

Balancing Energy Efficiency and Throughput Fairness in IEEE 802.11 WLANs

Andres Garcia-Saavedra^{*,a}, Pablo Serrano^a, Albert Banchs^{a,b}, Matthias
Hollick^c

^a*Depto. de Ingeniería Telemática. Universidad Carlos III de Madrid. Avda.
Universidad, 30. 28911 Leganés, Spain*

^b*IMDEA Networks. Avda. Mediterráneo, 22. 28918 Leganés, Spain*

^c*Secure Mobile Networking Lab. Technische Universität Darmstadt. Mornewegstr. 32.
64293 Darmstadt, Germany*

Abstract

The proliferation of wireless networks based on IEEE 802.11 has resulted in a heterogenous set of devices using a variety of applications to compete for the desired service performance. Most notably, the class of highly mobile and energy constrained devices is showing high growth rates. Yet, fairness of resource allocation is still only considered in terms of achievable throughput and without considering energy efficiency. In this paper we first show that performing an energy efficient and fair resource allocation in current IEEE 802.11 WLANs is challenging, given the diversity of power consumption figures among mobile devices. We then propose a criterion to objectively balance between the most energy-efficient configuration (where all resources are given to one station) and the throughput-fair allocation (where the power consumption is not considered). We derive a closed-form expression for the optimal configuration of 802.11 with respect to this criterion. Our analysis is validated through simulations, showing that our approach betters the prevalent allocation schemes discussed in literature in terms of energy efficiency,

*Corresponding author

Email addresses: agsaaved@it.uc3m.es (Andres Garcia-Saavedra),
pablo@it.uc3m.es (Pablo Serrano), banchs@it.uc3m.es (Albert Banchs),
matthias.hollick@seemoo.tu-darmstadt.de (Matthias Hollick)

This paper is an extended version of our paper “Energy-efficient Fair Channel Access for IEEE 802.11 WLANs” [1], which appeared in IEEE WoWMoM 2011 and has been selected for a fast-track publication in Elsevier PMC

while maintaining the notion of fairness among devices. Experimental results obtained in a real-world testbed confirm the main results derived from our analysis and simulations.

Key words: 802.11, Energy efficiency, Fairness

1. Introduction

Mobile devices are increasingly equipped with multiple radios to wirelessly access communication networks such as the Internet. The IEEE 802.11 wireless local area network (WLAN) technology is dominating and deployed at large, e.g., in public hotspots, campus, or home networks. Until recently, stations operated within these networks mostly belonged to the class of notebook computers; the stations of this class shared quite similar feature specifications with respect to energy supply. As one result, modeling and optimization has mainly been focusing on bandwidth efficiency and throughput-fair bandwidth allocation (see, e.g., [2, 3]), but has not looked into energy trade-offs.

However, along with a tremendous growth in numbers, we currently witness an increasing diversity in mobile computing devices that operate on battery power to allow for untethered operation populating 802.11 networks. This includes powerful notebook computers (which might be operated on AC power), slate or tablet computers such as the iPad, netbooks, smartphones and ebook readers, personal digital assistants such as Blackberries, or embedded systems and MP3 players. For this novel set of computers, wireless and battery powered operation is the norm rather than the exception. Compared with traditional notebooks, these new devices have a substantially different energy profile. Hence, energy efficiency as an optimization goal is of paramount importance but, to the best of our knowledge, its relation with throughput optimization has received little attention so far, [4] being one of the few works to analyze these various trade-offs in wireless networking.

1.1. On the relation between energy efficiency and throughput performance

Few works have analyzed the relation between energy efficiency and throughput performance in WLANs. In [5], Bruno et al. assumed an energy consumption model in which an interface alternates between a transmitting phase (with power consumption PTX) and a receiving phase (with power consumption PRX), and showed that under these assumptions the throughput-optimal configuration and the most energy efficient configuration coincide.

Table 1: Power consumption (in Watts) for a WLAN interface when in the transmission (ρ^{tx}), reception (ρ^{rx}) and idle (ρ^{id}) states.

#	Card	ρ^{tx}	ρ^{rx}	ρ^{id}
A	Lucent WaveLan	1.650	1.400	1.150
B	SoketCom CF	0.924	0.594	0.066
C	Intel PRO 2200	1.450	0.850	0.080
D	Agilent Card Test	1.188	1.138	1.108

In [6], we used a different power consumption model [7], which introduces a third ‘listening’ phase to account for the energy consumption during the carrier sensing and backoff operations. Using this extended model, which is the same that we assume in this paper, in [6] we showed that throughput and energy efficiency constitute different optimization objectives that results in different optimal configurations of the contention parameters.

We next introduce a simple numerical example to illustrate that the two optimization criteria considered, throughput and energy efficiency, result in different configurations of the contention parameters (the energy consumption model used to derive these figures is described in detail in Section 3.1). Let us consider a toy WLAN scenario consisting of one Access Point (AP) and two associated stations, which operate using the IEEE 802.11b physical layer. The maximal fair throughput allocation is tied to the minimum Contention Window (CW_{min}) and can be obtained using, e.g., a numerical search, giving the value of $CW_{min} = 17$. However, using this configuration, substantial amounts of energy might be consumed by collisions of frames from both stations. Indeed, when optimizing the network configuration with respect to energy efficiency, we obtain significantly different values for CW_{min} that depend on the energy parameters of the interface. For the case of the four interfaces considered in Table 1, selected from previous surveys [8, 9], the resulting optimal CW_{min} configurations are $CW_A = 20$, $CW_B = 59$, $CW_C = 68$ and $CW_D = 18$.

These results illustrate the differences between the configuration that maximizes energy efficiency and the one that maximizes throughput performance, which range from almost no difference (for the case of the relatively old Lucent WaveLan interface or the Agilent card) to a value approximately three times larger (for the case of the other two interfaces). In [6] we derived

the relation between these two CW configurations, which is given by

$$CW_{ee}/CW_{th} \approx \sqrt{\rho^{rx}/\rho^{id}},$$

where CW_{ee} denotes the optimal configuration for energy efficiency, CW_{th} denotes the throughput-optimal configuration, and ρ^{rx} and ρ^{id} are two of the three parameters that characterize the energy consumption of the interface.

According to the above, for the case of a *homogeneous* scenario in which all devices share the same power consumption behavior, the only decision is to agree on the criterion to optimize (either throughput performance or energy efficiency) and configure the contention parameters accordingly, following [6]. However, as we illustrate next, energy consumption heterogeneity is rather the norm than the exception, and therefore a different approach towards optimization is required when *heterogeneous* devices populate the WLAN. This is the key motivation behind this paper and the main contribution over our previous work [6].

1.2. On the heterogeneity of IEEE 802.11 devices

We have illustrated in Table 1 how different WLAN *interfaces* present quantitative and qualitatively different values for the parameters that characterize their energy consumption, as reported in [8, 9]. In addition to these figures, here we provide real-life measurements of the energy consumption of two different 802.11 *devices*, i.e., the consumption of all the hardware and not only the WLAN interface.

