Energy-Efficient Airtime Allocation in Multi-Rate Multi-Power-Level Wireless LANs *

Daji Qiao
Iowa State University
Ames, IA 50011
daji@iastate.edu

Kang G. Shin
The University of Michigan
Ann Arbor, MI 48109
kgshin@eecs.umich.edu

ABSTRACT

This paper considers the energy-conservation problem in multi-rate multi-power-level wireless local area networks (WLANs). This problem is addressed from a unique angle — the system-level fairness which is significantly different from most of current research that focuses on improving the performance of each individual wireless station. To emphasize fair energy-consumption among contending stations, we introduce a new fairness notion, called *energy-conservation fairness*, in contrast to the conventional throughput fairness and airtime fairness. Another contribution of the paper is an energy-efficient airtime allocation scheme, which allocates airtime shares to contending stations in such a way that the combined airtime and energy-conservation fairness is achieved. Our simulation results show that, when the energy-conservation fairness is considered, both aggregate system throughput and overall system energy-efficiency can be improved significantly with all contending stations consuming a similar amount of energy.

Categories and Subject Descriptors

C.2.1 [**Network Architecture and Design**]: Wireless Communication

General Terms

Algorithms, Design, Performance

Keywords

IEEE 802.11e, Energy-Conservation Fairness, Airtime Allocation

1. INTRODUCTION

Since most wireless stations are battery-powered and have a limited amount of energy, energy conservation has always been one of the most important issues in IEEE 802.11 WLANs (Wireless Local Area Networks) [1] and has continuously been drawing considerable attention. Despite the need for a holistic, system approach to

---

this issue, most of current research has been focusing on each individual station's energy consumption. To remedy this deficiency, we consider the energy-conservation problem from the perspective of *system-level* fairness in 802.11 WLANs and offer some interesting observations and insights.

The 802.11 PHYs (physical layers) provide multiple transmission rates by employing different modulation and channel coding schemes. Furthermore, an increasing number of commercial 802.11 products support multiple transmit-power levels. For example, the Cisco Aironet 350 Series Client Adaptor [27] is an 802.11b-based WLAN device and supports six transmit-power levels from 0 dBm to 20 dBm. In a typical 802.11 WLAN, different wireless stations may choose different transmission strategies — each of which consists of transmission rate, transmit power, and/or data payload length — to communicate with the AP. We call such a trend *transmission strategy diversity* in 802.11 WLANs.

Throughput fairness is one of the well-studied fairness notions in 802.11 WLANs. Its goal is to fairly allocate contending stations bandwidths in proportion to their associated weights. Recently, with throughput fairness as the design goal, a serious system performance degradation is found to be inevitable in the presence of the transmission-strategy diversity. In particular, a performance anomaly caused by the transmission-rate diversity (a special form of the transmission-strategy diversity) was first discovered experimentally in [9] and later studied in-depth via modeling and analysis in [30]. Since then, the concept of *airtime fairness* has been introduced and recognized as a more reasonable design goal than throughput fairness for multi-rate 802.11 WLANs [5]. Unfortunately, the transmit-power diversity — another form of the transmission-strategy diversity — was not considered in the airtime fairness notion. Therefore, even when the perfect airtime fairness is achieved, high-power stations consume much more energy than low-power stations.

In this paper, we first introduce a new fairness notion, called *energy-conservation fairness*, to emphasize fair energy-consumption by all contending stations in an 802.11 WLAN. Note that different fairness notions yield different airtime allocations. In practice, instead of requesting for airtime fairness or energy-conservation fairness alone, a typical fairness request from a wireless station would be in a hybrid form so as to achieve energy-conservation fairness subject to a minimum airtime share requirement. Determining the airtime allocation for such a hybrid scenario is not as straightforward as determing that for airtime or energy-conservation fairness alone. We then present a three-step scheduling algorithm to solve this constrained optimization problem. Note that it is difficult to control the airtime allocation in an 802.11 WLAN. Many researchers have proposed various ways to achieve this goal by adjusting the channel access parameters of contending stations, which

involves complicated computation and is difficult, if not impossible, to realize since the MAC layer is typically hard-coded in the device firmware. The 802.11e standard [10] is an extension to the current 802.11 MAC for QoS (Quality-of-Service) provisioning. One of the key improvements in 802.11e is to introduce a new concept called TXOP (Transmission Opportunity). A TXOP is a time interval during which a wireless station is allowed to transmit during each medium access. Finally, we propose to control the airtime allocation in an 802.11e EDCA (Enhanced Distributed Channel Access) [10] system by adjusting the TXOP limits for different contending stations.

How to determine an energy-efficient transmission strategy for individual wireless stations has been studied extensively [3,6,7,11, 12,20–22] and is *not* the focus of this paper. Numerous scheduling algorithms have been proposed to achieve the weighted throughput fairness [4,14–19,23,24,29] or airtime fairness [2,5,8,13,25,26,28, 31] among contending stations in a wireless network. By contrast, our airtime-allocation scheme is designed to deal with generic fairness requests (airtime fairness or energy-conservation fairness or a combination thereof), which has not been addressed elsewhere.

The rest of the paper is organized as follows. Section 2 presents three fairness notions for 802.11 WLANs. The proposed energy-efficient airtime-allocation scheme and related implementation issues are discussed in Section 3. Section 4 presents and evaluates the simulation results, and finally, the paper concludes with Section 5.

2. FAIRNESS NOTIONS FOR MULTI-RATE MULTI-POWER-LEVEL 802.11 WLANS

We first review two popular fairness notions in 802.11 WLANs, namely, *throughput fairness* and *airtime fairness*,[1] and then introduce a new fairness notion called *energy-conservation fairness*. Moreover, we will discuss how the transmission-strategy diversity may affect the system performance under each of the three fairness notions. Table 2 lists the notations to be used throughout this paper, where station i's transmission strategy consists of data payload length (L_i), transmission rate (R_i), and transmit power consumption (P_i). When station i is in the idle mode, its power consumption is denoted by O_i. The power factor (ω_i) will be discussed in the next section.

