UNIVERSITY CARLOS III OF MADRID

Department of Telematics Engineering

Master of Science Thesis

Energy-efficient fair channel access for IEEE 802.11 WLANs

Author: **Andrés García Saavedra**
Dipl. Eng. in Telecommunications

Supervisors: **Pablo Serrano Yáñez-Mingot, Ph.D.**
Albert Banchs Roca, Ph.D.

Leganés, July 2010

Abstract

The need to minimize the footprint of mobile communication protocols has lead to a shift in their design paradigm: the figure to optimize is no longer the number of bits transmitted per second, but instead the number of bits transmitted per joule. However, different mobile devices present very different power consumption figures (e.g., a laptop can consume ten times the power of a small access point), and therefore it is not clear how very-efficient devices should share the wireless resources with less-efficient devices. In this work we first propose a new energy-based criterion, the EF criterion, to define a trade-off between the most energy-efficient configuration (where all resources are given to a single device) and the throughput-fair allocation (where all devices evenly share the resources regardless of their power consumption). We then address the case of IEEE 802.11 wireless LANs, based on a performance analysis model for their energy consumption. We use this model to derive a closed-form expression for the configuration that optimizes performance according to the EF criterion, and validate it through simulations. We also present some preliminary experimental results characterizing the power consumption of wireless devices.

Table of Contents

List of Figures

Chapter 1

Introduction

Information and communication technology (ICT) holds one of the keys to the reduction of greenhouse gases produced worldwide. The consumption of energy in the ICT can be significantly reduced by increasing the energy efficiency of computing as well as networking. The importance of "greening the Internet" is thus recognized as a primary design goal of future global network infrastructures. Indeed, it is estimated that, today, the Internet already accounts for about 2% of total world energy consumption[1], and with the current trend of shifting offline services online, this percentage is expected to grow significantly in the next years. The energy consumption is to be further fueled by the forthcoming Internet-based platforms that require always-on connectivity.

Communication protocols, and in particular the technologies used in the access network, have been originally conceived to optimize metrics other than energy, such as throughput or delay. *Greening* these protocols thus represents a shift in the design paradigm, where energy instead of time is the most critical network resource. We no longer want to maximize the bits sent per time unit, but instead the bits the network can send per each joule consumed. Still, it is intuitively clear that this will not come for free, and there might be a price to pay in terms of throughput performance when developing sustainable and energy efficient architectures.

Indeed, these two performance parameters, throughput and energy efficiency, may constitute different objectives. In order to illustrate this, let us consider consider a simple WLAN scenario consisting on one Access Point (AP) and two associated stations (sta_1 and sta_2, respectively). We will assume that the only source of losses is frame collisions (i.e., ideal channel conditions), and the use of the IEEE 802.11b physical layer. In these circumstances, the minimum Contention Window (CW_{min}) that maximizes a fair throughput allocation can be obtained from, e.g., a numerical search on its space. This value results $CW_{min} = 17$, and obviously does not depend on the power consumption characteristics of the wireless LAN (WLAN) interface. However, if we want the CW_{min} configuration not to maximize the throughput but the energy efficiency (we well make a precise definition of this performance parameter in the next chapter) then, for the case of, e.g., SocketCom CF interfaces [1] the resulting configuration would be $CW_{min} = 56$, i.e., a value three times larger.

However, despite the ongoing concerns about the energy consumption of network devices, this relation between throughput maximization and energy-efficiency optimization

[1]As reported in "SMART 2020: Enabling the Low Carbon Economy in the Information Age", The Climate Group, available at `http://www.smart2020.org/`

has received relatively little attention. To the best of our knowledge, there have been two main contributions: on one hand, Bruno et al. [4] considered the case of p-persistent CSMA-based WLANs and proved that, based on a naïve energy consumption model, they could be jointly optimized; on the other hand, our previous work of [9], based on a more sophisticated energy consumption model, showed that they may constitute different different optimization objectives, this resulting in different configurations of the CW_{min} parameter as seen above for the simple case of a WLAN with two stations.

