transmissions). This feature of the protocol is often referred as “throughput-fairness,” where the throughput here is defined in terms of MSDUs/sec. Whenever the MSDU size is fixed for all the stations, the throughput perceived by each station becomes basically constant in terms of bits/sec.

2.2 IEEE 802.11e EDCA and TXOP

The emerging 802.11e defines Enhanced Distributed Channel Access (EDCA) in order to provide differentiated services among contending stations [10]. Two main features of EDCA includes: (1) the probability to win the contention can be differentiated among stations (or traffic types, exactly speaking); and (2) the channel access unit can be defined based on the channel access time. Four traffic types based on the priority are defined and referred to as four Access Categories (ACs). The EDCA channel access rule is still based on CSMA/CA, but the channel access parameters are defined per AC. That is, Arbitration Interframe Space (AIFS) defined for each AC is used instead of Distributed Interframe Space (DIFS) of DCF, and minimum/maximum contention window sizes (i.e., CW_{min} and CW_{max}) are defined again per AC basis. We here skip all the details of EDCA (e.g., multiple channel access functions per station, etc.) since they are out of scope of this paper.

Once a station wins the contention and starts its transmission, EDCA specifies new channel utilization operations based on the transmission opportunity (TXOP), which represents a time interval in which the station is authorized to access the channel. If we assume that all the stations generate the same AC traffic, EDCA still provides a fair channel access. However, each access does not result in a single MSDU transmission. During a TXOP, one or more MPDUs can be transmitted in a burst, separated by SIFS. In this way, the entire burst appears to be a single instance of the wireless channel activity to other stations (see Fig. 1 (a)). TXOP is defined by a starting time and the maximum duration. The starting time is represented by the time when the medium is determined to be available under the EDCA access rule; the maximum duration is limited by a threshold, called *TXOP limit*, which is determined per AC basis and distributed by the Access Point (AP) in the beacon frames along with other access parameters

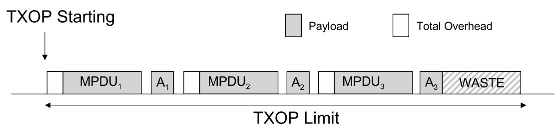

(a)

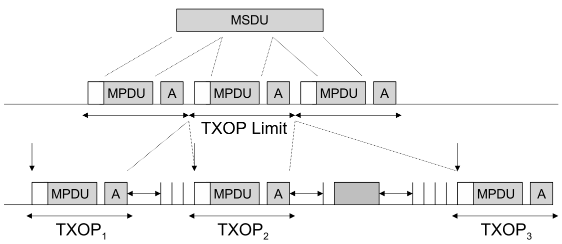

(b)

Figure 1. TXOP managing examples: (a) bursting where each MPDU in the burst is from an MSDU; and (b) fragmentation where an MSDU is fragmented to multiple MPDUs.

(i.e., AIFS, etc.).[1].

If the transmission duration of a pending MSDU exceeds the TXOP limit, it is mandatory to break the MSDU into smaller MPDUs. In fact, the TXOP limit cannot be violated in any case. For each fragment, the station has to contend for the channel access. Two or more TXOPs are required in order to complete the MSDU transmission (see Fig. 1 (b)). In order to use available TXOP duration limits as efficiently as possible, fragmentation can be also used after bursting, when the remaining TXOP duration is shorter than the time required to transmit the next whole MSDU. However, this type of fragmentation is optional.

2.3 TXOP Managing Policies

There can be various *TXOP managing policies* in terms of channel release and fragmentation rules. First, within a TXOP, channel release is arbitrary. That is, a station can utilize and stop transmission bursts whenever it wants within the TXOP limit. When fragmentation is used in order to fully exploit the TXOP limit (referred to as *full-time fragmentation* in this paper) or not to exceed the TXOP limit (referred to as *mandatory fragmentation*), stations can also choose the fragment sizes without any fixed threshold, different from

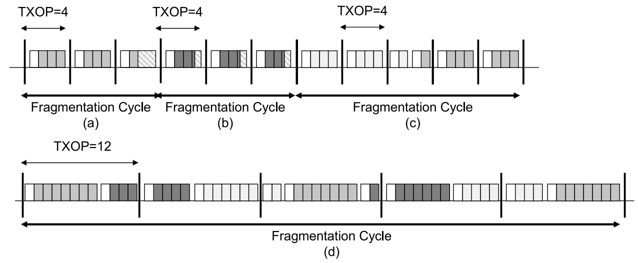

Figure 2. TXOP managing policy examples.

the legacy DCF fragmentation rule. This degree of freedom leads to different TXOP managing policies. Fig. 2 shows different channel release and fragmentation rules, where all of them are standard-compliant. Only consecutive TXOPs for a given station are plotted. That is, each vertical line in the figure actually corresponds to a time period, which potentially includes a number of backoff slots as well as frame transmissions by other stations. TXOP limit and MPDU sizes are expressed schematically in terms of transmission units, which correspond to the per-MPDU overhead time.[2] Successive MPDUs of the same color belongs to the same MSDU.

In Fig. 2 (a), we see that for an MSDU size of 7 transmission units and a TXOP time of 4 transmission units, we need at least three TXOPs to complete the transmission of the MSDU. In fact, due to the overhead, the time that can be used for the payload transmission is shorter than the whole TXOP limit (i.e., only 3 of the total 4 units). In the example, we observe that the first MSDU is broken into two MPDUs of the maximum size (i.e., 4 units) and a MPDU of a smaller size (i.e., 2 units). This rule is referred to as *maximum fragment rule*. Since the final MPDU transmission time is shorter than the TXOP limit, the residual TXOP time is wasted. However, the TXOP waste does not mean a channel idle period because after an AIFS interval from the end of the last MPDU transmission, other stations can access the channel. Indeed, it represents a waste only for the station which releases the channel since it does not fully exploit the granted channel access time.