Table 1: List of Notations

Notations	Comments
i	station ID
ϕ_i	associated weight
ω_i	power factor
L_i (octets)	data payload length
R_i (Mbps)	data transmission rate
P_i (mW)	transmit power consumption
O_i (mW)	idle power consumption
K_i	average # of medium accesses per unit time
N_i	# of data frame transmissions per medium access

2.1 Throughput ($\mathcal{B}$) Fairness

Throughput fairness is a well-studied fairness notion in 802.11 WLANs. Its goal is to allocate contending stations fair bandwidths in proportion to their associated weights:

$$\forall\, i, \ \ \mathcal{B}_i = L_i N_i K_i \propto \phi_i. \qquad (1)$$

1 *Throughput fairness* and *airtime fairness* are also often referred to as *bandwidth fairness* and *time-based fairness*, respectively.

The 802.11 DCF [1] is designed to offer equal transmission opportunities (or long-term equal medium access probabilities) to all contending stations, and each station is only allowed to attempt a single frame transmission for each medium access. In other words, with the DCF, we have

$$\forall\, i, j, \ \ K_i = K_j, \ \ N_i = N_j = 1. \qquad (2)$$

Therefore, when the stations transmit data frames with the payload size in proportion to their associated weights, i.e., $\forall\, i,\ L_i \propto \phi_i$, the DCF yields the desired throughput fairness.

2.1.1 Impact of Data-Payload-Size Diversity

From the above analysis, one can see that throughput fairness under the DCF is achieved only when the data payload sizes of contending stations are proportional to their associated weights.

2.1.2 Impact of Transmission-Rate Diversity

From Eq. (1), we can see that transmission-rate diversity has no impact on throughput fairness. This is surprising at the first sight but rather logical because, as long as a wireless station wins the contention to access the wireless medium, it is allowed to transmit its data frame regardless of the selected transmission rate.

On the other hand, under certain circumstances, the transmission-rate diversity may degrade significantly the aggregate system throughput. As pointed out in [9], in multi-rate 802.11 WLANs where different stations may transmit data at different rates, the DCF results in the following performance anomaly: *the throughput of stations transmitting at higher rates may be affected greatly by stations transmitting at lower rates, and hence, those high-rate stations may suffer an unexpected throughput degradation under the notion of throughput fairness.* This is because, according to the definition of throughput fairness, we have

$$\mathcal{B}_{high\_rate} = \mathcal{B}_{low\_rate} \cdot \frac{\phi_{high\_rate}}{\phi_{low\_rate}} < R_{low\_rate} \cdot \frac{\phi_{high\_rate}}{\phi_{low\_rate}}, \ \ (3)$$

meaning that the throughput of a high-rate station is bounded below a certain level determined by other low-rate contending stations, regardless of its own transmission rate. For example, under the uniform throughput fairness ($\phi_{high\_rate} = \phi_{low\_rate}$), the throughput of any contending station in the network is always bounded below the lowest transmission rate and, consequently, so is the aggregate system throughput.

2.1.3 Impact of Transmit-Power Diversity

It is clear from Eq. (1) that transmit-power diversity has no impact on throughput fairness.

2.2 Airtime ($\mathcal{A}$) Fairness

Recently, airtime fairness has been introduced to deal with the above-described performance anomaly, and is considered a more reasonable and intuitive fairness notion than throughput fairness for multi-rate 802.11 WLANs. Its goal is to allocate contending stations fair amounts of airtime (rather than bandwidths) in proportion to their associated weights:

$$\left. \begin{array}{c} \forall\, i, \mathcal{A}_i = \frac{L_i N_i K_i}{R_i} = \frac{\mathcal{B}_i}{R_i} \propto \phi_i \\ \sum_{i=1}^n \mathcal{A}_i = 1 \end{array} \right\} \Rightarrow \forall\, i, \mathcal{A}_i = \frac{\phi_i}{\sum_{i=1}^n \phi_i}. \ \ (4)$$

Under the notion of airtime fairness, each contending station is guaranteed to receive a certain percentage of the total airtime. This way, the throughput performance of an individual station is isolated from, and unaffected by, those of other stations. The authors of [28] showed that airtime fairness can make significant improvements in the aggregate system throughput in multi-rate 802.11 WLANs.

2.2.1 Impacts of Data-Payload-Size and Transmission-Rate Diversities

Unlike throughput fairness, both data-payload-size and transmission-rate diversities affect airtime fairness. Airtime fairness may be achieved when certain requirements of data payload sizes and transmission rates of contending stations are satisfied. For example, combining Eq. (4) with Eq. (2), we can see that airtime fairness under the DCF is achieved only when the stations' frame transmission durations are proportional to their associated weights, i.e., $\forall i, \frac{L_i}{R_i} \propto \phi_i$.

2.2.2 Impact of Transmit-Power Diversity

It is clear from Eq. (4) that transmit-power diversity has no impact on airtime fairness.