We consider two different commercial, off-the-shelf devices (the detailed description of our testbed is provided in Section 5), namely, a Soekris net4826 box equipped with an Atheros card and an Alix2d2 box equipped with a Broadcom card. For each device, we measured the power consumption in four different conditions: (*i*) when the wireless card is not connected, (*ii*) when the card is plugged in but the wireless driver is not loaded, (*iii*) when the wireless driver is loaded and the station is associated to an Access Point, but no traffic is sent, and (*iv*) when the device is sending 1470 bytes UDP packets at a rate of 400 frames per second. We denote these conditions as “no card”, “off”, “idle” and “sending”, respectively, and report the measured power consumption for each device and condition in Table 2, in which we also provide in parenthesis the increment of the power consumed of a given configuration as compared to the previous configuration

These results confirm the heterogeneity of devices and wireless interfaces: not only the *base* power consumption of each device is noticeably different

Table 2: Measured power consumption (in Watts) of two different commercial, off-the-shelf devices.

Experiment	Soekris net4826	Alix 2d2
no card	2.29	2.70
off	2.58 (+0.29)	2.76 (+0.06)
idle	3.51 (+0.93)	3.69 (+0.93)
sending	4.64 (+1.13)	4.03 (+0.34)

(regardless of whether the wireless interface is connected or not), but also the power consumption associated to wireless operations differs. Indeed, the table shows that, e.g., although associating to the AP increments consumption by approximately 1 W, the increment when sending UDP traffic is almost 4 times larger with the Soekris device than with the Alix device.

1.3. Contributions of the paper

This is an extended version of the paper presented in [1]. In the following we describe the key contributions of this paper as well as the extensions with respect to [1]:

1. First, we show that given the heterogeneity of existing 802.11 devices, an unrestrained optimization of energy efficiency leads to extreme throughput unfairness. Moreover, the optimization favors those devices with worst energy efficiency. In the previous version [1] we used the numerical figures provided in [8] to identify the heterogeneity of wireless interfaces, which we have experimentally confirmed in Table 2 .
2. Second, we propose a novel criterion to balance energy-efficiency and throughput fairness, namely, the energy-efficiency proportional fairness (EF) criterion.
3. Third, we derive the optimal configuration for 802.11 WLANs according to this criterion, and validate it against exhaustive searches on the configuration space considering four representative WLAN interfaces (in [1] we considered three) .
4. Finally, we confirm the above results, obtained analytically and via simulations, through experimentation in a small-sized testbed. As compared to the previous version of the paper [1], this contribution is entirely new .

The rest of the paper is organized as follows. In Section 2 we motivate the need to balance energy efficiency and throughput fairness, and propose a criterion for fair and energy efficient channel access, namely, the EF criterion. In Section 3 we introduce the energy consumption model used for the case of 802.11 WLANs, and derive a closed-form expression to optimize the network performance according to the EF criterion. The accuracy of the model and the effectiveness of the proposed configuration are extensively evaluated in Section 4. In Section 5 we report a set of experiments using a small-sized testbed composed of commercial, off-the-self devices (COTS), which confirm the results of our analysis. Finally, Section 6 concludes the paper.

2. Balancing energy efficiency and throughput fairness: the EF criterion

We have seen that, in an homogeneous deployment, the only challenge is to select the performance parameter to optimize, namely throughput or energy efficiency, and configure the WLAN accordingly. In this section we show that, given the existing heterogeneity of wireless devices reported above, performing a *good* allocation of wireless resources among devices is challenging. We then propose a criterion to balance between the lack of energy considerations of a throughput-based allocation and the extreme unfairness of a purely energy-based allocation. Throughout the paper we will denote with η the energy efficiency of the WLAN, and with η_i the energy efficiency of a given station i , i.e.,

$$\eta = \frac{\text{throughput}(WLAN)}{\text{power}(WLAN)} \quad \eta_i = \frac{\text{throughput}_i}{\text{power}_i}. \quad (1)$$

2.1. The need to balance energy efficiency and throughput fairness

Let us consider the same WLAN scenario as in the previous section with one AP and two stations, namely STA₁ and STA₂. The power consumption of STA₁ (STA₂) is modeled after the parameters of interface A (B) from Table 1. We denote with CW_1 (CW_2) the CW_{min} configuration used by STA₁ (STA₂), and consider two different strategies to configure these parameters:

- “Throughput Strategy”: We set $CW_1 = CW_2$, in order to have a fair share of the wireless resources, and perform a sweep on the $CW = \{8, 1024\}$ parameter space to choose the value that maximizes throughput.




Figure 1: Performance of a simple heterogeneous WLAN

- “Efficiency Strategy”: We let CW_1 and CW_2 diverge, and perform a sweep on the $CW = \{8, 1024\}$ parameter space to find the configuration that maximizes the energy efficiency η of the WLAN.

In order to obtain the configurations resulting from the above strategies we use our event-driven simulator of [10], extended with the energy consumption model described in Section 3.1, and perform numerical searches on the parameter space to identify the best-performing scenario (a more detailed description of the simulation tool is provided in Section 4). For the first strategy the resulting optimal CW value, as already provided in the previous section, is $CW = 17$. For the second strategy, the resulting configuration is $CW = \{CW_1, CW_2\} = \{8, 1024\}$. We report the obtained values of throughput and energy efficiency in Fig. 1, with the following main results:

- The first strategy provides a bandwidth-fair allocation where both stations receive the same throughput, while the overall energy efficiency is 3.48 Megabit per Joule (MbpJ).
- The second strategy results in an energy-efficiency improvement of

approximately 10%, while the resulting throughput allocation is extremely unfair, as STA₂ is practically starved.

The fact that the most energy-efficient allocation is obtained using an extremely unfair allocation is caused by the CSMA-based channel access scheme, as choking one interface will prevent the energy wastage caused by collisions. The price to pay for increasing the efficiency is then to introduce unfairness. However, it is interesting to observe that with this strategy *the starved station is the one with the most efficient interface*. Although this could be striking at first, it can be easily explained as follows. Given that each interface consumes a minimum power as given by the ρ^{id} parameter, there are two possible configurations with no collisions: (i) to choke STA₁, resulting in the approximately power consumption of: $\rho_B^{tx} + \rho_A^{id} \approx 2.1 W$, (ii) to choke STA₂, with the power consumption: $\rho_A^{tx} + \rho_B^{id} \approx 1.7 W$. The resulting optimal configuration, then, penalizes the more efficient station in order to provide the largest energy savings.

This simple scenario serves to illustrate the risks of using a naïve strategy to optimize the overall energy efficiency: not only it results in an extremely unfair throughput allocation, but also it *penalizes the most energy-efficient interface*. On the other hand, it is clear that the use of throughput-only allocation criteria, while resulting in throughput-fair allocations, do not consider energy efficiency at all as they do not take into account the power consumption of the interface.