One key limitation of these previous approaches is that they only consider *homogeneous* scenarios, where all WLAN devices share the same power consumption characteristics. This actually constitute a non-realistic scenario as, indeed, WLAN devices show very different power consumption figures, as illustrated in Table 2.1 for just three different interfaces. We argue that any configuration that aims at optimizing the energy efficiency of a wireless network needs to take into account the diversity of the power consumption interfaces. This is the challenge that we tackle in this work: the configuration of *heterogeneous* scenarios, where different WLAN stations have different WLAN power consumption figures. We will start by first addressing the following key question: in case a station consumes twice as much power as another station, should they get the same bandwidth? Should there be any difference? Once tackled this question, we will consider the case of 802.11 based WLANs, and provide a closed-form expression for their configuration.

The rest of the work is organized as follows. In Chapter 2 we first illustrate the challenge of configuring heterogeneous WLAN scenarios, and then we will propose a criterion to address it. In Chapter 3 we present the energy consumption model that we use to predict the consumption in the WLAN and its validation, while in Chapter 4 we use this model to derive a closed-form expression to achieve the optimal configuration. We validate this configuration in Chapter 5 through extensive simulations and numerical searches, and we describe some experimental results in Chapter 6. Finally in Chapter 7 we conclude the work and present ongoing research activities.

Chapter 2

An energy efficiency-based criterion

The work in [9] defined an optimization criterion for homogeneous WLANs, that is, a network of wireless stations with identical energy properties. It provided the configuration (i.e., CW_{min}) that maximizes the overall energy efficiency of the system, thus obtaining the optimal point of operation: highest ratio between the throughput obtained and the power consumed (i.e, the energy efficiency η).

$$\eta = \frac{throughput}{power}$$

In that work we proved that, for different interfaces (with different power consumption figures), the configuration to be used by the stations varied. We have already seen before that for a simple scenario, the minimum contention window (CW_{min}) that optimizes the overall throughput performance is $CW_{min} = 17$, while the value that optimizes the number of bits sent per joule consumed is $CW_{min} = 56$. Nevertheless the work of [9] is based on a non-realistic assumption: *Homogeneity* among stations. We will show in the following why a heterogeneous scenario constitutes a more challenging scenario that requires the definition of a carefully-designed optimization criteria.

2.1 The need for a fairness criterion

In order to illustrate the *risks* of using a configuration computed to provide maximum overall efficiency without further considerations, we will run a simple experiment with two stations and one Access Point using 802.11b.

Table 2.1: Power consumption parameters while transmitting (ρ_t), receiving (ρ_r) and idling (ρ_i) as reported in [1] (in watts)

#	Card	ρ_t	ρ_r	ρ_i
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

One station will be modeled after a high-consuming interface, A (from Table 2.1), whereas the other station will use an interface, B, with lower power consumption requirements. We will run a series of simulations:

- In the first simulations, we set $CW_1 = CW_2$, in order to have a fair share of the wireless resources, and perform a sweep on the contention windows space, $CW = \{8, 1024\}$, to choose the value that maximizes the global throughput performance.
- In the second series of simulations, we let CW_1 and CW_2 be different, and we perform another sweep to find the configuration that maximizes the overall energy efficiency, η, of the WLAN.

For the first experiment (named "Throughput"), we obtain $CW_1 = CW_2 = 17$; whereas for the second experiment (named "Efficiency") we obtain $CW_1 = 8$, $CW_2 = 1024$. Figure 2.1 shows the obtained values for the per-station throughput and overall energy efficiency for both simulations. The "Throughput" configuration, as expected, offers a fair share of bandwidth for both stations ($3.76 Mbps$) but a lower energy efficiency value than the other approach ($3.48 bpJ$). On the other hand, the "Efficiency" configuration, despite giving higher overall energy efficiency (3.75 bpJ, approximately a 10% improvement), it is extremely unfair and practically *chokes* one of the stations (station B).