2.3 Energy-Conservation (ε) Fairness

Transmit-power diversity is a unique phenomenon in the emerging multi-rate multi-power-level 802.11 WLANs, which is not considered in airtime fairness. To understand the impact of transmit-power diversity, let's consider the following scenario. In a typical 802.11 WLAN, some wireless stations may be far away from the AP, and hence, transmit at high power levels to overcome the low quality of their radio transmissions, while some stations may be near the AP and can communicate with AP at low transmit-power levels. In this scenario, even when the perfect airtime fairness is achieved and all the contending stations have an equal share of the total airtime usage, high-power stations consume much more energy than low-power stations. Based on the above observations, we introduce a new fairness notion, called *energy-conservation fairness*, to emphasize fair energy consumption by all contending stations in proportion to their associated weights:[2]

$$\forall i, \ \mathcal{E}_i = \frac{L_i N_i K_i (P_i - O_i)}{R_i} = \frac{\mathcal{B}_i (P_i - O_i)}{R_i}$$
$$= \mathcal{A}_i (P_i - O_i) \propto \phi_i \tag{5}$$
$$\implies \forall i, \ \mathcal{A}_i \propto \frac{\phi_i}{P_i - O_i}.$$

As a result, we have

$$\left. \begin{array}{c} \forall i, \ \mathcal{A}_i \propto \frac{\phi_i}{P_i - O_i} \\ \sum_{i=1}^n \mathcal{A}_i = 1 \end{array} \right\} \implies \forall i, \ \mathcal{A}_i = \frac{\frac{\phi_i}{P_i - O_i}}{\sum_{i=1}^n \frac{\phi_i}{P_i - O_i}}. \tag{6}$$

Thus, energy-conservation fairness (i) requires each contending station to receive an airtime share proportional to the ratio of its associated weight to the difference between its transmit power consumption and its idle power consumption, (ii) is equivalent to airtime fairness in single-power-level 802.11 WLANs, and (iii) is equivalent to both airtime and throughput fairness in single-rate single-power-level 802.11 WLANs.

2.3.1 Impacts of Data-Payload-Size, Transmission-Rate, and Transmit-Power Diversities

From Eq. (5), we can see that energy-conservation fairness is affected by all three types of diversity. It may be achieved only when certain requirements about data payload sizes, transmission rates, and transmit-power consumptions of contending stations are satisfied.

---

[2]In this paper, we assume no power-saving mode, meaning that, when an 802.11 device is not actively transmitting or receiving, it remains in the idle mode and keeps sensing the channel. Furthermore, we assume that, when an 802.11 device is in the idle mode, it consumes the same amount of power as when it is receiving data.

2.4 Fairness Indices

We use three fairness indices, $\mathcal{I}_\mathcal{B}$, $\mathcal{I}_\mathcal{A}$, and $\mathcal{I}_\mathcal{E}$, to quantitatively evaluate throughput, airtime, and energy-conservation fairness, respectively, of an airtime-allocation scheme. For example, the energy-conservation fairness index may be calculated as:

$$\mathcal{I}_\mathcal{E} = \frac{\left(\sum_{i=1}^n \frac{\mathcal{A}_i (P_i - O_i)}{\phi_i} \right)^2}{n \cdot \sum_{i=1}^n \left(\frac{\mathcal{A}_i (P_i - O_i)}{\phi_i} \right)^2}. \tag{7}$$

When all contending stations have the same energy-consumption performance, the perfect energy-conservation fairness is achieved and $\mathcal{I}_\mathcal{E}$ is 1. In general, $\mathcal{I}_\mathcal{E}$ is between 0 and 1, and the closer to 1 the $\mathcal{I}_\mathcal{E}$ value, the fairer an airtime-allocation scheme gets in terms of energy conservation. Similarly, we can calculate $\mathcal{I}_\mathcal{B}$ and $\mathcal{I}_\mathcal{A}$.

3. ENERGY-EFFICIENT AIRTIME ALLOCATION

Instead of requesting for airtime fairness or energy-conservation fairness alone, a typical fairness request from a wireless station may be in a hybrid form. Here we focus on the *minimum-airtime-share constrained energy-conservation fairness problem*, which combines airtime fairness with energy-conservation fairness. In order to characterize a wireless station's desired minimum airtime share, we introduce the following notation. As shown in Table 2, each station is associated with a *power factor*, denoted by ω, in addition to its associated weight ϕ. The minimum airtime share of station i in this hybrid scenario is then $\omega_i \mathcal{A}_i$ where $\mathcal{A}_i$ is calculated according to Eq. (4) — definition of airtime fairness.

Determining the airtime allocation for this hybrid problem is not as straightforward as that for airtime fairness or energy-conservation fairness alone. We now describe a three-step energy-efficient airtime-allocation scheme as a possible solution. Moreover, to help better understand how our scheme works, we will show how each step of our scheme is applied to the following simple example scenario:

> **EXAMPLE SCENARIO.** *Four stations are contending for the shared wireless medium, and their associated weight, power factor, and transmit power information are listed in Table 2. The objective is to find the airtime allocation that yields the best energy-conservation fairness while meeting the minimum-airtime-share constraint for each station.*

Table 2: Example Scenario: Four Contending Stations

i (station ID)	1	2	3	4
ϕ_i	1	1	1	1
ω_i	1	1	$\frac{1}{4}$	$\frac{1}{2}$
$(P_i - O_i)$	P_{min}	$3P_{min}$	$4P_{min}$	$4P_{min}$

3.1 Step I: Calculate Original Airtime Shares

The first step of our scheme is to find the original values ($\mathcal{A}^{or}$) of the airtime shares, which are obtained by considering airtime fairness only. Hence, according to Eq. (4), we have

$$\forall i, \ \mathcal{A}_i^{or} = \frac{\phi_i}{\sum_{i=1}^n \phi_i}. \tag{8}$$

Now, we apply *Step I* to our example scenario. Consider station 3, for instance. Then, we have $\mathcal{A}_3^{or} = \frac{1}{1+1+1+1} = \frac{1}{4}$. Similarly, we can calculate the original airtime shares for other contending stations, and the results are listed in Table 3.

Table 3: Example Scenario: Original Airtime Shares

i (Station ID)	1	2	3	4
ϕ_i	1	1	1	1
ω_i	1	1	$\frac{1}{4}$	$\frac{1}{2}$
$(P_i - O_i)$	P_{min}	$3P_{min}$	$4P_{min}$	$4P_{min}$
$\mathcal{A}_i^{or}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$

3.2 Step II: Determine Lower Bounds of Final Airtime Shares

The second step is to determine the lower bounds ($\mathcal{A}^{bd}$) of final airtime shares based on the power-factor information as well as the energy-conservation fairness requirement.