Based on the above, we claim that a trade-off between energy efficiency maximization and throughput fairness is needed. In the following we present our proposed criterion to define this trade-off, namely, the EF criterion.²

2.2. The EF criterion

The use of network-wide energy efficiency figures, as we have seen in the previous section, is not well suited to properly address general (i.e., hetero-

²Note that we have only considered power consumption figures, and not parameters such as, e.g., the remaining battery capacity. Although such battery parameters have been considered before in energy-related scenarios (e.g., in [11]), they are not well suited for the scenarios that we envision. Indeed, the approach that we propose provides an incentive to energy-efficient devices by favoring them over inefficient ones. In contrast, a solution that favored battery constrained devices would incentivize battery limited devices, which would harm the overall performance. Following this reasoning, in this paper we only focus on the energy efficiency of the different wireless interfaces implementing the MAC protocol.

geneous) scenarios. The use of throughput-based approaches, on the other hand, does not consider the impact of the different power consumption parameters and therefore may result in energy wastage. We argue that a trade-off between these two is needed.

In order to specify this trade-off, we build on the per-station energy efficiency η_i defined in (1). Note that η_i provides the throughput the station i is successfully transmitting weighted by the energy the station has to spend and, therefore, it partially takes into account if a station is being choked. This way, for our toy example of the previous section (Figure 1), the resulting values for the first configuration strategy are $\eta_i = \{5.54, 2.54\}$ MbpJ, while for the case of the second strategy the values are $\eta_i = \{5.02, 0.11\}$ MbpJ.

Based on these η_i variables, our challenge is to define an appropriate criterion for their configuration. To this aim, note that we have a two-fold objective: on one hand, we want to maximize the overall efficiency (denoted as η) in the WLAN; on the other hand, we want to preserve some degree of fairness between the η_i 's, thus avoiding that any station is starved. In order to solve this tradeoff, Kelly's *proportional fairness* criterion [12] is well accepted in the literature for similar scenarios. This criterion was originally proposed in the context of wired networks, and has been widely used to address a variety of throughput fairness issues [13] including other fairness problems of wireless packet networks [14, 15]. This criterion is defined as follows. A throughput allocation $\{r_1, \dots, r_n\}$ is proportionally fair if it is feasible, and for any other feasible allocation $\{r_1^*, \dots, r_n^*\}$ the aggregate of proportional changes is not positive, i.e.,

$$\sum_i \frac{r_i^* - r_i}{r_i} \leq 0 \quad (2)$$

Note that, with the above definition, in a two station scenario the throughput of one station would be decreased by say 10% only as long as this allowed an increase in the throughput of the other station of more than 10%, which represents a balance between two extreme allocations (i.e., throughput is equally shared, or throughput is given to the most efficient station). To investigate the proportional fair allocation further, we consider a small perturbation around the proportional fair allocation $r_i \rightarrow r_i + dr_i$. From (2),

$$\sum_i \frac{dr_i}{r_i} \leq 0 \implies \sum_i (\log(r_i))' dr_i \leq 0$$

It follows from the above that the proportional fair allocation represents a local maximum of the function $\sum \log(r_i)$. Since this is a concave function, it has only one maximum, and therefore the local maximum is also the global maximum. We can identify the proportional fair (PF) allocation with the one that maximizes the sum of the logarithms:

$$PF \iff \max \sum \log(r_i)$$

In this paper, following the previous works of [13, 14, 15] we advocate for the use of the PF criterion to solve the fairness issue that arises in a WLAN with heterogeneous stations. In particular, we propose to use the energy-efficiency proportional fairness (EF) criterion, based on the maximization of the sum of the per-station energy efficiency, i.e.,

$$EF \iff \max \sum \log(\eta_i) \tag{3}$$

To illustrate why the use of the EF criterion prevents extremely unfair allocations while supporting energy-efficient configurations, we consider the same heterogeneous scenario with one AP and two different stations modeled after the power consumption figures of Interfaces A and B from Table 1. In order to analyze different configurations of the CW , we set $CW_2 = kCW_1$ with k ranging from 0.4 to 1.6, and for each k value we perform a sweep on the $CW_1 = \{1, 4096\}$ to obtain the configuration that maximizes the overall efficiency. For each resulting configuration we compute the throughput of each station and the EF value given by (3).³ Results are shown in Fig. 2, and can be summarized as follows:

- Large CW_2/CW_1 ratios increase the overall efficiency η , but lead to the starvation of STA₂, as can be seen from the R_2/R_1 ratio. This is the result that we have seen in the previous section, namely, that the most energy-efficient configuration is the one that chokes the most efficient interface.
- However, the value of EF is not maximized for such extremely unfair allocations, but instead the maximum is reached when $k \approx 1.15$. From

³Note that, for the sake of readability, throughout the paper we use EF to refer *both* to the quantity $\sum \log(\eta_i)$ resulting from a particular configuration, and to the criterion that maximizes this value. The distinction will be clear based on the context.




Figure 2: Performance of a WLAN with two stations for different CW configurations.

this point on, the relative increase in η_1 (η_2) is not compensated by the relative decrease of η_2 (η_1) and, therefore, the allocation is not EF-optimal.

In our toy example, the η -optimal allocation is given by the configuration $CW = \{3, 384\}$, which provides an overall efficiency $\eta = 3.82$ MbpJ and a throughput allocation $R_i = \{8.23, 0.06\}$ Mbps. On the other hand, the EF-optimal configuration $CW = \{26, 30\}$ results in the following performance figures: $\eta = 3.49$ MbpJ and $R_i = \{3.97, 3.47\}$ Mbps. For this case, then, the EF-optimal configuration trade-offs an 8.6% reduction in the overall efficiency in order to improve throughput fairness from 0.51 to 0.995, computed according to Jain's fairness index [16]. Furthermore, the EF-optimal configuration of the CW is not only different from the maximum throughput allocation ($CW = 17$), but also from the the case of maximum energy efficiency for homogeneous scenarios ($CW = 59$ for the case of interface B, $CW = 20$ for interface A).

From these results we conclude that our EF-criterion defines a trade-off between a fair throughput allocation and an energy-efficiency configuration.

Although the rest of the paper is devoted to the case of 802.11 WLANs, we note that the criterion could be applied to any scenario with heterogeneous interfaces.

3. EF-optimal configuration for 802.11 WLANs

In this section we first introduce an accurate yet complex model to characterize the energy consumption of a WLAN in which stations have different contention parameters. We subsequently present a simpler model that sacrifices accuracy for analytical tractability, and then derive the EF-optimal configuration based on this simplified model.

3.1. Energy consumption model

Our model assumes an IEEE 802.11 WLAN with N stations sharing the wireless channel, all of them using the same modulation and coding scheme (MCS). We assume saturation conditions, i.e., stations always have a frame ready for transmission, in order to analyze the most stringent scenario in terms of fairness (we will discuss the impact of heterogeneous coding schemes and non-saturation conditions in Section 3.3). We also assume that the only reason for frame loss is a collision, and that upon accessing the channel stations transmit a frame of fixed size L .