Different fragmentation rules are possible. In Fig. 2 (b), the MSDU is broken into three equal-sized MPDUs. In this case, the TXOP waste is distributed among three TXOPs. This rule is referred to as *constant fragment rule*. In the last case in Fig. 2 (c), we observe that, in order to minimize the TXOP waste, the TXOP residual time is used to partially transmit the subsequent MSDU, i.e., the first fragment of the MSDU. In this case, the channel is released only after the TXOP time expiration, and fragment sizes are always the maximum

[1]TXOP limit is expressed in a 11-bit field in the unit of $32\mu s$, and the maximum TXOP limit is defined to be $65536\mu s$. However, in order to better understand the general trend, we often extend the performance evaluation considering up to $100ms$.

[2]The overhead is due to the MAC/PHY header/trailer and ACK transmission, and is fixed per MPDU transmitted at a given transmission rate. We use the terms "overhead" and "protocol overhead" interchangeably.

Proceedings of the Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM'05)
0-7695-2342-0/05 $20.00 © 2005 IEEE

possible. Since the overall protocol overheads in a given TXOP depend only on the number of conveyed MPDUs, the number of total payload bits per TXOP is not constant. In the case of TXOP limit exceeding an MSDU transmission time, burst transmissions can be employed. Stations can hold the channel for a given number of complete MSDU transmissions, and optionally transmit an extra MPDU to partially send the subsequent MSDU, in order to fully utilize the residual TXOP as shown in Fig. 2 (d). As a consequence of the fragmentation, the head of burst (HOB) MPDU in each TXOP is smaller than the original MSDU, and hence the collision times become shorter. However, the per-MSDU overhead increases. Note that, in general, TXOP wastes cannot be completely avoided, since residual times shorter than the per-MPDU overhead cannot be utilized at all. If the MSDUs have the same size, given a fragmentation rule, the MPDU sizes assume the same values periodically. We define the smallest number of TXOPs required to complete an integer number of MSDU transmissions as a *fragmentation cycle*, where it depends on the employed fragmentation rule. The fragmentation cycles are illustrated for all four cases in Fig. 2.

3 Analytical Framework

We consider an infrastructure Basic Service Set (BSS), which is composed of an AP and associated stations,[3] which employ different data rates. We assume that all the stations generate the traffic of the same AC (and hence, use the same access parameters). This also corresponds to adopting the same TXOP value by all the stations, set by the AP. Indeed, our model can easily take into account heterogeneous TXOP settings. We analyze the saturation performance, i.e., the performance perceived when the transmission queues of all the stations are never empty, without considering finite retry limits.

Following [1], we assume that channel accesses are slotted. In each slot, three different events can occur: (1) no transmission is originated, (2) a single station accesses the channel (i.e., successful transmission), or (3) two or more stations attempt to access (i.e., collision). Given the per-slot access probability τ and the collision probability p (which are fixed for all the stations), the average number of per-slot TXOPs received by each station is equal to the successful access probability $P_s = \tau(1-p)$. If $E[slot]$ is the average slot duration, $\tau(1-p)/E[slot]$ represents the average number of per-second TXOPs. Given the number N of total stations and the probability $P_{idle} = (1-\tau)^N$ that there is no transmission in the current slot, the average slot duration becomes:

$$E[slot] = P_{idle}\sigma + P_s E[T_s] + (1 - P_s - P_{idle})E[T_c]$$

where T_s is the time duration when the channel is busy due to a successful transmission, T_c is the average time duration when the channel is busy due to a collision, and σ is the duration of backoff slot-time, respectively. Then, for target station i, the throughput S_i and the average inter-service interval D_i between two successive TXOPs can be easily expressed as:

$$S_i = \frac{\tau(1-p)}{E[slot]} E[P_{T_i}] \qquad (1)$$

$$D_i = \frac{1}{S_i/E[P_{T_i}]} = \frac{E[slot]}{\tau(1-p)} \qquad (2)$$

where P_{T_i} is the total payload bits that are transmitted by the target station in a TXOP time. In the following, we detail the derivation of the average payload, transmission and collision times, while we compute τ and p by simply adopting the equations provided in [1].

Then, the key aspect of the TXOP operation modeling in a multi-rate WLAN is the computation of transmission and collision times. Note that when different data rates are employed in the network, frame length does not represent the actual channel occupation. Therefore, in the following, we express each channel activity due to a frame transmission or collision in terms of transmission time.

Let r_i be the data rate used by station i, r_i^* the corresponding ACK rate (i.e., the maximum BSS basic rate equal or lower to the data rate), and r_{PHY} the physical header rate, respectively. The time intervals $I_i^*(P)$ and $I_i(P)$ required to transmit a single MPDU with payload of P bits or to transmit and acknowledge the MPDU, at rate r_i, are respectively given by:

$$I_i^*(P) = \frac{P + H_{MAC}}{r_i} + \frac{H_{PHY}}{r_{PHY}}$$

$$I_i(P) = I_i^*(P) + \frac{A}{r_i^*} + \frac{H_{PHY}}{r_{PHY}} + SIFS$$

where H_{MAC} is the overhead length including the MAC header/trailer, H_{PHY} is the physical layer header (including preamble), and A is the ACK length. For a given transmission rate r_i, each frame transmission times can be divided into two different components: the payload transmission time P/r_i and the fixed per-MPDU overhead time $I_i(0)$.