THEOREM. *Let P_{min} denote the difference between the minimum possible transmit-power consumption and the idle power consumption, then $\mathcal{A}_i^{bd} \equiv \mathcal{A}_i^{or} \cdot \max\left\{\omega_i, \frac{P_{min}}{P_i - O_i}\right\}$ is the lower bound of station i's final airtime share $\mathcal{A}_i$.*

PROOF: Omitted due to space limitation. ∎

It is interesting to observe that, when all the wireless stations specify their power factors to be 1, we have

$$\begin{cases} \forall\, i, \ \mathcal{A}_i \geqslant \mathcal{A}_i^{bd} = \mathcal{A}_i^{or} \cdot \max\left\{1, \frac{P_{min}}{P_i - O_i}\right\} = \mathcal{A}_i^{or} \\ \sum_{i=1}^{n} \mathcal{A}_i = 1 \end{cases} \quad (9)$$
$$\implies \forall\, i, \ \mathcal{A}_i = \mathcal{A}_i^{bd} = \mathcal{A}_i^{or}.$$

In this case, airtime fairness is achieved. On the other hand, choosing power factors to be 0 means that energy-conservation fairness is the only concern. Any value between 0 and 1 for the power factor represents a compromise between airtime fairness and energy-conservation fairness.

Now, we apply *Step II* to our example scenario. Consider station 3, for instance. Then, we have $\mathcal{A}_3^{bd} \equiv \mathcal{A}_3^{or} \cdot \max\left\{\omega_3, \frac{P_{min}}{P_3 - O_3}\right\}$ $= \frac{1}{16}$. Similarly, we can calculate lower bounds of airtime shares for other contending stations, and the results are listed in Table 4.

Table 4: Example Scenario: Lower Bounds of Final Airtime Shares

i (Station ID)	1	2	3	4
ϕ_i	1	1	1	1
ω_i	1	1	$\frac{1}{4}$	$\frac{1}{2}$
$(P_i - O_i)$	P_{min}	$3P_{min}$	$4P_{min}$	$4P_{min}$
$\mathcal{A}_i^{or}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
$\mathcal{A}_i^{bd}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{16}$	$\frac{1}{8}$

3.3 Step III: Generate the Final Airtime Allocation

The third step is to generate the final airtime allocation based on the lower bound information and, again, the energy-conservation fairness requirement. Fig. 1 shows the pseudo-coded algorithm for *Step III*. The algorithm starts by using the lower bounds ($\mathcal{A}_i^{bd}$) as the initial values of final airtime shares, and stops when all ($1 - \sum_{i=1}^{n} \mathcal{A}_i^{bd}$) shares of the total airtime usage have been allocated to contending stations to yield the best energy-conservation fairness. As shown in the pseudo code, the values of $\mathcal{A}_i$'s are increased by rounds. At each round, the normalized energy consumptions ($\bar{\mathcal{E}}$) of all contending stations are first sorted in ascending order. Then, the stations with minimum $\bar{\mathcal{E}}$ values are allocated extra shares of

airtime ($\delta_{\mathcal{A}}$) in proportion to $\frac{P-O}{\phi}$, until either their normalized energy consumptions are raised to the second smallest amount or all the available airtime shares have been allocated, whichever occurs first.

Initialization: $\forall\, i, \ \mathcal{A}_i = \mathcal{A}_i^{bd}$;
while $\left(\sum_{i=1}^{n} \mathcal{A}_i < 1\right)$ {
    for $(i = 1 : n)$ $\bar{\mathcal{E}}_i = \mathcal{A}_i \cdot \frac{P_i - O_i}{\phi_i}$;
    Sort $\bar{\mathcal{E}}_i$'s in ascending order;
    $\bar{\mathcal{E}}_{min} = \min(\bar{\mathcal{E}}_i)$;
    $\bar{\mathcal{E}}_{next\_min}$ = the second-smallest $\bar{\mathcal{E}}_i$ value;
    if $\left(\bar{\mathcal{E}}_{next\_min} == \bar{\mathcal{E}}_{min}\right)$ $\bar{\mathcal{E}}_{next\_min} = \infty$;
    $\Upsilon = \left\{i, \text{ where } \bar{\mathcal{E}}_i = \bar{\mathcal{E}}_{min}\right\}$;
    for (each $i \in \Upsilon$) {
        $\delta_{\mathcal{A}_i} = \min\left\{ \frac{\left(1 - \sum_{i=1}^{n} \mathcal{A}_i\right) \cdot \frac{\phi_i}{P_i - O_i}}{\sum_{i \in \Upsilon} \frac{\phi_i}{P_i - O_i}}, \ \frac{\bar{\mathcal{E}}_{next\_min} - \bar{\mathcal{E}}_{min}}{\frac{P_i - O_i}{\phi_i}} \right\}$;
        $\mathcal{A}_i = \mathcal{A}_i + \delta_{\mathcal{A}_i}$;
    }
}

Figure 1: Pseudo-code of *Step III*.

Now, we apply *Step III* to our example scenario, and the results are listed in Table 5. One can see that, in this example, final airtime

Table 5: Example Scenario: Generate Final Airtime Allocation

i (Station ID)		1	2	3	4
ϕ_i		1	1	1	1
ω_i		1	1	$\frac{1}{4}$	$\frac{1}{2}$
$(P_i - O_i)$		P_{min}	$3P_{min}$	$4P_{min}$	$4P_{min}$
$\mathcal{A}_i^{or}$		$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{4}$
$\mathcal{A}_i^{bd}$		$\frac{1}{4}$	$\frac{1}{4}$	$\frac{1}{16}$	$\frac{1}{8}$
Round 1	$\delta_{\mathcal{A}_i}$	$\frac{1}{4}$	-	$\frac{1}{16}$	-
	$\mathcal{A}_i$	$\frac{1}{2}$	$\frac{1}{4}$	$\frac{1}{8}$	$\frac{1}{8}$

shares are obtained after one round of increments. We only recap the computation details for $\mathcal{A}_3$ as follows:
[Round 1]

$$\sum_{i=1}^{4} \mathcal{A}_i = \frac{11}{16} < 1, \quad \Upsilon = \{1, 3\}$$

$$\implies \delta_{\mathcal{A}_3} = \min\left\{ \frac{\left(1 - \frac{11}{16}\right) \cdot \frac{1}{4}}{1 + \frac{1}{4}}, \ \frac{\frac{1}{2} - \frac{1}{4}}{4} \right\} = \frac{1}{16}.$$