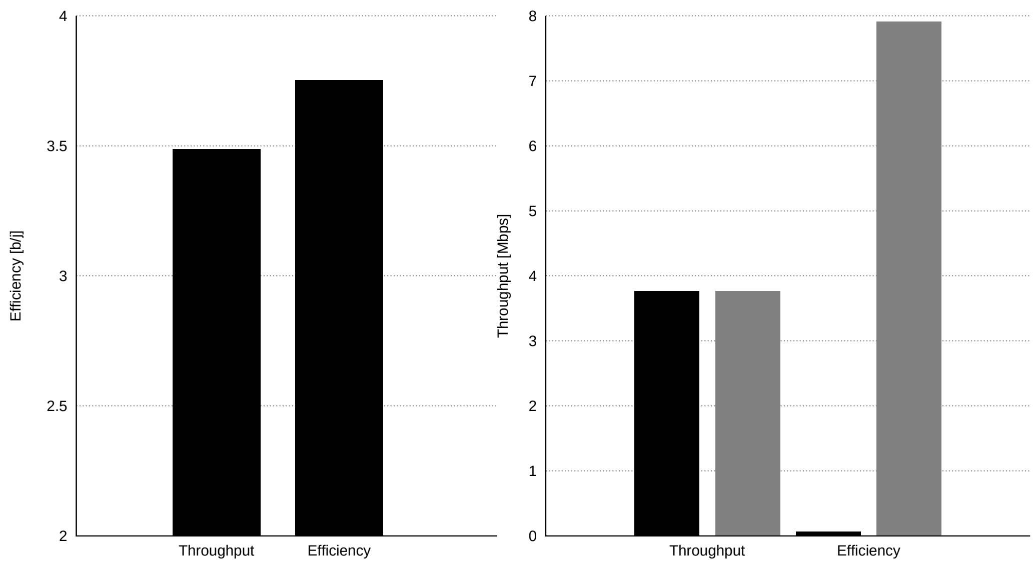

Figure 2.1: Throughput and Energy-Efficiency performance of a WLAN with two stations modeled after interfaces A (grey) and B (black) from Table 2.1.

The explanation to this behavior is as follows. If we can configure different parameters to two groups of stations without fairness considerations, choking one of them will minimize the number of collisions; which is not so obvious is which group of station(s) should suffer from starvation. At a simple glance we might, erroneously, give a larger share of bandwidth to the more efficient stations. If we have a closer look at Figure 2.1, we will realize that the best performance (in terms of overall energy efficiency) is achieved when penalizing the most efficient station! The reason is that a station cannot be deactivated and therefore, such interface has to unavoidably consume a minimum portion of energy by idling. The configuration algorithm will rather give all the share to that interface whose difference between the power consumed when idling and the power consumed when transmitted is smaller (these values being higher than the ones for the other group).

This simple scenario helps us to understand the challenges behind configuring a heterogeneous wireless network: Using a *naïve* mechanism that simply computes the maximum overall efficiency, we may penalize the *greener* interfaces! Besides, if we do not consider the power properties of the different interfaces we will obtain a throughput-optimized configuration that may result in energy wastage. Therefore, **a trade-off between overall energy-efficiency and throughput fairness is needed**.

2.2 The energy-efficiency fair criterion

The use of *overall* energy efficiency figures, as we have seen in the previous section, is not well suited to properly address realistic (i.e., heterogeneous) scenarios, as it may result in configurations with extreme unfairness across stations. The use of throughput-based approaches, on the other hand, do not consider the impact of the different power consumption parameters and therefore may result in energy wastage. We argue that a tradeoff between these two approaches is needed.

In order to define a trade-off between these two different optimization objectives, we first define the *per-station* energy efficiency η_i as the ratio between the throughput and the power consumed by a given station i, i.e.,

$$\eta_i = \frac{throughput_i}{power_i}$$

Note that η_i provides the throughput the station i is successfully transmitting over the energy the station has to spend, and therefore can also be used to measure situations of extremely unfairness across station, e.g., in the previous scenario, the resulting values for the "Througput" configuration is $\eta = \{5.54, 2.54\}\ bpJ$, while for the case of the "Energy" configuration is $\eta = \{5.02, 0.11\}\ bpJ$.