3.1 Average Collision and Success Times

We group all the stations which perceive the same performance into *transmission classes*, where a transmission class is specified by the employed data rate and

[3] The reason why our work cannot be applied to the ad-hoc mode of 802.11e is that the TXOP limit value cannot be adjusted in the ad-hoc mode [10].

the MSDU size distribution. Let c be the total number of transmission classes, n_i the total number of stations belonging to transmission class i, and $N = n_1 + \cdots n_c$ the total number of contending stations in the network. Let $F_i(x)$ be the cumulative distribution of the time required to transmit the head of the burst (HOB) MPDU, at the beginning of a transmission attempt, for transmission class i. The collision duration is a random variable, which depends on the *longest duration frame*, requiring the longest transmission time among those involved in the collision, due to long payload size and/or low transmission rate. Let x_i be the number of class-i stations, which access the channel in the current slot, and $X = x_1 + \cdots + x_c$ the total number of stations accessing in the current slot. Whenever X is greater than 1, a collision occurs. Taking the conditional expectation on the number of stations belonging to each transmission class, we can express the average collision period $E[T_c]$ as follows:

$$E[T_c] = E[E[T_c|x_1, \cdots, x_c]|X > 1] =$$
$$= \sum_{\substack{x_1=0 \\ x_1+\cdots+x_c>1}}^{n_1} \cdots \sum_{x_c=0}^{n_c} Pr(x_1, \cdots, x_c) E[T_c|x_1, \cdots, x_c]$$

The probability to have $x_1, \cdots, x_c$ stations of classes $1, \cdots, c$, involved in the collision, respectively, is[4]:

$$Pr(x_1, \cdots, x_c) = \frac{\binom{n_1}{x_1} \cdots \binom{n_c}{x_c} \tau^X (1-\tau)^{N-X}}{1-(1-\tau)^N - N\tau(1-\tau)^{N-1}}$$

The average conditional value of the collision period is given by:

$$E[T_c|x_1, \cdots, x_c] = DIFS$$
$$+ \int_0^\infty [1 - F_1(x)^{x_1} \cdots F_c(x)^{x_c}]\, dx$$

Whenever a single station accesses the channel, a successful transmission is originated. The station can grab the channel until its TXOP expires. According to the employed TXOP managing rule and transmission rate, different amount of information bits can be transmitted within a TXOP. Let T_{s_i} be the successful transmission time for a class-i station (including protocol overhead times for all the MPDUs transmitted in the TXOP). In this case, we can easily compute the average successful time as:

$$E[T_s] = E[E[T_s|x_1, \cdots, x_c]|X = 1] = \sum_{i=0}^{c} \frac{n_i}{N} E[T_{s_i}]$$

[4]Note that, in the general case where different transmission classes belong to different ACs, the above formula can be easily adapted considering a different access probability τ_i for each AC. The computation of τ_i is out of the scope of the present paper and can rely on several results found in the literature.

In the next section, we derive $E[T_{s_i}]$, $F_i(x)$, and $E[P_{T_i}]$ for a class-i station since they are needed for the throughput and average inter-service interval calculations.

4 MPDU Size Distributions

The MPDU size distribution $F_i(x)$ is a function of the MSDU size distribution and employed TXOP managing policy. In the following, we assume that MSDU sizes are fixed to P for all the stations, and derive the MPDU size distributions (in terms of the number of payload bits per MPDU) for different TXOP managing policies. Since we derive the MPDU size distribution for a given transmission class i, for sake of presentation, we do not explicitly indicate the index, i.

4.1 Burst Transmissions without Fragmentation

When fragmentation is not used, the MPDU sizes are constant and equal to the MSDU size P. Let d be the number of MSDU transmissions that can be transmitted in a TXOP interval:

$$d = \left\lfloor \frac{TXOP + SIFS}{I(P) + SIFS} \right\rfloor$$

If $d = 0$, no transmission can be originated because the MPDU transmission time exceeds the TXOP limit. If $d > 0$, in each successful TXOP, both of the average aggregate payload size, P_T, and the average channel occupation time, T_s, are constant. Moreover, the HOB MPDU transmission times, needed to compute the average collision time, are fixed at $I^*(P)$. Then, the following equations can be easily derived:

$$\begin{aligned} E[P_T] &= d \cdot P \qquad (3) \\ E[T_s] &= (d-1)(I(P) + SIFS) + I(P) + DIFS \\ &= d \cdot (I(P) + SIFS) + 2\sigma \\ F(x) &= u(x - I^*(P)) \end{aligned}$$

where $u(x)$ is the unit step function.

4.2 Transmissions with Fragmentation

When we use fragmentation, different MPDU sizes can be generated in each TXOP. Given the fragmentation rule and the TXOP time, it is easy to convince that MPDU sizes assume a finite set of payload values in the range $[1, P]$ bits, and such values are repeated periodically (see Fig. 2 (d)). The cycle ends whenever the final MPDU transmitted within the TXOP time completes the pending MSDU transmission, and restarts in the next TXOP with the transmission of the first fragment of a new MSDU. Therefore, within a cycle, an integer number of MSDUs are transmitted. If L is the length of the cycle in terms of consecutive TXOPs, the total number of MSDUs transmitted within a cycle (or L TXOPs) can be expressed by $d \cdot L + l$, where $d \cdot L$ represents the number of MSDUs that can be fitted into the L TXOPs without fragmentation, and l represents the number of further MSDUs gained with the fragmentation, respectively. The, the following equations can be derived:

$$\begin{aligned} E[P_T] &= d \cdot P + \frac{l}{L} P \qquad (4) \\ E[T_s] &= d \cdot (I(P) + SIFS) + 2\sigma + \frac{lP}{Lr} + \\ & \quad (I(0) + SIFS)(1 + (l-1)/L) \\ F(x) &= \frac{1}{L} \sum_{i=1}^{L} u(x - I^*(P(i))) \end{aligned}$$

where $P(i)$ is the HOB MPDU size in the i-th TXOP within a cycle. Due to the cyclic nature of the fragmentation process, each HOB fragment of size $P(i)$ is transmitted with probability $1/L$. Note that $d+1+(l-1)/L$ represents the average number of MPDUs per TXOP. As shown in Fig. 2, such a number in general is not constant, and can be either $d+1$ or $d+2$ depending on the size of the HOB MPDU. In the following, we derive L, l, and $P(i)$ for different fragmentation rules.