The final airtime shares of other contending stations can be obtained by performing similar computations. Energy-conservation fairness index of our airtime allocation is

$$\mathcal{I}_{\mathcal{E}}' = \frac{\left(\frac{1}{2} + \frac{1}{4} \cdot 3 + \frac{1}{8} \cdot 4 + \frac{1}{8} \cdot 4\right)^2}{4\left[\left(\frac{1}{2}\right)^2 + \left(\frac{1}{4} \cdot 3\right)^2 + \left(\frac{1}{8} \cdot 4\right)^2 + \left(\frac{1}{8} \cdot 4\right)^2\right]} = 0.9643,$$

while $\mathcal{I}_{\mathcal{E}}$ of the airtime-allocation scheme based on airtime fairness only is

$$\mathcal{I}_{\mathcal{E}}'' = \frac{\left(\frac{1}{4} + \frac{1}{4} \cdot 3 + \frac{1}{4} \cdot 4 + \frac{1}{4} \cdot 4\right)^2}{4\left[\left(\frac{1}{4}\right)^2 + \left(\frac{1}{4} \cdot 3\right)^2 + \left(\frac{1}{4} \cdot 4\right)^2 + \left(\frac{1}{4} \cdot 4\right)^2\right]} = 0.8571.$$

Clearly, our scheme performs better in terms of energy-conservation fairness. Notice, however, that our scheme does not achieve the perfect energy-conservation fairness. This is due to the minimum airtime shares required by the contending stations. It is easy to calculate that the following airtime allocation $\left\{\frac{6}{11}, \frac{2}{11}, \frac{3}{22}, \frac{3}{22}\right\}$ achieves

the perfect energy-conservation fairness, where station 2's airtime share ($\frac{2}{11}$) is lower than its desired minimum level ($\frac{1}{4}$).

3.4 Implementation Issues

In an 802.11e EDCA (Enhanced Distributed Channel Access) system, there are two ways of allocating airtime shares to contending stations: (i) via controlling the channel access parameters of each wireless station, or (ii) via controlling the TXOP limit[3] of each wireless station. With the first method, each station occupies the medium for the same amount of time during each access but has a different medium access frequency. With the second method, all contending stations will use the same channel access parameters, but each station will occupy the shared wireless medium for a different amount of time during its access.

Since the idea of controlling the channel access parameters has been used by the EDCA to provide service differentiation among ACs (Access Categories), if we re-use the same idea for the purpose of providing the desired fairness among contending stations belonging to the same AC, it will inevitably introduce ambiguities between coarse (among ACs) and fine (among contending stations belonging to the same AC) levels of service differentiation in an 802.11e EDCA system. Hence, to avoid unnecessary confusion, we choose to provide user-desired fairness by controlling the TXOP limits of contending stations as follows.

As listed in Table 2, L_i and R_i represent the data payload length and the transmission rate of station i, respectively. Hence, the transmission duration of a single data frame (excluding the physical- and MAC-layer overheads) by station i is $D_i = \frac{L_i}{R_i}$. Let m denote the station index such that $D_m = \max_{1 \leqslant i \leqslant n} D_i$. Then, in order to achieve the desired airtime allocation among contending stations, the number of frame transmissions per medium access by station i is $N_i = \frac{D_m}{D_i} \cdot \frac{A_i}{A_m}$. Consequently, the TXOP limit for station i can be calculated as

$$
\begin{aligned}
TXOP_i = & N_i \left(tPLCPoverhead + \frac{L_i + aMACheader}{R_i} \right) \\
& + (2N_i - 1)tSIFStime + N_i \left(tPLCPoverhead + \frac{L_{ack}}{R_{ack}} \right),
\end{aligned}
\tag{10}
$$

where L_{ack} is the Ack frame size, and R_{ack} is the Ack transmission rate. Apparently, this scheme works perfectly when N_i is an integer. When any of the N_i values is not an integer, frames must be fragmented to achieve precise airtime control.

In order to achieve the desired fairness via controlling the TXOP limits, the AP needs to (i) collect the associated weight and power-factor information as well as the transmission-strategy information from all contending stations, (ii) determine the airtime allocation and calculate the corresponding TXOP limits, and (iii) convey the TXOP limit to each station. It has been shown in [5] that the wireless stations and the AP may be able to exchange those pieces of information with the help of two newly-defined 802.11e elements: the *TSPEC (Traffic Specification)* element and the *EDCA Parameter Set Information* element.

4. PERFORMANCE EVALUATION

This section evaluates the effectiveness of our proposed energy-efficient airtime-allocation scheme via simulation.

---

[3]TXOP (Transmission Opportunity) is a new concept introduced in 802.11e. A TXOP is a time interval during which a particular station is allowed to occupy the wireless medium and transmit after it wins the medium contention. The maximum duration of a TXOP is called the TXOP limit.

4.1 Simulation Setup

We assume that each 802.11e wireless station is equipped with an 802.11b wireless network interface. Based on the power characteristics of the Cisco Aironet 350 Series Client Adaptor [27], Fig. 2 illustrates the available options for the transmission strategy of the simulated wireless network interface and the corresponding station-to-AP distances. For example, this figure shows that, when a station is 40–50m away from the AP, it transmits at 11 Mbps with 20 dBm power.
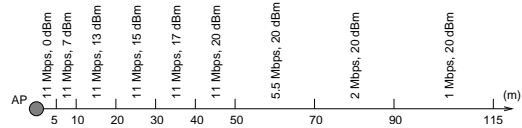

Figure 2: Relation between a station's transmission strategy and the station-to-AP distance.