We denote with CW_{min}^i the CW_{min} used by station i . We first obtain the probability that a station i with minimum contention window CW_{min}^i transmits upon a backoff counter decrement τ_i by means of the following equation given by [17]

$$\tau_i = \frac{2}{1 + CW_{min}^i + p_i CW_{min}^i \sum_{j=0}^{m-1} (2p_i)^j},$$

where m is a parameter that specifies the maximum size of the CW ($CW_{max} = 2^m CW_{min}$) and p_i is the probability that a transmission attempt of station i collides. This probability can be computed as

$$p_i = 1 - \prod_{j \neq i} (1 - \tau_j).$$

The above constitutes a system of non-linear equations that can be solved numerically (see [18] for more details), giving the values of the τ_i 's.

To model the energy consumption of the WLAN we follow a similar approach to the one of [7], extending our previous model of [6] to account for the heterogeneity of the scenario, with station i having the set of power consumption figures $\{\rho_i^{tx}, \rho_i^{rx}, \rho_i^{id}\}$. These parameters represent the power consumption when the interface is transmitting, receiving or in the idle state, respectively. Based on the transmission probabilities τ_i 's, we compute the energy consumed per slot by station i , denoted by e_i , by applying the total probability theorem as follows:

$$e_i = \sum_{j \in \Theta} E_i(j)p(j) \quad (4)$$

where Θ is the set of events that can take place in a single timeslot⁴, while $E_i(j)$ and $p(j)$ are the energy consumed in case of event j and its probability, respectively. The set Θ of events and their probabilities is listed as follows:

- The slot is empty, $p(e)$
- There is a success from the considered station, $p(s, i)$
- There is a success from another station, $p(s, \neg i)$
- There is a collision and the considered station is involved, $p(c, i)$
- There is a collision but the considered station is not involved, $p(c, \neg i)$

The probability of each event can be computed based on τ_i 's as follows

$$\begin{aligned} p(e) &= \prod (1 - \tau_j) & p(s, i) &= \tau_i \prod_{j \neq i} (1 - \tau_j) \\ p(s, \neg i) &= \sum_{j \neq i} \tau_j \prod_{k \neq j} (1 - \tau_k) & p(c, i) &= \tau_i (1 - \prod_{j \neq i} (1 - \tau_j)) \\ p(c, \neg i) &= 1 - \tau_i - p_e - p_{s, \neg i} \end{aligned}$$

While the energy consumed by station i for each of the previous events can be computed as

$$E_i(e) = \rho_i^{id} T_e$$

⁴A timeslot is defined as the amount of time between two backoff counter decrements of a station, see [17].

Table 3: Power consumed (in mJ) per event for the interfaces of Table 1 and 802.11b

#	$E(e)$	$E(s, i)$	$E(c, i)$	$E(s, \neg i)$	$E(c, \neg i)$
A	0.0230	2.2834	2.2454	1.9801	1.9421
B	0.0013	1.2151	1.1349	0.8148	0.7346
C	0.0016	1.8930	1.7759	1.1651	1.0481
D	0.0222	1.6811	1.6766	1.6207	1.6162

$$E_i(s, i) = \rho_i^{tx} T_s + \rho_i^{rx} T_{ack} + \rho_i^{id} (SIFS + DIFS)$$

$$E_i(s, \neg i) = \rho_i^{rx} (T_s + T_{ack}) + \rho_i^{id} (SIFS + DIFS)$$

$$E_i(c, i) = \rho_i^{tx} T_s + \rho_i^{id} EIFS$$

$$E_i(c, \neg i) = \rho_i^{rx} T_s + \rho_i^{id} EIFS$$

where T_e is the duration of an empty slot time, $SIFS$, $DIFS$ and $EIFS$ are constants defined by the 802.11 standard, and T_s and T_{ack} are the transmission durations of a frame of size L and the acknowledgement frame, respectively, which can be computed as

$$T_s = T_{PLCP} + \frac{H + L}{C} \quad T_{ack} = T_{PLCP} + \frac{ACK}{C}$$

where T_{PLCP} is the length of the frame preamble, H is the frame header, C the transmission rate being used, and ACK represents the length of an acknowledgement frame.

Given the above expression for the energy consumption of station i in a timeslot, we can express the energy efficiency of station i as the ratio between the bits successfully transmitted over the energy consumed in a slot time:

$$\eta_i = \frac{p(s, i)L}{e_i} \quad (5)$$

It can be seen that the full expression for η_i consists of the sum of several terms that non-linearly depend on the τ_i 's. In order to improve the analytical tractability of the model, we quantify the energy consumed per timeslot for the three interfaces we consider in Table 3. Based on the observed results, we make the following approximations:

$$E(s, i) \approx E(c, i) \quad E(s, \neg i) \approx E(c, \neg i),$$

which supports the following approximate expression for (4)

$$\hat{e}_i = p_e E_i(e) + \tau_i E_i(s, i) + (1 - p_e - \tau_i) E_i(s, \neg i). \quad (6)$$

Note that the use of (6) results in an overestimation of the power consumed, as for the terms being approximated we take the largest of them. We further rearrange (6) as

$$\hat{e}_i = E_i(s, \neg i)(1 - \alpha_i p(e) + \beta_i \tau_i)$$

where we introduce the (non-negative) parameters α_i and β_i , used to quantify the relative energy consumed when idling or transmitting over the case when there is a transmission from a station different from i , i.e.,

$$\alpha_i = 1 - \frac{E_i(e)}{E_i(s, \neg i)} \quad \beta_i = \frac{E_i(s, i)}{E_i(s, \neg i)} - 1$$

Note that we denote with η_i the energy efficiency as computed with the use of (4) and with $\hat{\eta}_i$ the efficiency computed using (6). In Section 4.1 we assess the accuracy of both expressions to model the energy consumption and efficiency in a heterogeneous WLAN.

3.2. EF configuration

Based on the energy consumption model presented in the previous section, in this section we derive the configuration that optimizes WLAN performance according to our EF criterion. We start with the following expression for the energy efficiency $\hat{\eta}_i$ as derived in the previous section:

$$\hat{\eta}_i = \frac{L}{E_i(s, \neg i)} \frac{p(s, i)}{1 - \alpha_i p(e) + \beta_i \tau_i}$$

Computing the EF-optimal configuration requires to find the τ_i 's maximizing the efficiency fairness, i.e.,

$$\max \sum_i \log \hat{\eta}_i$$

To find this configuration, we first perform the following partial derivatives and set them to zero

$$\frac{\partial}{\partial \tau_k} \sum_i \log \hat{\eta}_i = 0 \quad , \quad \forall k,$$

that results in the following expression

$$\frac{1}{\tau_k} - \frac{N-1}{1-\tau_k} - \frac{\alpha_k \prod_{j \neq k} (1-\tau_j) + \beta_k}{1 - \alpha_k p(e) + \beta_k \tau_k} - \sum_{i \neq k} \frac{\alpha_i \prod_{j \neq k} (1-\tau_j)}{1 - \alpha_i p(e) + \beta_i \tau_i} = 0$$