Based on these η_i variables, the challenge remains on defining an appropriate criterion for their configuration. Note that, had we had a throughput allocation problem with r_i being the throughput station i receives, we could have used, e.g., Kelly's *proportional fairness* (PF) criterion [8] to define the proper trade-off to configure the throughput allocation vector, this being the one that maximizes the sum of the rates' logarithms, i.e.,

$$PF \Longleftrightarrow \max \sum \log r_i$$

Based on this well-adopted throughput allocation criterion, in this work we advocate for the use of the energy-efficiency proportional fairness criterion (hereafter the EF criterion), based on the maximization of the sum of the per-station energy efficiency, i.e.,

$$EF \Longleftrightarrow \max \sum \log \eta_i \tag{2.1}$$

In order to illustrate why the use of the EF criterion prevents extremely unfair allocations while supporting energy-efficient configurations, let us consider the same scenario we used in Section 2.1; that is, one Access Point and two stations, one modeled after interface A (station 1) while the other is modeled after interface B (station 2). We will run different simulations for different configurations while computing the values of EF[1] performance (as defined in 2.1), throughput, and overall energy efficiency. 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 η.

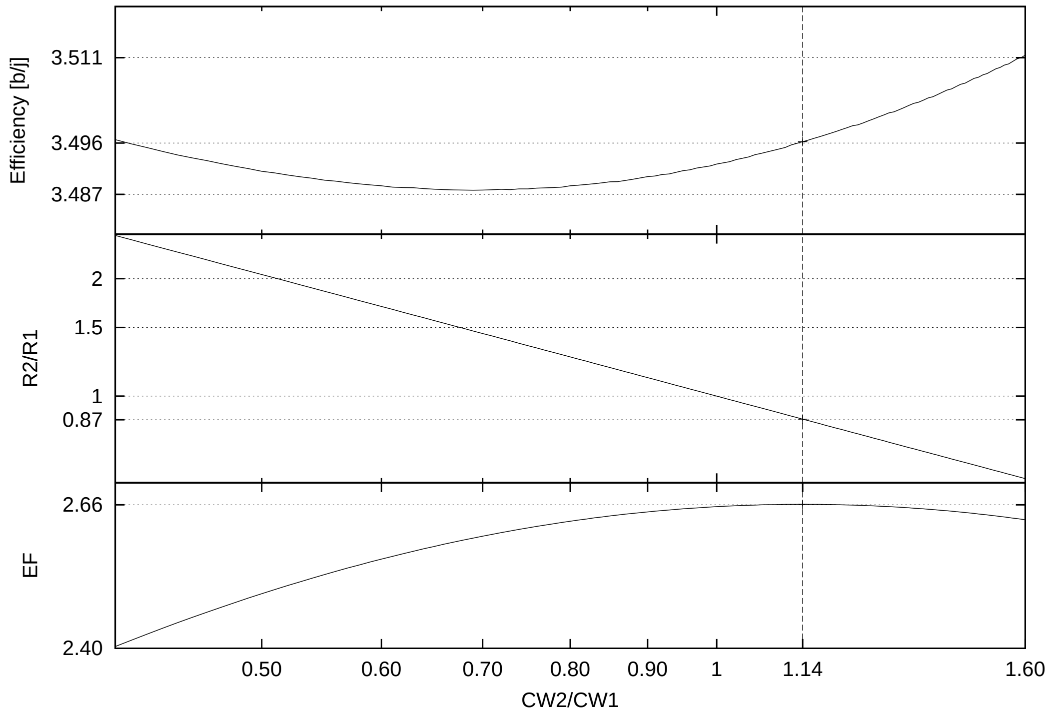

Figure 2.2: Total efficiency, throughput ratio and EF performance of a WLAN with two stations for different CW configurations.