We evaluate the following two schemes: (i) 802.11 DCF with which each station contends with the same channel access parameters and is only allowed to transmit one data frame per channel access, and (ii) the proposed energy-efficient airtime-allocation scheme with varying power factor ω from 1.0 to 0.0. They are compared with each other using aggregate system throughput (in Mbps), overall system energy-efficiency (in Mbits/Joule), and fairness. We conduct simulation with different numbers of contending stations and various network topologies. Each simulation run lasts for 15 minutes. Each station transmits in a greedy mode, i.e., its data queue never gets empty.

4.2 Results with Two Contending Stations

We first compare the schemes under consideration with a simple network configuration in which only two stations (STA1 and STA2) contend for the shared wireless medium and communicate with the AP. The associated weights for both stations are $\phi_1 = \phi_2 = 1$. As shown in Fig. 3, STA1 is static and 45m away from the AP. Hence, according to Fig. 2, STA1 always transmits at 11 Mbps with 20 dBm power. In contrast, STA2 is mobile and moves towards the AP. The starting and ending points are 115m and 5m from the AP, respectively. Correspondingly, STA2 adapts its transmission strategy along with its movement, from 1 Mbps rate and 20 dBm power at the starting point, to 20 Mbps rate and 0 dBm power at the ending point.
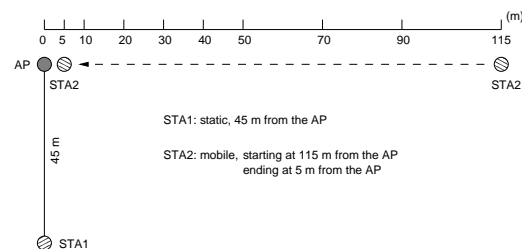

Figure 3: Two-station network configuration.

4.2.1 Results with Different STA2 Locations

As shown in Figs. 4(a) and (c), when STA2 gets closer to the AP, both the aggregate system throughput and the system energy-efficiency improve for all schemes, because STA2 can transmit at a higher rate or with a lower power level as it moves closer to the
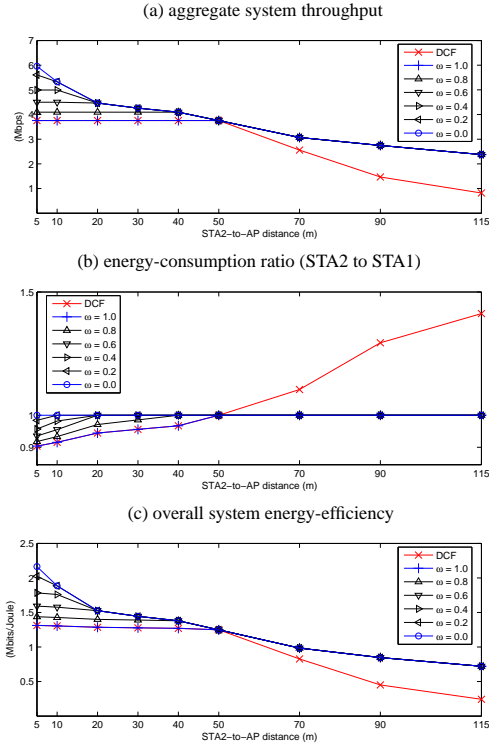
(a) aggregate system throughput

(b) energy-consumption ratio (STA2 to STA1)

(c) overall system energy-efficiency

Figure 4: Comparison for two-station network configuration.

(a) aggregate system throughput

(b) overall system energy-efficiency

Figure 5: Comparison for two-station network configuration measured over the entire STA2 movement duration.

AP. However, different schemes show different improvement patterns dictated by their respective design philosophies, which are discussed next.

We make four observations. First, the 802.11 DCF (cross points in the figure) achieves the the perfect throughput fairness regardless of the distance between STA2 and the AP. However, since DCF is not designed to handle the transmission-strategy diversity, it yields poor aggregate throughput when STA2 is far away from the AP and transmits at a lower rate than STA1.

Second, when $\omega = 1.0$ (plus points in the figure), our scheme allocates airtime shares to STA1 and STA2 in proportion to their associated weights. Since we use $\phi_1 = \phi_2 = 1$ in the simulation, both stations are assigned an equal airtime share. The airtime fairness guarantees that the throughput of the high-rate STA1 is not affected by the low transmission rate of STA2. As a result, the aggregate system throughput is improved significantly.

Third, when $\omega = 0.0$ (circle points in the figure), our scheme allocates airtime shares to STA1 and STA2 in such a way that the perfect energy-conservation fairness is achieved (see its unity energy-consumption ratio in Fig. 4(b)). In contrast, since transmit-power information is not considered in either throughput or airtime fairness, unbalanced power-consumptions can be observed in Fig. 4(b) for both schemes as STA2 moves closer to the AP and starts transmission at lower power levels than STA1. One key advantage of introducing the energy-conservation fairness is that the system energy-efficiency may be improved significantly. As one can see from Fig. 4(c), the system may deliver about 70% more data per unit of energy consumption than when throughput or airtime fairness is used. This is because, with energy-conservation fairness, more airtime shares are allocated to low-power stations that typically transmit at higher rates, as shown in Fig. 2.