4.2.1 Mandatory Fragmentation

Whenever $I(0) < TXOP < I(P)$, it is not possible to send a whole MSDU within the TXOP limit (i.e., $d = 0$), but the available time is enough to send a smaller payload. In this case, fragmentation is mandatory in order to access the channel. The maximum number of payload bits that can be transmitted in the TXOP is:

$$Q = (TXOP - I(0))r$$

Then, the minimum number of fragments required to complete the MSDU transmission is $\lceil P/Q \rceil$. In order to minimize the overall protocol overheads, we assume that each station uses the minimum number of fragments per MSDU. However, under this constraint, different fragment sizes can be resulted. A fragmentation cycle starts with the first fragment of the pending MSDU and ends with the last fragment which completes the MSDU transmission. See Fig. 2 (a) and (b) as examples. Therefore, the number l of MSDUs transmitted in a fragmentation cycle is just 1, and $L = \lceil P/Q \rceil$.

Regarding the generation of L fragments, we consider two different rules. Under the *maximum fragment rule*, the payload size of the first $L-1$ fragments is Q while the last fragment has the payload of size $P-(L-1)Q$ $(\le Q)$. See Fig. 2 (a). On the other hand, under the *constant fragment rule*, all the L fragments have basically the same payload size equal to $\lceil P/L \rceil$. See Fig. 2 (b). Note that, given the length L of the fragmentation cycle, $E[P_T]$ and $E[T_s]$ do not depend on the specific fragment sizes. The fragmentation rule affects only the collision time through different fragment sequences and $F(x)$ functions.

4.2.2 Full Time Fragmentation

Fragmentation can be also activated in order to fully utilize the TXOP limit whenever the residual TXOP, after d (≥ 0) MSDU transmissions, is greater than the per-MPDU overhead time. The number of available payload bits for the last fragment is:

$$Q = (TXOP - d(I(P) + SIFS) - I(0))r$$

Let's first consider the case of $d > 0$. In the subsequent TXOP, the HOB MPDU (or fragment) size is $P - Q$. Since the TXOP interval can accommodate d complete MSDU transmissions and Q residual bits in a burst of $d+1$ MPDUs, the last fragment of the burst is $2Q$ long. This corresponds to an HOB fragment of $P - 2Q$ size in the next TXOP and the last fragment of $3Q$ long, and so forth. If P is an integer multiple of Q, after P/Q TXOPs, the last fragment of the burst is equal to a whole MSDU, and hence a fragmentation cycle ends. Then, it is $L = P/Q$, $l = 1$ and $P(i) = P - i \cdot Q$ for $i = 0, \cdots, L-1$.

If P is not an integer multiple of Q, after $\lceil P/Q \rceil - 1$ TXOPs, the HOB MPDU size is $P\%Q$ $(< Q)$. In such a condition, the TXOP limit does not expire after d entire MSDU transmissions (since the overall TXOP accommodates the payload of size $dP + Q$). If the available time is not sufficient to accommodate a further per-MPDU overhead, the fragmentation cycle ends with a residual TXOP waste. Otherwise, another fragment can be transmitted within the TXOP. Note that in this case, the TXOP ends up including $(d + 2)$ MPDUs and the overall payload size becomes $dP + Q^*$, where $Q^* < Q$ because of the extra per-MPDU overhead, and is given by

$$Q^* = Q - (I(0) + SIFS)r$$

Therefore, the last fragment within the TXOP has the size of $Q^* - P\%Q$, and the following HOB MPDU size is $P-(Q^*-P\%Q)$. From this moment, in each subsequent TXOP, the HOB MPDU size is reduced by Q bits with respect to the preceding one, until a fragment of shorter than Q is produced again. In other words, we can consider the complete fragmentation cycle as the union of *sub-cycles*, which end with an HOB MPDU shorter than Q and start with a different first HOB MPDU size from other sub-cycles.

The fragmentation cycle does not depend on d, and even the case of $d = 0$ can be included, just considering that the fragment size is limited by the available TXOP

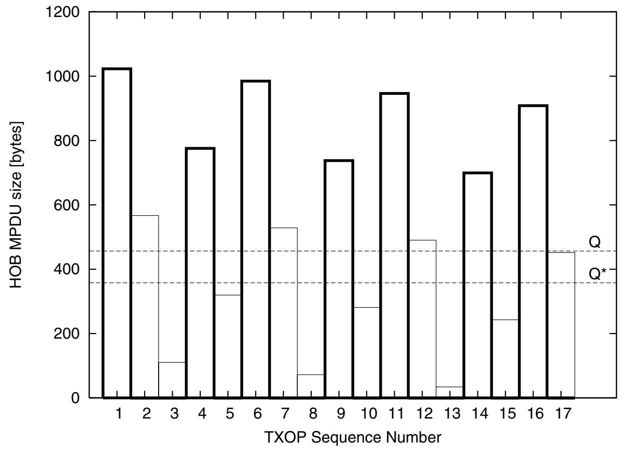

Figure 3. Example of HOB MPDU size evolution: TXOP=13.4ms, P=1024 bytes

or by the residual MSDU payload. Given $\Delta = Q^* - P\%Q$, the size of the HOB MPDU in the j-th TXOP of the i-th sub-cycle is determined by

$$\min\{P - (i\Delta)\%Q - jQ, dP + Q\}$$

$$i = 0, \cdots, l-1, \quad j = 0, \cdots, \left\lceil \frac{P - (i\Delta)\%Q}{Q} \right\rceil - 1$$

where the number l of MSDUs gained with the fragmentation in the cycle, is the lowest integer which solves the equation $(l\Delta)\%Q \geq Q^*$, and $(l-1)$ represents the number of $(d+2)$ MPDU bursts within the cycle. Accordingly, l also represents the number of sub-cycles within the fragmentation cycle. Even if the previous equation does not have a trivial solution (it requires some considerations about the number theory), the fragmentation cycle can be solved without knowing l, by stopping the cycle when an HOB MPDU size belongs to the range $[Q^*, Q]$. If this condition is met, the residual TXOP time after the bursting is shorter than the per-MPDU overhead time and no other fragment can be transmitted.