Multiplying both sides by $(1 - \tau_k)$ and re-arranging terms results in the following expression

$$\frac{1}{\tau_k} = \frac{\beta_k(1-\tau_k)}{1 - \alpha_k p(e) + \beta_k \tau_k} + \sum_{\forall i} \frac{1 + \beta_i \tau_i}{1 - \alpha_i p(e) + \beta_i \tau_i} \approx \sum_{\forall i} \frac{1 + \beta_i \tau_i}{1 - \alpha_i p(e) + \beta_i \tau_i}.$$

Therefore, the τ_k that provides the EF-optimal configuration does not depend on the k , but it is the same for all stations. We have therefore one first result stating that, in order to achieve an EF-optimal configuration in 802.11 WLANs, stations have to fairly share the channel,⁵ i.e.,

$$\tau_i \approx \tau_k \quad \forall i, k \quad (7)$$

The remaining challenge is therefore to compute the optimal transmission probability (from now on we will write $\tau_i = \tau \forall i$). Because of the logarithm's properties, the maximization problem can be reformulated with the product of each station's efficiency, i.e.,

$$\max \sum_i \log \eta_i \iff \max \prod_i \eta_i$$

Under the assumptions (i) $\tau \ll 1$, which is reasonable in optimal operation as large τ values would lead to a high collision probability, and (ii) $\beta_i < 1$, which is also reasonable given the values from Table 3, we can approximate $\hat{\eta}_i$ as

$$\hat{\eta}_i = \frac{L}{E_i(s, \neg i)} \frac{\tau(1-\tau)^{N-1}}{1 - \alpha_i p(e) + \beta_i \tau} \approx \frac{L}{E_i(s, \neg i)} \frac{\tau(1-\tau)^{N-1}}{1 - \alpha_i p(e)}$$

By making the approximation

$$\prod_i (1 - \alpha_i p(e)) \approx \left(1 - \frac{\sum \alpha_i}{N} p(e)\right)^N$$

⁵Note that we already saw for the case of two stations that the optimal ratio between CW was $k \approx 1.15$.

the EF-optimal configuration can be computed by maximizing

$$\max \prod_i \eta_i \iff \max \frac{(\tau(1-\tau)^{N-1})^N L^N}{(\prod_i E_i(s, \neg i)) (1 - p_e \frac{\sum_i \alpha_i}{N})^N}$$

Therefore, the optimal configuration for the τ 's can be obtained by maximizing the following expression

$$\max \frac{\tau(1-\tau)^{N-1}}{1 - p_e \frac{\sum_i \alpha_i}{N}}$$

Performing the derivative and making it equal to zero yields

$$\begin{aligned} ((1-\tau)^{N-1} - (N-1)\tau(1-\tau)^{N-2})(1 - (1-\tau)^N \frac{\sum_i \alpha_i}{N}) = \\ N(1-\tau)^{N-1} \frac{\sum_i \alpha_i}{N} \tau(1-\tau)^N \end{aligned}$$

The above can be solved using a second-order Taylor expansion of $(1-\tau)^N$, that results in the following approximate solution for τ^*

$$\tau^* \approx \frac{1}{N} \sqrt{2 \left(\frac{N}{\sum \alpha_i} - 1 \right)} \approx \frac{1}{N} \sqrt{2 \frac{T_e}{T_s} \left(\frac{1}{N} \sum \frac{\rho_i^{id}}{\rho_i^{rx}} \right)} \quad (8)$$

Therefore, an AP that gathers the ρ parameters of all N stations in the WLAN could compute the CW (with $CW_{min} = CW_{max}$) that provides the optimal energy-fair configuration as follows:

$$CW^* = \frac{2}{\tau^*} - 1$$

Remark 1: For the case of *homogeneous* WLANs, where all stations have the same set of ρ parameters, (8) results in the expression that we already derived in [6]:

$$\tau^* \approx \frac{1}{N} \sqrt{\frac{2\rho^{id}T_e}{\rho^{rx}T_s}}$$

Remark 2: One major disadvantages of (8) is that it requires fetching the $\{\rho^{id}, \rho^{rx}\}$ parameters of all WLAN stations, which could be difficult in practice. In order to tackle this, we make the following *coarse* approximation

$$\sqrt{\rho^{id}/\rho^{rx}} \approx 1,$$

which results in the following approximate expression for the optimal τ

$$\hat{\tau}^* \approx \frac{1}{N} \sqrt{\frac{2T_e}{T_s}}. \quad (9)$$

In the next section, after the performance validation of the energy consumption model, we assess the EF performance of a WLAN configured using (8) and (9), and compare it against exhaustive searches in the CW space as well as the default standard configuration.

3.3. Non-saturation and heterogeneous modulation rates

Given that our aim is to compute the optimal CW configuration, we have assumed saturated stations throughout the previous sections, as in these conditions the CW has the most noticeable impact on performance. Indeed, when the network load is not high and stations are not saturated, we have seen in our previous work [10] that the transmission probabilities τ 's, and hence the energy consumption, are independent of the CW setting, and therefore optimizing the CW configuration is not critical in this scenario. For a mixed scenario in which some stations are saturated and other stations are not saturated, a first approximation to the optimal configuration would consist on using the expression of (8) but taking into account the saturated stations only, since (following our findings in [10]) the other stations have a much smaller impact on the overall behavior. In this way, an Access Point could e.g. periodically estimate the number of saturated stations in the WLAN, and compute the corresponding CW configuration.

We have also assumed that all stations use the same modulation rate. We note, though, that the EF criterion specified by (3) is not tied to this rate-homogeneous scenario, and therefore it could also be applied in a multi-rate scenario.⁶ Based on the expression for the EF criterion, we argue that a station using a low MCS will see its throughput decreased, as its longer transmission time has a two-fold impact: (i) its energy consumption increases due to the increased T_s , and (ii) the average slot time seen by the other stations is also longer, thus reducing the overall performance. Still, given that the EF criterion prevents starvation for any station, the resulting configuration will lead to a satisfactory allocation, achieving a good tradeoff between efficiency

⁶Indeed, in our previous work [15] we used Kelly's proportional fairness criterion to achieve a good allocation in a multi-rate WLAN suffering from the *performance anomaly*.

and fairness. In fact, if we consider the case of a homogeneous scenario with respect to energy consumption parameters, but with different modulation rates, we would obtain a similar allocation to the one of [15]. We leave for future work the derivation of the closed-form expressions for the optimal EF configuration in the above conditions.

4. Simulation-based performance evaluation

In this section we assess the accuracy of the energy consumption model, as well as the performance obtained using the configuration strategies derived in the previous section. For this purpose we have extended the simulator used in [10], which is an event-driven simulator that models the details of the 802.11 MAC protocol with high accuracy for each station, with the energy consumption model presented in Section 3.1. The simulations are performed for a WLAN with the MAC layer parameters of IEEE 802.11b, assuming a channel in which frames are only lost due to collisions, and all stations always have a $L=1470$ byte frame ready for transmission. For each considered scenario we provide the average of 10 simulation runs.