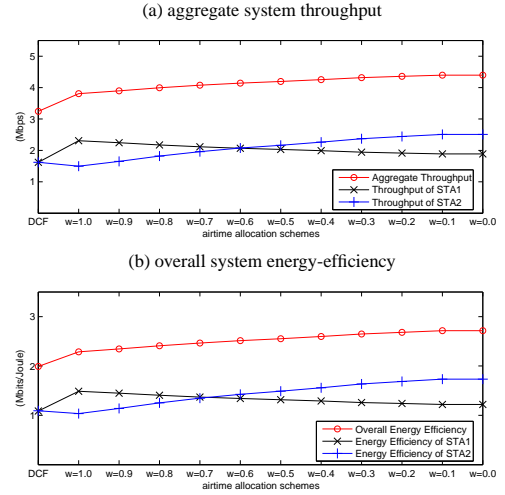Finally, any ω value between 1.0 and 0.0 represents a compro-

mise between airtime fairness and energy-conservation fairness. As discussed earlier, all three fairness notions and combinations thereof are equivalent to one another if there is no transmission-strategy diversity in the network. This can be seen from the figure that all the performance curves intersect at the point when STA2 moves to the 50m boundary and operates with the same transmission strategy as STA1.

4.2.2   Results for the Entire STA2 Movement

In this section, we compare the schemes under consideration over the entire STA2 movement duration, and results are plotted in Fig. 5. It is clear that by considering the transmission-strategy diversity in fairness, the aggregate system throughput and the overall system energy-efficiency are improved significantly.

Let's look at STA2's throughput and energy-efficiency performance, shown as plus points in the figure. We can see that, when energy-conservation fairness is the only concern ($\omega = 1.0$), STA2 performs the best. It is allowed to transmit less and "bank" its energy when it is far away from the AP. Later on, as it moves closer to the AP and the link quality improves, it may use the banked energy more efficiently by transmitting at higher rates and lower power levels. The amount of improvement varies with the network configuration, such as STA2's speed and trajectory, as well as STA1's location. In comparison, STA1 presents the best performance in terms of both throughput and energy-efficiency when the airtime fairness is considered. In this case, STA1 is guaranteed to receive its airtime share. With all other schemes, STA1's airtime share is reduced due to either the low transmission rate of STA2 (with DCF) or the low transmission power level of STA2 (with schemes when $\omega < 1$).

4.3   Results with 50 Contending Stations

The second part of our simulation considers a more realistic scenario where 50 stations are randomly placed within a circle around the AP with a 115m radius. All stations are static with an equal weight. The simulation results are plotted in Fig. 6, where each point represents an average over 100 simulation runs.

In Fig. 6(a), we compare the three fairness indices of the schemes under consideration. Obviously, DCF achieves the perfect throughput fairness but performs poorly in terms of airtime and energy conservation fairness. The perfect airtime fairness is achieved when

(a) three fairness indices

(b) aggregate system throughput

(c) overall system energy-efficiency
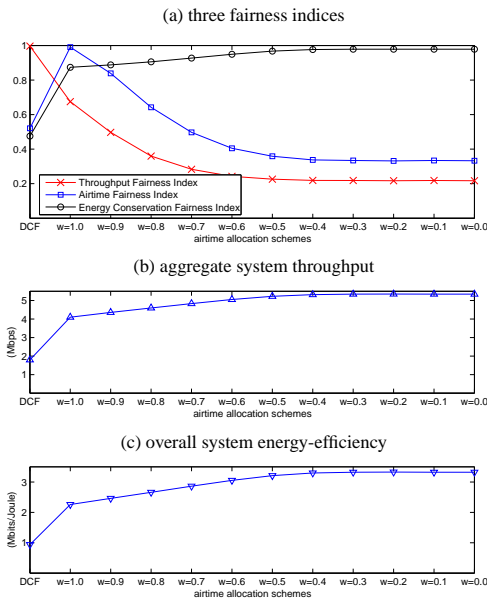
Figure 6: Comparison for random-topology network configuration.

$\omega = 1.0$, and the power factor of 0.0 yields the best energy-conservation fairness, which all conform to their respective design goals.

Energy-conservation fairness improves both aggregate system throughput and overall system energy-efficiency significantly, because its airtime allocation is determined by considering both aspects of the transmission-strategy diversity: transmission-rate diversity and transmit-power diversity. In particular, it improves the aggregate system throughput by more than 25% over that of airtime fairness and more than 2.5 times that of the DCF. It improves energy-efficiency by about 40% over that of airtime fairness and more than 3 times that of the DCF.

5. CONCLUSION

In a multi-rate multi-power-level 802.11 WLAN, different wireless stations may communicate with the AP using different transmission rates, different transmit-power levels, and/or with different data payload sizes. This phenomenon is often referred to as *transmission-strategy diversity*. In this paper, we introduced a new fairness notion, called *energy-conservation fairness*, to deal with transmission-strategy diversity, and particularly, transmit-power diversity, in multi-rate multi-power-level 802.11 WLANs. It emphasizes fair energy consumption by all contending stations in the same network. We also presented an energy-efficient airtime-allocation scheme to meet generic fairness requirements that combine airtime fairness with energy-conservation fairness. Our extensive simulation results show that, when the energy-conservation fairness is considered, both aggregate system throughput and overall system energy-efficiency can be improved significantly while all contending stations consume a similar amount of energy.

6. REFERENCES

[1] I. 802.11. *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*. Standard, IEEE, Aug. 1999.

[2] I. Aad and C. Castelluccia. Differentiation Mechanisms for IEEE 802.11. In *Proc. IEEE INFOCOM'01*, volume 1, pages 209–218, Apr. 2001.

[3] S. Agarwal, S. V. Krishnamurthy, R. K. Katz, and S. K. Dao. Distributed Power Control in Ad-Hoc Wireless Networks. In *Proc. IEEE PIMRC'01*, 2001.

[4] V. Bharghavan, A. Demers, S. Shenker, and L. Zhang. MACAW: A Media Access Protocol for Wireless LAN's. In *Proc. ACM SIGCOMM'94*, pages 212–225, London, England UK, Aug. 1994.

[5] C.-T. Chou, K. G. Shin, and S. Shankar. Contention-Based Airtime Usage Control in Multirate IEEE 802.11 Wireless LANs. *IEEE/ACM Transactions on Networking*, 16(6):1179–1192, Dec. 2006.