We present the results in Figure 2.2 that can be summarized as follows:

- The highest values of overall efficiency (as we already saw in the previous section) are achieved for large values of the ratio CW_2/CW_2. That implies starvation of the least-consuming station, as it can be seen from the R_2/R_1 ratio.
- The EF value is not maximized for that extremely unfair configuration, but instead it is achieved for $k \approx 1.14$. From this point on, the increase in $\eta_1 (\eta_2)$ is not compensated

[1]Note that, for the sake of readability, throughout the work we will 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.

by the decrease of $\eta_2(\eta_1)$ and, hence, the allocation is not EF-optimal.

For this set of simulations, the efficiency-optimal WLAN is achieved with a configuration $CW = \{3, 384\}$, this providing an overall efficiency $\eta = 3.82 bpJ$ and a throughput allocation of $R = \{8.23, 0.06\}$ Mbps, i.e., extreme unfairness. On the other hand, the EF-optimal configuration is given by $CW = \{26, 30\}$ offering an overall efficiency of $\eta = 3.49 bpJ$ and throughput of $R = \{3.97, 3.47\}$ Mbps. After these numbers, we can see that the EF-optimal configuration gives us an 8.6% lower overall efficiency in exchange for achieving an increase of 95% in throughput fairness (from 0.51 to 0.995)[2]. Note that the EF-optimal configuration of the CW is not only different from the maximum throughput allocation (that would be $CW = \{17, 17\}$), but also from the one in case both stations were modeled after the same WLAN interface ($CW = 56$ for the case of interface B, $CW = 19$ for the case of interface A).

These results show how the EF-criterion provides a proper trade-off between fair bandwidth allocation and energy efficiency. Note that, despite the rest of this work is devoted to 802.11 WLANs, the proposed criterion EF could also be applied for any other technology as an objective to optimize when looking for a fair and energy-efficient configuration.

[2]These being computed using Jain's fairness index [7].

Chapter 3

Energy consumption model for 802.11 WLANs

In the previous section we have illustrated why heterogeneous WLANs constitute a challenging scenario, and we have proposed a criterion to achieve a trade-off between throughput fairness and energy efficiency. In order to apply this criterion to derive the EF-optimal configuration for heterogeneous WLANs, in this section we will present the model to characterize the energy consumption in a WLAN. First we will introduce an accurate but complex model, and then we will present a simpler model that sacrifices accuracy for analytical tractability. Finally, we will validate their accuracy using simulations.

3.1 Energy model

The model considers an 802.11 WLAN scenario, in the assumption of ideal channel conditions (i.e., no hidden terminals and capture effect). In this scenario N stations, sharing the wireless channel, operate in saturated conditions; that is, each station has always a packet available to transmit.

We follow Bianchi's seminal work [3], where a station i with minimum contention window CW^i_{min} has a probability τ_i to attempt transmission upon a backoff counter decrement (i.e, a timeslot). In turn, p_i is the probability that a transmission attempt by station i collides. The relation between both probabilities is given by the well-known equations:

$$
\begin{aligned}
\tau_i &= \frac{2}{1 + CW^i_{min} + p_i CW^i_{min} \sum_{j=0}^{m-1} (2p_i)^j} \\
p_i &= 1 - \prod_{j \neq i} (1 - \tau_j)
\end{aligned}
$$

The above constitutes a system of non-linear equations that can be solved numerically (see [2]), giving the values for the τ_i's. Note that for the case of $CW^i_{min} = CW^i_{max} = CW^i$ the computation of the transmission probability is simplified as

$$
\tau_i = \frac{2}{CW^i + 1}
$$

In order to model the energy consumption of a WLAN we follow a similar approach to the one of [5] which defines three parameters to model the power consumption information of an 802.11 interface:

- $\rho_i^{tx} \equiv$ Power consumption of interface i while transmitting.
- $\rho_i^{rx} \equiv$ Power consumption of interface i while receiving.
- $\rho_i^{id} \equiv$ Power consumption of interface i while neither transmitting nor receiving, but idling.