Fig. 3 shows an example of the HOB MPDU size evolution. The figure plots the HOB MPDU payload size vs. the TXOP sequence number within a fragmentation cycle, in the case of data rate equal to 1 Mbps, $P = 1024$ bytes, and $TXOP = 13.4ms$. Note that in this case, $d = 1$, i.e., a single complete MSDU can be transmitted during a TXOP. The Q and Q^* values are indicated as dashed lines. Each HOB MPDU can be numbered according to the TXOP in which it is generated. The first HOB MPDU corresponds to a whole MSDU. Because of the fragmentation, the following HOB MPDUs 2 and 3 are reduced by Q and $2Q$ bits, respectively. Since HOB MPDU 3 is shorter than Q (specifically, it is equal to $P\%Q$), in the third TXOP a sub-cycle is concluded, and the next HOB MPDU size is $P - \Delta$. Again, the following HOB MPDU is reduced by Q bits, and, since the resulting size is lower than Q, the sub-cycle ends. HOB MPDU 6 has the size of $P - (2\Delta)\%Q$, and so on. In the figure, the HOB MPDU corresponding to the starting of a sub-cycle are marked with bold lines. The fragmentation cycle ends when the HOB MPDU size results in the range $[Q^*, Q]$. In the example, this happens after 17 TXOPs. It results in that L is equal to 17 and l is equal to 7 (i.e., there are 7 sub-cycles).

5 Performance Evaluation

In this section, in order to validate our model, we compare our analytical results with the simulation results. We have used the ns-2 framework, and extended the 802.11 MAC module in order to include the TXOP operation and multiple data-rate capability. We assume the 802.11b as the underlying PHY while the BSS basic rate set is assumed to include 1 and 2 Mbps. Unless specified otherwise, we use a constant MSDU size of 1024 bytes. We plot analytical results with lines and simulation results with points. Each simulation run lasts for 100 s. For the sake of presentation, simulation results are plotted only in some figures and for some parameter settings. However, analytical and simulation results extremely well match in all the considered simulation scenarios. It should be noted that in many cases, the x-axis is drawn in a log scale.

5.1 Bursting and Mandatory Fragmentation

In this subsection, we first evaluate the effects of fragmentation and burst transmissions in the case of uniform transmission rates among the stations. The fragmentation mechanism is inherently not efficient since it introduces an extra overhead for each fragment. We here consider only the mandatory fragmentation case, i.e., an MSDU is fragmented only when its transmission time exceeds the TXOP limit. On the other hand, bursting is utilized whenever possible.

Fig. 4 shows the throughput performance as the TXOP increases for the cases of 5 and 50 stations with data rate equal to 1 Mbps. In the right part of the figure (corresponding to the case of $d > 0$[5]), fragmentation is not utilized. The throughput performance stays constant for the TXOP range corresponding to the same number of MSDU transmissions. Whenever the TXOP value becomes large enough for an extra MSDU transmission, a step improvement is observed. The step becomes more significant as the collision probability becomes high (in the figure, for the 50-station scenario,

[5] The figure is vertically divided into two cases; the left part corresponds to $d = 0$, and the right part corresponds to $d > 0$. A vertical line is drawn to show the boundary between two cases.

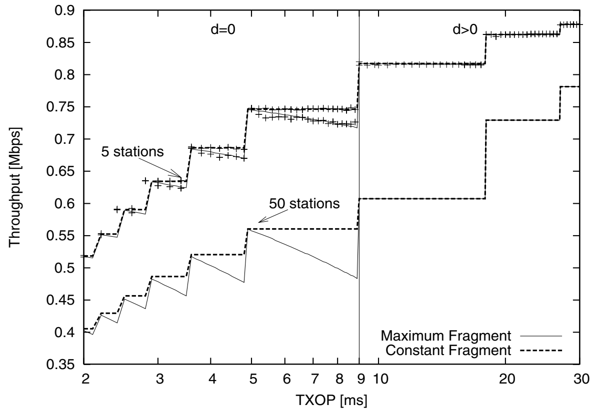

Figure 4. Impact of TXOP and different fragmentation rules

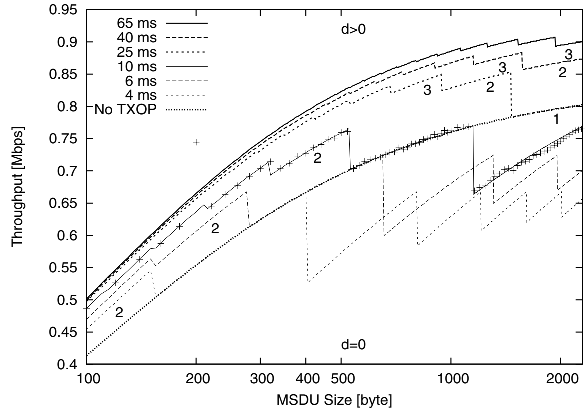

Figure 5. Impact of different MSDU sizes with constant fragment rule

the improvement is more evident than for the 5-station scenario), and less evident as the number of per-TXOP MPDU increases (i.e., as the TXOP limit increases).