[6] J.-P. Ebert and A. Wolisz. Combined Tuning of RF Power and Medium Access Control for WLANs. *Mobile Networks & Applications*, 5(6):417–426, Sept. 2001.

[7] J. Gomez, A. T. Campbell, M. Naghshineh, and C. Bisdikian. Conserving Transmission Power in Wireless Ad Hoc Networks. In *Proc. IEEE ICNP'01*, pages 24–34, Nov. 2001.

[8] J. He, L. Zheng, Z. Yang, and C. T. Chou. Performance Analysis and Service Differentiation in IEEE 802.11 WLAN. In *Proc. LCN'03*, Oct. 2003.

[9] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda. Performance Anomaly of 802.11b. In *Proc. IEEE INFOCOM'03*, pages 836–843, San Francisco, CA, Mar. 2003.

[10] IEEE 802.11e. *Part 11: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements*. Supplement to IEEE 802.11 Standard-1999 Edition, 2005.

[11] E.-S. Jung and N. H. Vaidya. A Power Control MAC Protocol for Ad Hoc Networks. In *Proc. ACM MobiCom'02*, pages 36–47, Atlanta, GA, Sept. 2002.

[12] P. Lettieri and M. B. Srivastava. Adaptive Frame Length Control for Improving Wireless Link Throughput, Range, and Energy Efficiency. In *Proc. IEEE INFOCOM'98*, San Francisco, California, Mar. 1998.

[13] X. Liu, E. K. P. Chong, and N. B. Shroff. Transmission Scheduling for Efficient Wireless Utilization. In *Proc. IEEE INFOCOM'01*, Apr. 2001.

[14] S. Lu, T. Nandagopal, and V. Bharghavan. A Wireless Fair Service Algorithm for Packet Cellular Networks. In *Proc. ACM MobiCom'98*, pages 10–20, Dalla, TX, Oct. 1998.

[15] H. Luo, S. Lu, and V. Bharghavan. A New Model for Packet Scheduling in Multihop Wireless Networks. In *Proc. ACM MobiCom'00*, pages 76–86, Boston, MA, Aug. 2000.

[16] T. Nandagopal, T.-E. Kim, X. Gao, and V. Bharghavan. Achieving MAC Layer Fairness in Wireless Packet Networks. In *Proc. ACM MobiCom'00*, pages 87–98, Boston, MA, Aug. 2000.

[17] T. Nandagopal, S. Lu, and V. Bharghavan. A Unified Architecture for the Design and Evaluation of Wireless Fair Queueing Algorithms. In *Proc. ACM MobiCom'99*, pages 132–142, Seattle, WA, Aug. 1999.

[18] T. S. E. Ng, I. Stoica, and H. Zhang. Packet Fair Queueing Algorithms for Wireless Networks with Location-Dependent Errors. In *Proc. IEEE INFOCOM'98*, pages 1103–1111, San Francisco, CA, Mar. 1998.

[19] T. Ozugur, M. Naghshineh, P. Kermani, C. M. Olsen, B. Rezvani, and J. A. Copeland. Balanced Media Access Methods for Wireless Networks. In *Proc. ACM MobiCom'98*, pages 21–32, Dallas, TX, Oct. 1998.

[20] M. Pursley, H. Russell, and J. Wysocarski. Energy-Efficient Routing in Frequency-Hop Networks with Adaptive Transmission. In *Proc. IEEE MILCOM'99*, Nov. 1999.

[21] M. Pursley, H. Russell, and J. Wysocarski. Energy-Efficient Transmission and Routing Protocols for Wireless Multiple-Hop Networks and Spread-Spectrum Radios. In *Proc. EuroComm'00*, pages 1–5, 2000.

[22] D. Qiao, S. Choi, A. Jain, and K. G. Shin. MiSer: An Optimal Low-Energy Transmission Strategy for IEEE 802.11a/h. In *Proc. ACM MobiCom'03*, pages 161–175, San Diego, CA, Sept. 2003.

[23] D. Qiao and K. G. Shin. Achieving Efficient Channel Utilization and Weighted Fairness for Data Communications in IEEE 802.11 WLAN under the DCF. In *Proc. IWQoS'02*, pages 227–236, Miami Beach, FL, May 2002.

[24] P. Ramanathan and P. Agrawal. Adapting Packet Fair Queueing Algorithms to Wireless Networks. In *Proc. ACM MobiCom'98*, Dalla, TX, Oct. 1998.

[25] L. Romdhani, Q. Ni, and T. Turletti. Adaptive EDCF: Enhanced Service Differentiation for IEEE 802.11 Wireless Ad Hoc Networks. In *Proc. WCNC'03*, volume 2, pages 1373–1378, Mar. 2003.

[26] B. Sadeghi, V. Kanodia, A. Sabharwal, and E. Knightly. Opportunistic Media Access for Multirate Ad Hoc Networks. In *Proc. ACM MobiCom'02*, Sept. 2002.

[27] C. Systems. Cisco Aironet 350 Wireless LAN Client Adapter. Online Data Sheet.

[28] G. Tan and J. Guttag. Time-based Fairness Improves Performance in Multi-rate WLANs. In *Proc. USENIX'04*, Boston, MA, June 2004.

[29] N. H. Vaidya, P. Bahl, and S. Gupta. Distributed Fair Scheduling in a Wireless LAN. In *Proc. ACM MobiCom'00*, pages 167–178, Boston, MA, Aug. 2000.

[30] D.-Y. Yang, T.-J. Lee, K. Jang, J.-B. Chang, and S. Choi. Performance Enhancement of Multi-Rate IEEE 802.11 WLANs with Geographically-Scattered Stations. *IEEE Transactions on Mobile Computing*, 5(7):906–919, July 2005.

[31] J. Zhao, Z. Guo, Q. Zhang, and W. Zhu. Performance Study of MAC for Service Differentiation in IEEE 802.11. In *Proc. GLOBECOM'02*, volume 1, pages 778–782, Nov. 2002.