4.1. Validation of the energy consumption analysis

To validate the accuracy of the analytical models we first consider a WLAN using the standard DCF configuration with N stations, where one fourth of the stations is modeled after interface A, B, C and D of Table 1, respectively. We compute the total energy efficiency as given by simulations (“Simulation”), the analytical model of (5) (“Model”) and the use of the approximate expression \hat{e}_i (6) (“Approx.”), with the results represented in Fig. 3.

The figure shows that both models are able to predict WLAN energy behavior, as analytical results closely follow those from simulations. It can be seen as well that the energy efficiency η rapidly decreases with N (note that the y-axis is in log scale), a result caused by the increase in the number of collisions for the static DCF configuration, and that the approximate model slightly underestimates the overall efficiency, because it overestimates the energy consumed in a timeslot.

Despite the accuracy of both models, it should be noted that our aim is not to predict the WLAN behavior in terms of energy consumption, but to derive the configuration that maximizes the EF performance. To validate if the models are well suited to this aim, we perform the following experiment:




Figure 3: Overall energy efficiency of a WLAN with N stations.

for a varying number N of stations, we set $CW_{min} = CW_{max}$ and perform a search on the CW of stations A, B, C and D (denoted with CW_A , CW_B , CW_C and CW_D , respectively) to find the configuration that maximizes EF performance. This search is done (i) using simulations and (ii) by means of the approximate energy consumption model given by (6). The results are depicted in Fig. 4, where we also plot for comparison purposes the CW that optimizes throughput performance.

These results further confirm that the throughput-optimal and the EF-optimal configuration are obtained with significantly different values of the CW . Furthermore, we confirm that the approximate model for the energy consumption can be used to derive the configuration that maximizes the EF performance, as simulations and numerical searches provides very similar CW values. Note that the results from Fig. 4 also validate the relation obtained in (7), as the resulting CW 's values for the four different interfaces are very similar.




Figure 4: Optimal CW configuration for a heterogeneous WLAN.

4.2. Validation of the proposed EF configurations

We next validate the performance of our configuration for a heterogeneous WLAN scenario (note that we already addressed in detail homogeneous scenarios in [6]) with different mixtures of the interfaces listed in Table 1. We denote with N_A , N_B , N_C and N_D the number of WLAN stations with the power characteristics of interfaces A, B, C and D from Table 1, respectively. In order to gain insight the trade-off defined by the EF criterion, we first consider a topology with $N = 20$ stations, with 5 station per considered interface, and compute the throughput per station and the overall energy efficiency for the following configurations:

- The default standard configuration, denoted as “DCF”.
- The configuration given by (8), denoted as “EF”.
- The configuration maximizing the overall energy efficiency in the WLAN, denoted as “Max. Efficiency” (this is obtained through a numerical search on the CW parameter space).




Figure 5: Throughput and energy efficiency performance in a heterogeneous scenario with 20 stations.

The results are illustrated in Fig 5. The figure confirms that, on the one hand, “DCF” provides a fair channel access but with poor energy efficiency performance, while on the other hand the “Max. Efficiency” configuration boosts the bits per Joule ratio, but at the expense of a fair distribution of the channel access time. Our EF configuration sits in between these two extremes, increasing the energy efficiency of the network while providing a fair throughput distribution.

We next assess the performance of the two proposed configuration rules, namely (8) and (9), in terms of the EF value as given by (3). To this aim, in addition to the “DCF” and “EF” configurations described above, we consider the following two configurations:

- The configuration given by (9), denoted as “Approx.”.
- The maximum achievable EF performance resulting from an exhaustive search on the CW parameter space, denoted as “Exhaustive”.

We consider various scenarios with different mixtures of the number of stations equipped with a specific interface. The resulting EF values for the four considered configurations are given in Table 4, and can be summarized as follows:

Table 4: EF Performance of different approaches.

Scenario				EF Performance			
N_A	N_B	N_C	N_D	DCF	Approx.	EF	Exhaustive
5	5	5	5	-29.02	-18.49	-18.28	-18.27
5	5	5	10	-47.58	-29.78	-29.59	-29.58
5	5	10	5	-45.75	-27.89	-27.55	-27.54
5	5	10	10	-67.39	-40.18	-39.88	-39.87
5	10	5	5	-43.95	-26.09	-25.75	-25.74
5	10	5	10	-65.63	-38.39	-38.09	-38.07
5	10	10	5	-63.75	-36.49	-35.99	-35.99
5	10	10	10	-88.45	-49.62	-49.19	-49.17
10	5	5	5	-48.50	-30.73	-30.53	-30.52
10	5	5	10	-70.20	-43.04	-42.85	-42.83
10	5	10	5	-68.35	-41.14	-40.81	-40.80
10	5	10	10	-93.07	-54.28	-53.98	-53.96
10	10	5	5	-66.54	-39.34	-39.02	-39.01
10	10	5	10	-91.29	-52.49	-52.19	-52.17
10	10	10	5	-89.41	-50.58	-50.11	-50.10
10	10	10	10	-117.15	-64.44	-64.02	-64.00

- The performance of the default standard configuration is poor, worsening as the total number of stations increases, as most of the resources are wasted in collisions.
- Our configuration provides EF values very close to the ones achievable by means of the exhaustive search. Indeed, in all cases the differences between “EF” and “Exhaustive” are almost negligible, this way proving the ability of (8) to drive the WLAN to the EF-optimal point of operation.
- When the energy consumption information is not available, a WLAN configured according to (9) provides performance values that, although smaller than the maximum achievable ones, significantly outperforms the ones derived from the use of the standard configuration. The larger N_A the better this approximation results, given that ρ^{rx} and ρ^{id} are very similar for this interface.

We conclude that these results confirm the validity of (8) to provide the EF-optimal configuration in heterogeneous WLANs, as well the good performance from the approximate expression (9).

5. Testbed-based performance evaluation

In the previous section we have considered only simulations to quantify throughput and energy efficiency for a variety of scenarios, including different interfaces and different configurations of the contention parameters. In this section we present experimental results from a real-world testbed composed of 802.11g COTS devices that confirm the main results of our work, namely:

- Existing 802.11 COTS devices, in particular PC-boxes, present different characteristics with respect to their power consumption behavior.
- Maximizing the overall energy efficiency of the WLAN, oblivious to any other consideration, results in extreme unfairness.
- The EF-optimal configuration provides an adequate trade-off between the standard, throughput-fair allocation and the most energy-efficient configuration.

We present our testbed and the measurement methodology in the next subsection, and then provide the results from our experiments in the following subsection.

5.1. Testbed description

In the following we describe the testbed and measurement methodology used to obtain our experimental results, including the two different commercial, off-the-shelf devices, the device used to measure power consumption, and the handling of the uncertainties introduced in the measurement process.