In the left part of the figure (i.e., the case of $d = 0$), fragmentation is enforced, and two different fragmentation rules, i.e., the maximum and constant fragment rules per the definition in Section 2.3, are explored. In both cases, we observe that, as the TXOP limit increases, whenever the fragment number is decremented, there is an evident step improvement in the overall performance, due to the reduction of the per-MSDU overhead. For a TXOP limit value, which exactly corresponds to the fragment number decrement, we observe that the two fragmentation rules are practically equivalent. In fact, in such conditions, the MSDU size is an integer multiple of the maximum fragment payload size, $(TXOP - I(0))r$, and hence the constant fragment rule is equivalent with the maximum fragment one. However, in general, the constant fragment rule results in a better throughput performance than the maximum fragment one. This is because, for a given number of fragments, the constant fragment rule minimizes the average collision time, $E[T_c]$. Note that a collision time is determined by the longest frame involved with the collision. An interesting consequence of the independence of $E[P]$ and $E[T_s]$ on the fragmentation rules is that stations employing different fragmentation strategies perceive the same performance.

Dually, for a given TXOP value, the throughput performance strongly depend on the MSDU size. It is well known that in the absence of fragmentation, the system efficiency increases as the payload size increases. In the case of TXOP bursting/fragmentation operations, this trend is not true in general. That is, the throughput performance does not grow monotonically as the MSDU size increases. Fig. 5 presents the throughput performance when there are 10 stations transmitting at 1 Mbps for various TXOP limit values. The case of a single MSDU transmission without the TXOP operation is plotted as a reference. The curves above the reference without TXOP represents the cases of bursting transmissions, while the curves below the reference represents the cases of fragmentation. The number of per-TXOP MPDUs is indicated in some cases in the figure. Every time that such a number is reduced, the throughput degrades abruptly. Given the number of per-TXOP MPDUs, the throughput improves as the MSDU size increases. In the case of fragmentation, since the maximum possible MPDU size is fixed, and this maximum is obtained every time that the MSDU size fits exactly into an integer number of TXOPs, the throughput curve fluctuates as the MSDU size increases with the period corresponding to the maximum MPDU size determined by the TXOP limit.

5.2 Full-Time vs. Mandatory Fragmentation

In this subsection, we analyze the full-time fragmentation defined in Section 2.3, i.e., each TXOP is fully utilized by employing fragmentation if needed. On one side, this operation allows to maximize the payload bits per TXOP, and, in general, reduce the collision times. However, on the other side, the channel time waste increases due to the increased per-MSDU overhead. Figs. 6 and 7 compare the full-time fragmentation with the previously-considered constant fragment strategy, where fragmentation is employed only when an MSDU cannot be fitted into a given TXOP limit, for the cases of 5 and 50 stations, data rate equal to 1 Mbps, and MSDU size equal to 1024 bytes.

From Fig. 6, we observe that there is not a general result about the best TXOP usage in terms of throughput maximization. The optimum choice depends on the ra-

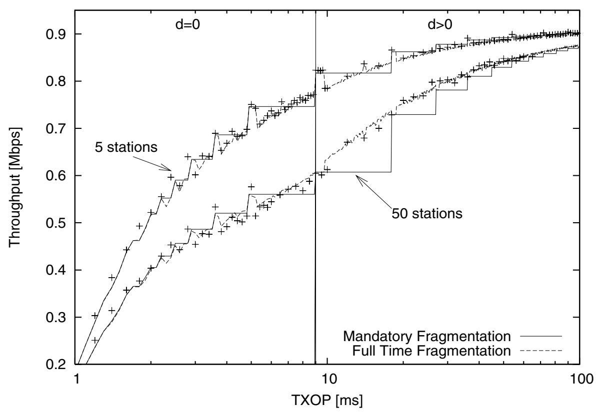

Figure 6. Comparison between different TXOP managing policies

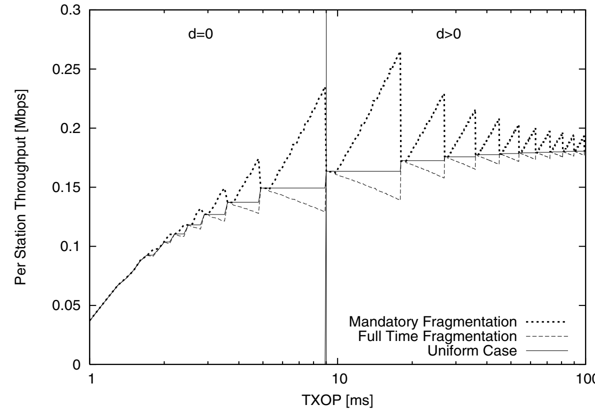

Figure 8. Resource repartition in the case of different TXOP usages.

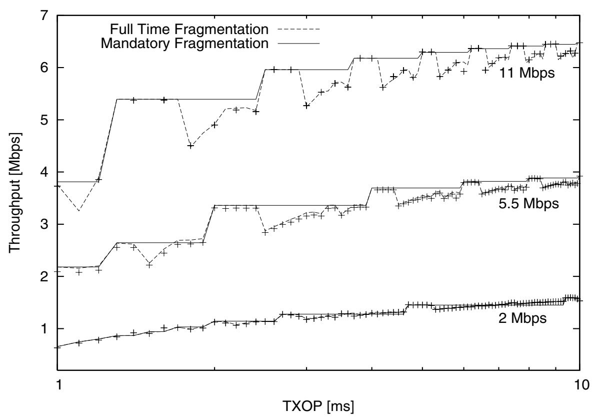

Figure 7. Impact of different data rates on TXOP utilization

tio between MPDU size and TXOP limit and on the collision probability experienced in the network. As the collision probability is high (e.g., 50-station scenario), the full-time fragmentation is superior more often. However, as the number of MSDU transmissions in each TXOP interval increases (i.e., large TXOP limit) or as the number of per-MSDU fragments increases (i.e., small TXOP limit), the full-time fragmentation becomes not useful. Moreover, Fig. 7 shows the throughput perceived by 10 stations employing 2, 5.5, and 11 Mbps. We observe that for higher data rates, full-time fragmentation is not effective, due to the relatively-longer per-MPDU overhead time compared to the payload transmission time. Although the system efficiency considerably improves as the TXOP limit increases, we have also to consider that the service latency degrades consequently.