We will assume that stations only transmit data to the Access Point, and that all frames have a fixed length L. This way, the energy consumed by station i in a timeslot is given by:

$$e_i = \sum_{j \in \Theta} E_i(j) p(j) \tag{3.1}$$

where,

$\Theta \equiv$ Set of events that can take place within one timeslot.
$E_i(j) \equiv$ Energy consumed by station i in case of event j.
$p(j) \equiv$ Probability that event j occurs.

The set of events and their probability are:

- 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, $p(c, \neg i)$

while these probabilities can be computed as:

$$\begin{aligned}
p(e) &= \prod (1 - \tau_j) \\
p(s, i) &= \tau_i \prod_{j \neq i} (1 - \tau_j) \\
p(s, \neg i) &= \sum_{j \neq i} \tau_j \prod_{k \neq j} (1 - \tau_k) \\
p(c, i) &= \tau_i (1 - \prod_{j \neq i} (1 - \tau_j)) \\
p(c, \neg i) &= 1 - \tau_i - p_e - p_{s, \neg i}
\end{aligned} \tag{3.2}$$

Table 3.1: Power consumed (in mJ) per event for the interfaces of of Table 2.1

#	$E(e)$	$E(s,i)$	$E(s,\neg i)$	$E(c,i)$	$E(c,\neg i)$	α_i	β_i
A	0.0230	2.2834	1.9801	2.2454	1.9421	0.9884	0.1532
B	0.0013	1.2151	0.8148	1.1349	0.7346	0.9984	0.4913
C	0.0016	1.8930	1.1651	1.7759	1.0481	0.9986	0.6247

This way, we can easily compute the energy consumed by a station i for all the possible events.

$$\begin{aligned} &E_i(e) = \rho_i^{id} T_e \\ &E_i(s,i) = \rho_i^{tx} T_s + \rho_i^{rx} T_{ack} + \rho_i^{id}(SIFS + DIFS) \\ &E_i(s,\neg i) = \rho_i^{rx}\left(T_s + T_{ack}\right) + \rho_i^{id}(SIFS + DIFS) \\ &E_i(c,i) = \rho_i^{tx} T_s + \rho_i^{id} EIFS \\ &E_i(c,\neg i) = \rho_i^{rx} T_s + \rho_i^{id} EIFS \end{aligned} \tag{3.3}$$

where T_e is the duration of an empty slot time, $SIFS$, $DIFS$ and $EIFS$ are constants defined by the 802.11 standard [6], 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

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

where T_{PLCP} is the length of the frame preamble, H is the frame header, C the modulation rate being used, and ACK represents the length of an acknowledgement frame. We can expand (3.1) using the information in (3.3) and (3.2) to obtain the following expression:

$$\begin{aligned} e_i \;=\; & E_i(e)\cdot p(e) + \\ + \; & E_i(s,i)\cdot p(s,i) + \\ + \; & E_i(s,\neg i)\cdot p(s,\neg i) + \\ + \; & E_i(c,i)\cdot p(c,i) + \\ + \; & E_i(c,\neg i)\cdot p(c,\neg i) \end{aligned} \tag{3.4}$$

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 as the ratio between the bits successfully transmitted over the energy consumed in a slot time:

$$\eta_i = \frac{p(s,i)L}{e_i} \tag{3.5}$$

As it can be seen from (3.1), the full expression for e_i consists of the sum of several terms that non-linearly depend on the τ's. In order to improve the analytical tractability, we quantify the energy consumed per timeslot for the different interfaces in Table 3.1 to make the following approximations:

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

With the above, we make the following approximation for (3.1)

$$\hat{e}_i = p_e E_i(e) + \tau_i E_i(s,i) + (1 - p_e - \tau_i) E_i(s,\neg i) \tag{3.6}$$

Note that the use of (3.6) will result in an overestimation of the power consumed, as for the two terms being approximated we always take the largest. We will further rearrange (3.6) 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.,

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

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

Note that we will denote with η the energy efficiency as computed with the use of (3.1) and with $\hat{\eta}$ the efficiency computed using the approximate expression (3.6). In the next section we will assess the accuracy of both expressions to model the energy consumption and efficiency in a heterogeneous WLAN.