As systems under test, we used the following two devices:

- **Soekris.** The Soekris net4826-48 device⁷ is a low-power, low-cost computer equipped with 233MHz AMD Geode SC1100 CPU, 2 Mini-PCI sockets, 128 Mbyte SDRAM and 256 Mbyte compact flash circuits for data storage, which we extend with a 2 GB USB drive. We installed an Atheros AR5414-based 802.11a/b/g card as a wireless interface. As software platform we installed Gentoo 10.0 Linux (kernel 2.6.24) and the popular MadWifi open-source WLAN driver⁸ (version v0.9.4).
- **Alix.** The Alix2d2 device⁹ is another popular low-cost computer. It is equipped with a Geode LX800 AMD 500 MHz CPU, 256 MB DDR DRAM, 2 Mini-PCI sockets and a CompactFlash socket where we attached a 4 GB card. Its wireless interface is a Broadcom BCM94318MPG-based 802.11b/g MiniPCI card. As software platform we installed Ubuntu 9.10 Linux (kernel 2.6.29), using its b43 WLAN driver.

In order to measure their energy consumption, we use a PCE PA-6000 power meter device.¹⁰ The power meter is connected in series between an AC or DC power source and the system under test, to provide instantaneous values of current, voltage and power factor (among other parameters), at a sample rate of approximately 1 sample/second. In addition to its graphical display, the PA-6000 features an RS-232 interface to support automated gathering of the measured values for later processing.

Our testbed configuration is illustrated in Fig. 6. The systems under test are powered through the PCA PA-6000 device, which is connected in series to a standard 12-V battery and whose internal circuits are powered by regular

⁷<http://www.soekris.com/>

⁸<http://madwifi-project.org/>

⁹<http://www.pcengines.ch/>

¹⁰<http://www.industrial-needs.com/technical-data/power-analyser-PCE-PA-6000.htm>

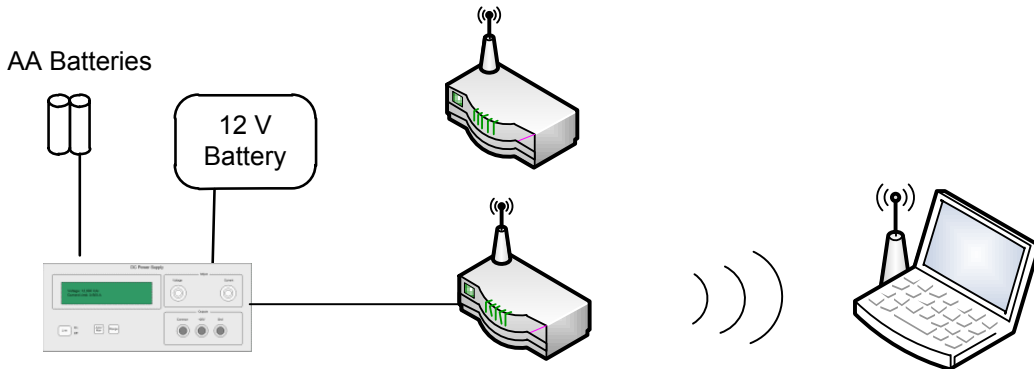


Figure 6: Testbed deployment: configuration of the devices.

AA batteries. We chose this particular setting after extensive measurements using other configurations that did not provide the required accuracy (e.g., connecting the PCA PA-6000 to the electrical wall plug provides less accurate results). The figure also illustrates one additional device operating as Access Point, which is a standard laptop with a built-in Atheros PCMCIA card running Ubuntu 9.10 with the `ath5k` driver, running the `hostapd` software. We provide in Fig. 7 a picture of our testbed, which is deployed in the basement of our building, where no other WLAN traffic could be detected.

5.2. Experimental evaluation

We configure the AP and the two stations with a fixed rate of 36 Mbps. In order to confirm our previous results, we run extensive tests with different configurations of the CW_{min} of the Alix and the Soekris device, denoted as CW_{AL} and CW_{SO} , respectively, both sending UDP packets of 1470 bytes using `iperf`¹¹ towards the AP at the maximum achievable rate (i.e., saturation conditions). For each configuration we measure the individual throughput and power consumption during 60 seconds and provide the average of 5 repetitions. As we have only one power meter, for each configuration we individually measure the power consumed by each device (needed to compute the EF value), and then measure the power consumed by the two devices together, in order to have a good accuracy in the measurements. The repeatability of experiments, which is required to follow this methodology, is possible thanks to the lack of WLAN activity in our basement deployment (e.g., for a given

¹¹<http://iperf.sourceforge.net/>

Figure 7: Picture of the deployed testbed.

configuration, the variations in terms of frames transmitted, received and collided among the 5 repetitions were smaller than 0.1%).

In addition to the default DCF configuration for 802.11g, namely $CW_{min} = 16$, $CW_{max} = 1024$, we performed a sweep in the contention parameters by setting $CW_{min} = CW_{max}$ with $CW_{min} \in \{2, 4, 8, \dots, 1024\}$ (note that the device only support powers of two as CW values). We present in Table 5 only a representative set of the obtained results for the sake of space. For each row representing a different $\{CW_{AL}, CW_{SO}\}$ configuration, we provide the throughput obtained by each station (R_{AL} and R_{SO}), the total throughput R , the corresponding value of Jain’s fairness index (JFI), the total power consumed P by the devices and the resulting overall energy-efficiency η . We also highlight in bold font the configuration that maximizes performance according to the EF criterion. The results can be summarized as follows:

- The default DCF configuration provides a fair bandwidth allocation ($JFI = 0.99$), but its energy efficiency performance (2.06 MbpJ) is not among the best values. Still, it should be noted that the default DCF configuration outperforms other fair configurations with either overly

small contention window values ($\{CW_{AL}, CW_{SO}\} = \{2, 2\}$) or overly large contention window values $\{CW_{AL}, CW_{SO}\} = \{1024, 1024\}$.

- As we described in Section 2, energy efficiency is maximized by giving all resources to one station and choking the other. Our results confirm that the $\{CW_{AL}, CW_{SO}\} = \{2, 1024\}$ and $\{CW_{AL}, CW_{SO}\} = \{1024, 2\}$ configurations provide the best values with respect to this metric (2.58 and 2.31 MbpJ, respectively), with an improvement of up to 25% as compared to DCF performance.
- Furthermore, out of these two extreme cases, the most energy efficient configuration is the one where all resources are given to the device that consumes the most energy in the idle mode, i.e., the Alix device (see Table 2). In this way, the most energy efficient device in idle mode (the Soekris box) suffers from starvation.
- Finally, the configuration that maximizes the EF criterion, in bold, provides a good trade-off between the DCF configuration and the most energy efficient configuration. Indeed, throughput distribution is kept fair, while the overall energy efficiency is improved by approximately 5%. Note that for larger testbeds, DCF performance will worsen due to the increased collision rate, and therefore this improvement in terms of energy efficiency is expected to grow further.