Finally, note that if the stations in the network use different channel release rules, the bandwidth repartition is, in general, not fair. Fig. 8 shows the per-station performance in the case of 4 stations employing the mandatory fragmentation and 1 station employing the full-time fragmentation. We observe that the performance differences can be very remarkable. For example, for a TXOP value of 18 ms, the station which employs full-time fragmentation obtains almost two times the bandwidth obtained by the other stations. In fact, it is able to transmit an amount of payload bits larger than other stations, despite of the overall efficiency reduction due to the increment of the average slot time duration. The gap becomes smaller and smaller as the number of per-MSDU fragments (i.e., short TXOP) or per-TXOP MSDUs (i.e., long TXOP) increases.

5.3 Channel Holding Time Equalization

TXOP operation can equalize the channel holding time among the stations, thus providing the "temporal fairness." In this subsection, we investigate the effects of uniform TXOP settings on the resource repartition among stations with heterogeneous traffic sources and data rates.

5.3.1 Heterogeneous MSDU sizes

Fig. 9 plots the ratio of aggregated throughputs between two groups of stations: (1) 10 stations, which transmit fixed-size MSDUs of 1024 bytes, and (2) 10 stations, which transmit MSDUs of varying sizes. The x-axis represents the varying MSDU size of the second group. The performance is plotted for the cases without the TXOP operation (i.e., legacy DCF), mandatory fragmentation, and full-time fragmentation. When the legacy DCF access is utilized (labelled by "No TXOP"), the throughput perceived by each group is fair in terms of MPDUs/sec. Therefore, stations transmitting longer

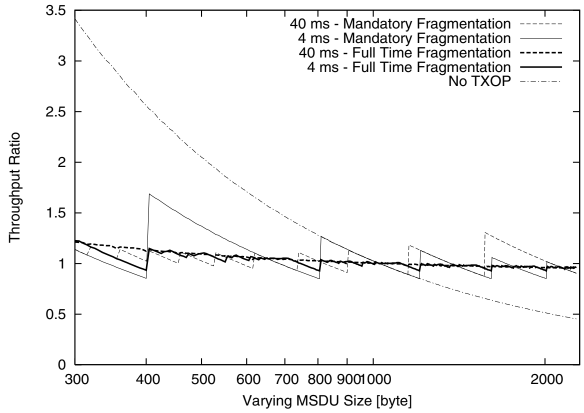

Figure 9. Bandwidth repartition among stations with different traffic sources

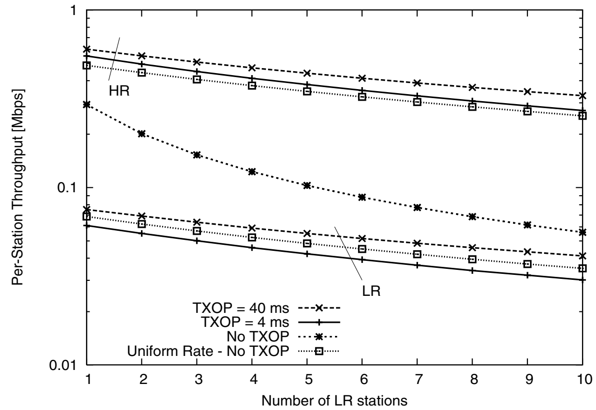

Figure 10. Bandwidth repartition among stations with different rates

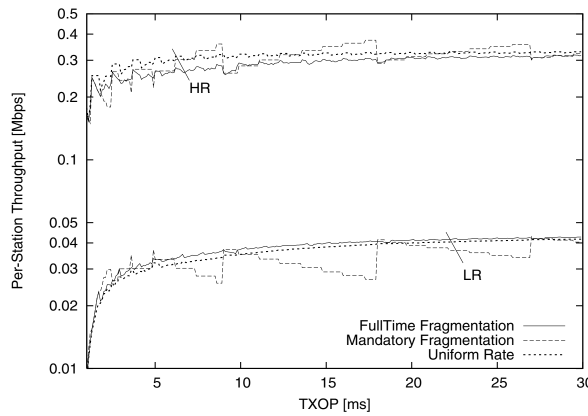

Figure 11. Performance of HR stations for varying TXOP

MPDUs achieve more throughput. That is, since the number of stations for each group is constant, the aggregated throughput ratio between two groups is equal to the payload ratio (i.e., for each x value, the throughput ratio is $1024/x$).

The introduction of the TXOP operation completely changes the bandwidth repartition criteria. When a given TXOP interval conveys an integer number of MSDUs for both groups, the aggregate throughput ratio between two groups of stations is given by the ratio between the channel utilization efficiency. That is, if O is the per-MPDU overhead, it is given by $1024/(1024+O)\cdot(x+O)/x$, where x is the given MSDU size. Note that this value is almost 1. Generally, the throughput ratio fluctuates around 1. Full-time fragmentation achieves the performance closer to the ideal temporal fairness. Moreover, this becomes more true as the TXOP limit value increases. Surprisingly, it is not true that stations with longer MSDUs always receive larger amount of bandwidth. For example, for a TXOP value of 4 ms, stations employing an MSDU value of 400 bytes receive more bandwidth than stations employing 1024 bytes. This is due to the different MPDU sizes and TXOP wastes generated by the fragmentation.