3.2 Validation

In this section we assess the accuracy of both the analytical and the approximate model. To do this we will perform three experiments: We will first show how the energy efficiency varies depending on the number of nodes within a heterogeneous scenario, we will validate the per-station efficiency for different 802.11 configurations, and finally we will numerically search the optimal EF configuration for using the approximate model and simulations.

3.2.1 Energy efficiency of a heterogeneous WLAN

Let us consider a scenario with N stations using the standard DCF configuration, where one third of the stations are modeled after interface A, another third after interface B, and the rest after interface C. We then compute the overall energy efficiency using the analytical model (3.5) ("Model"), using the approximate expression $\hat{e}_i$ (3.6) ("Approx."), and compare them against results from simulations ("Simulation"). These are plot in Fig. 3.1.

The figure shows that both models are able to predict WLAN energy behavior, as analytical results closely follows 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

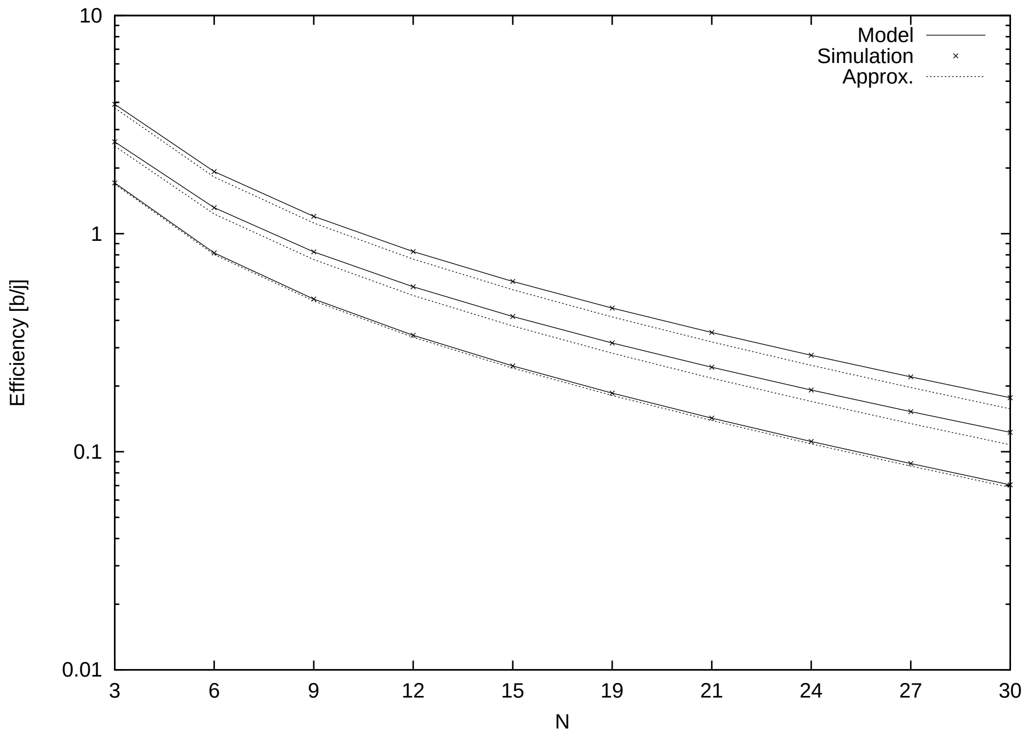

Figure 3.1: Overall energy efficiency η of a heterogeneous WLAN with N stations.

by the increase in the number of collisions for the static DCF configuration, and that the approximate model overestimates the energy consumed in a timeslot, this way underestimating the overall efficiency.