Table 5: Throughput performance (in Mbps), power consumption (in Watts) and energy efficiency (in MbpJ) of our deployment for different configurations of the CW .

CW_{AL}	CW_{SO}	R_{AL}	R_{SO}	R	JFI	P	η
2	2	9.85	8.56	18.41	0.99	9.86	1.86
8	8	10.50	10.80	21.30	0.99	9.86	2.16
DCF	DCF	10.30	10.10	20.40	0.99	9.87	2.06
1024	1024	1.59	1.58	3.17	0.99	8.02	0.39
2	1024	21.00	0.12	21.12	0.50	8.19	2.58
8	1024	19.50	0.20	19.70	0.51	8.56	2.30
2	8	15.50	5.04	20.54	0.79	8.97	2.28
8	2	5.93	14.30	20.23	0.85	9.85	2.05
1024	8	0.35	19.90	20.25	0.52	9.54	2.12
1024	2	0.08	22.00	22.08	0.50	9.53	2.31

These experimental results confirm the main conclusions derived from our analysis and simulations, namely, that maximizing the energy-efficiency of WLAN without any other consideration could result in extreme unfairness, and that our EF criterion provides a good trade-off between throughput performance and energy efficiency.

6. Conclusions

Energy-efficient operation of mobile devices has been recognized as a key challenge for the design of future communication systems, which comprises the optimization of the energy consumption of wireless communications. Yet the performance of the dominant wireless local area network standard IEEE 802.11 is optimized with respect to throughput fairness only, thus neglecting the aspect of energy fairness. Switching from the “information per unit of time” metric to the “information per unit of energy” metric to facilitate energy efficiency is relatively straightforward for the case of all wireless devices being homogeneous in their power consumption behavior, as there is a well-defined (energy) performance figure to optimize. However, nowadays WLANs are populated with highly diverse devices with respect to power consumption, and therefore the proper definition of the figure of merit to optimize is challenging. Indeed, via both simulations and experimentation we have shown that the optimization of overall energy efficiency, oblivious to any other consideration, derives in extremely unfair resource allocations.

In order to circumvent this, we have proposed the energy-efficiency proportionally fair (EF) criterion to achieve a tradeoff between energy efficiency and throughput fairness. For the case of 802.11 WLANs, we have analytically derived the closed-form expression of the configuration that optimizes performance according to the EF criterion. The proposed configuration has been validated through extensive simulations, and has been shown to perform very similarly to the maximum achievable values derived from exhaustive searches on the configuration space. We have also deployed a small-sized testbed consisting on three 802.11g nodes, and confirmed the main results of our analysis: *(i)* the unrestrained optimization of the overall energy efficiency in heterogeneous WLANs leads to extreme throughput unfairness, which in addition favors the less efficient interfaces; *(ii)* our EF criterion finds a good trade-off between this unfair configuration and a throughput-fair but inefficient configuration.

Acknowledgments

This work has been supported by the European Community's Seventh Framework Programme (FP7-ICT-2009-5) under grant agreement n. 257263 (FLAVIA project). It has been also supported by the Spanish Government, MICINN, under research grant TIN2010-20136-C03. The authors wish to thank the anonymous referees for their helpful and constructive comments.

References

- [1] A. Garcia-Saavedra, P. Serrano, A. Banchs, M. Hollick, Energy-efficient fair channel access for IEEE 802.11 WLANs, in: World of Wireless, Mobile and Multimedia Networks (WoWMoM), 2011 IEEE International Symposium on a, 2011.
- [2] A. Banchs, X. Perez, Providing throughput guarantees in IEEE 802.11 Wireless LAN, in: Wireless Communications and Networking Conference, 2002. WCNC2002. 2002 IEEE, Vol. 1, 2002, pp. 130–138 vol.1.
- [3] A. Banchs, User fair queuing: fair allocation of bandwidth for users, in: INFOCOM 2002. Twenty-First Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE, Vol. 3, 2002, pp. 1668–1677 vol.3.
- [4] Y. Chen, S. Zhang, S. Xu, G. Li, Fundamental trade-offs on green wireless networks, *Communications Magazine*, IEEE 49 (6) (2011) 30–37.
- [5] R. Bruno, M. Conti, E. Gregori, Optimization of efficiency and energy consumption in p-persistent CSMA-based wireless LANs, *IEEE TMC* 1 (1) (2002) 10–31.
- [6] P. Serrano, A. Garcia-Saavedra, M. Hollick, A. Banchs, On the Energy Efficiency of IEEE 802.11 WLANs, in: European Wireless 2010, 2010.
- [7] M. Ergen, P. Varaiya, Decomposition of Energy Consumption in IEEE 802.11, in: IEEE ICC 2007, 2007, pp. 403–408.
- [8] V. Baiamonte, C.-F. Chiasserini, Saving energy during channel contention in 802.11 WLANs, *Mob. Netw. Appl.* 11 (2) (2006) 287–296.

- [9] S. Chiaravalloti, F. Idzikowski, L. Budzisz, Power consumption of WLAN network elements, TKN Technical Report Series TKN-11-002, Telecommunication Networks Group, Technical University Berlin (Aug. 2011).
- [10] P. Serrano, A. Banchs, P. Patras, A. Azcorra, Optimal Configuration of 802.11e EDCA for Real-Time and Data Traffic, *IEEE Transactions on Vehicular Technology* 59 (5) (2010) 2511–2528.
- [11] H. Wang, N. Agoulmine, M. Ma, Y. Jin, Network lifetime optimization in wireless sensor networks, *Selected Areas in Communications, IEEE Journal on* 28 (7) (2010) 1127–1137.
- [12] F. Kelly, A. Maulloo, D. Tan, Rate control in communication networks: shadow prices, proportional fairness and stability, in: *Journal of the Operational Research Society*, Vol. 49, 1998.
- [13] L. Massoulié, J. Roberts, Bandwidth sharing and admission control for elastic traffic, *Telecommunication Systems* 15 (2000) 185–201.
- [14] T. Nandagopal, T.-E. Kim, X. Gao, V. Bharghavan, Achieving MAC layer fairness in wireless packet networks, *MobiCom '00*, ACM, New York, USA, 2000, pp. 87–98.
- [15] A. Banchs, P. Serrano, H. Oliver, Proportional fair throughput allocation in multirate IEEE 802.11e Wireless LANs, *Wireless Networks* 13 (5).
- [16] R. Jain, W. Hawe, D. Chiu, A quantitative measure of fairness and discrimination for resource allocation in shared computer systems, in: *DEC-TR-301*, 1984.
- [17] G. Bianchi, Performance analysis of the IEEE 802.11 distributed coordination function, *IEEE JSAC* 18 (3) (2000) 535–547.
- [18] A. Banchs, P. Serrano, Revisiting 802.11e EDCA performance analysis, *Wirel. Pers. Commun.* 43 (4) (2007) 1145–1149.