5.3.2 Heterogeneous data rates

We consider the case of two data rate classes, which employ fixed size MSDU (1024 bytes) and different data rates. In particular, we focus on the most critical case when 11 Mbps stations (referred to as High Rate or HR stations) share the channel with 1 Mbps stations (referred to as Low Rate or LR stations). Fig. 10 plots the average per-station throughput perceived by each station belonging to the same class. We run simulations for 10 HR stations and a varying number of LR stations. This number is represented in the x-axis. We compare the throughput with the one obtained when all the stations of the network employ the same data rate. We call this reference scenario as the uniform rate case. For each x value, the uniform rate curves refer to the per-station throughput obtained by the considered class, in the case of $10+x$ stations, all employing the same rate. We plot results in the case of the legacy DCF and for two different values of TXOP, i.e., 4 ms (which requires MSDUs fragmentation for LR stations) and 40 ms (in which multiple transmissions occurs for both the transmission classes). With the legacy DCF, the average throughput perceived by each station of the network is constant. In fact, since the throughput is constant in terms of MSDUs/s, and since the MSDU size is fixed for both groups, the bandwidth repartition does not depend on the data rate. Therefore, HR and LR curves overlap, i.e., all the stations attain the same throughput independently from the employed data rate. This phenomenon is often referred in the literature as performance anomaly. The introduction of TXOP is very effective to cope with this anomaly, by making the performance dependent only on the number of contending stations. As we can see from the figure, the performances for both classes are quite close to the one obtained when a single data transmission rate is employed. The difference is due to the resulting efficiency in the MPDU transmissions, due to the fragmentation or multiple frame bursting.

Since this efficiency depends on the TXOP value, Fig. 11 plots the per-station throughput as a function of the TXOP limit value for the case of 10 LR and 10 HR stations. If the mandatory fragmentation is employed, the resource repartition presents some oscillations, due to the rough equalization of the channel holding times. Note how the transmission inefficiency of LR stations corresponds to a throughput gain for the HR stations. A very interesting result is represented by the curves obtained in the case of full-time fragmentation. Although in most cases this type of fragmentation is less efficient than simple multiple transmissions, it allows to better equalize the channel access times. In fact, the fluctuations are reduced, and the system performance is almost constant for all the TXOP values. In other words, this type of TXOP managing allows to make system performance less sensitive to the MPDU transmission time/TXOP ratio.

6 Concluding Remarks

In today's multi-rate supporting 802.11 WLANs, the "throughput fairness" characteristic of the standard DCF appears undesirable. We, in this paper, analyze and evaluate the "temporal fairness" provisioning capability of the emerging IEEE 802.11e. We develop an analytical framework to evaluate different TXOP managing policies in terms of channel release and fragmentation rules. Through the model, we prove that the system throughput is optimized if MPDUs exceeding the TXOP limit are divided into *equal-sized* fragments. Moreover, we show that, based on the network condition, it can be more efficient to release the channel before the TXOP expiration without activating fragmentation to fully exploit the TXOP. We also discuss the effects of temporal fairness in the case of heterogeneous traffic sources and transmission rates. Since the channel holding time equalization requires, in general, fragmentation, which in turns introduces a large amount of overhead, we also show that there exists some tradeoff between fairness and system efficiency.

References

[1] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE Journal of Selected Areas in Communications*, vol. 18, no. 3, March 2000.

[2] G. Bianchi and I. Tinnirello, "Remarks on IEEE 802.11 DCF Performance Analysis," to appear in *IEEE Communications Letters*, 2004.

[3] S. Choi *et al.*, "Multiple Frame Exchanges during EDCF TXOP," IEEE 802.11-01/566r3, Jan. 2002.

[4] S. Choi *et al.*, "IEEE 802.11e Contention-Based Channel Access (EDCF) Performance Evaluation," in *Proc. IEEE ICC'03*, May 2003.

[5] J. del Prado Pavon and S. Choi, "Link Adaptation Strategy for IEEE 802.11 WLAN via Received Signal Strength Measurement," in *Proc. IEEE ICC'2003*, vol. 2, May 2003.

[6] J. del Prado Pavon and S. Shankar N, "Impact of Frame Size, Number of Stations and Mobility on the Throughput Performance of IEEE 802.11e," in *Proc. IEEE WCNC'2004*, March 2004.

[7] M. Heusse, F. Rousseu, G. Berger-Sabbatel, and A. Duda, "Performance Anomaly of 802.11b," in *Proc. IEEE INFOCOM'2003*, Mar. 2003.

[8] G. Holland, N. Vaidya, and P. Bahl, "A Rate-Adaptive MAC Protocol for Multi-Hop Wireless Networks," in *Proc. ACM SIGMOBILE'2001*, July 2001, pp. 236–251.

[9] IEEE 802.11b-1999, Supplement to Part 11: Wireless LAN MAC and PHY specifications: Higher-speed Physical Layer Extension in the 2.4 GHz Band, 1999.

[10] IEEE 802.11e/D8.0, Draft Supplement to Part 11: Wireless MAC and PHY specifications: MAC Enhancements for Quality of Service (QoS), Feb. 2004.

[11] D. Qiao, S. Choi, and K. G. Shin, "Goodput Analysis and Link Adaptation for IEEE 802.11a Wireless LANs," *IEEE Trans. on Mobile Computing*, vol. 1, no. 4, 2002, pp. 278–292.

[12] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly, "Opportunistic Media Access for Multirate Ad Hoc Networks," in *Proc. ACM MOBICOM'2002*, Atlanta, GA, September 2002.

[13] Y. Xiao, "An Analysis for Differentiated Services in IEEE 802.11 and IEEE 802.11e Wireless LANs," in *Proc. IEEE ICDCS'2004*, March 2004.

[14] J. Zhao *et al.*, "Performance Study of MAC for Service Differentiation in IEEE 802.11," in *Proc. IEEE GLOBECOM'2002*, Taipei, Nov. 2002.

Proceedings of the Sixth IEEE International Symposium on a World of Wireless Mobile and Multimedia Networks (WoWMoM'05)
0-7695-2342-0/05 $20.00 © 2005 **IEEE**