3.2.2 Per-station efficiency for a WLAN

In order to validate the per-station efficiency η_i we consider $N = 30$ stations in the WLAN, modeling one third of the total number of stations after the ρ parameters of Interfaces A, B and C of Table 2.1, respectively. Then we set $CW_1 = CW_2 = CW_3$, and perform a sweep on the CW, computing the resulting per-station energy efficiency η as given by simulations, the analytical model and its approximation, with the results being represented in Fig. 3.2. We have again that the analytical model closely follows the results from simulations, while the approximate model slightly deviates from them, underestimating the energy efficiency. However, note that this difference between the simulation results and those derived from the approximated model are very small and, most importantly, that the maximum for the simulation results and the approximate model (which we denote with a circle in Fig. 3.2) are located around similar CW values. This ability of the approximate model to capture the behavior of η_i supports the derivation of the EF-optimal configuration, that we will address in the next section.

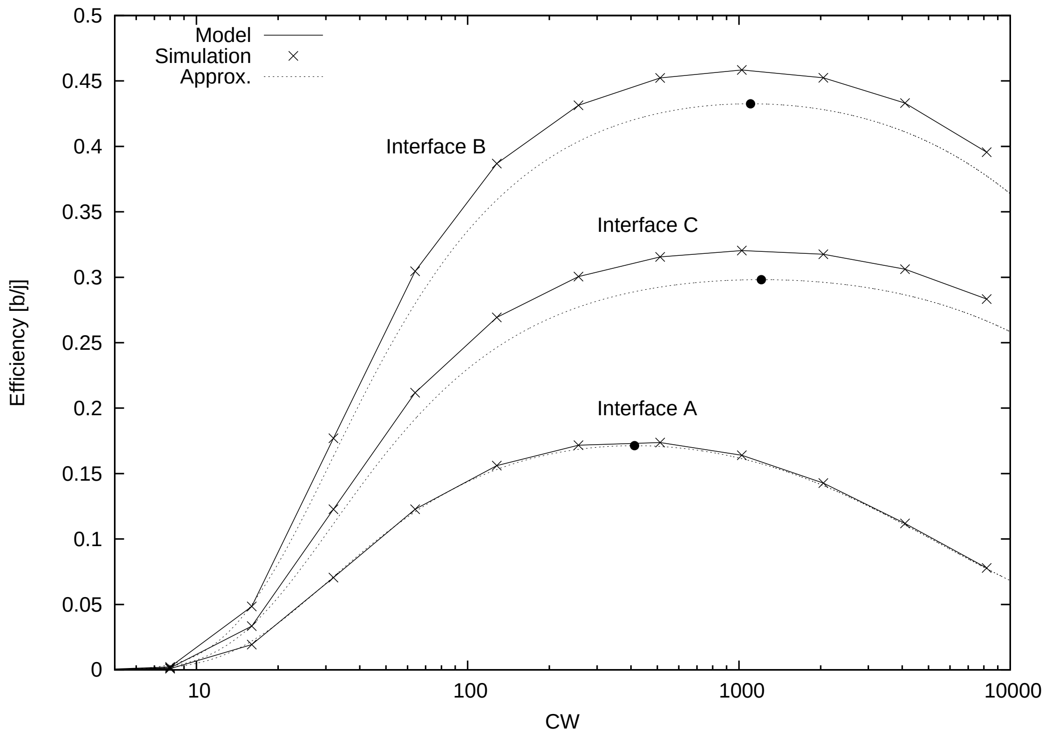

Figure 3.2: Per-station efficiency η_i for a WLAN with $N = 30$ stations.

3.2.3 EF-optimal configuration for 802.11

We have seen that the approximate model overestimate the energy consumption, but is able to follow the behavior of the actual η_i values. To further validate this, we will consider the same scenario as in the previous section. We will then perform an exhaustive search for the best set of the 802.11's CW_{min} parameters that gives us the highest EF value using $i)$ the approximate model, and $ii)$ the analytical model. The results, seen in figure 3.3, demonstrate that the optimum optimal configuration resulting from simulations is very similar to the one obtained using the approximate model, which confirms its validity to derive the EF configuration.

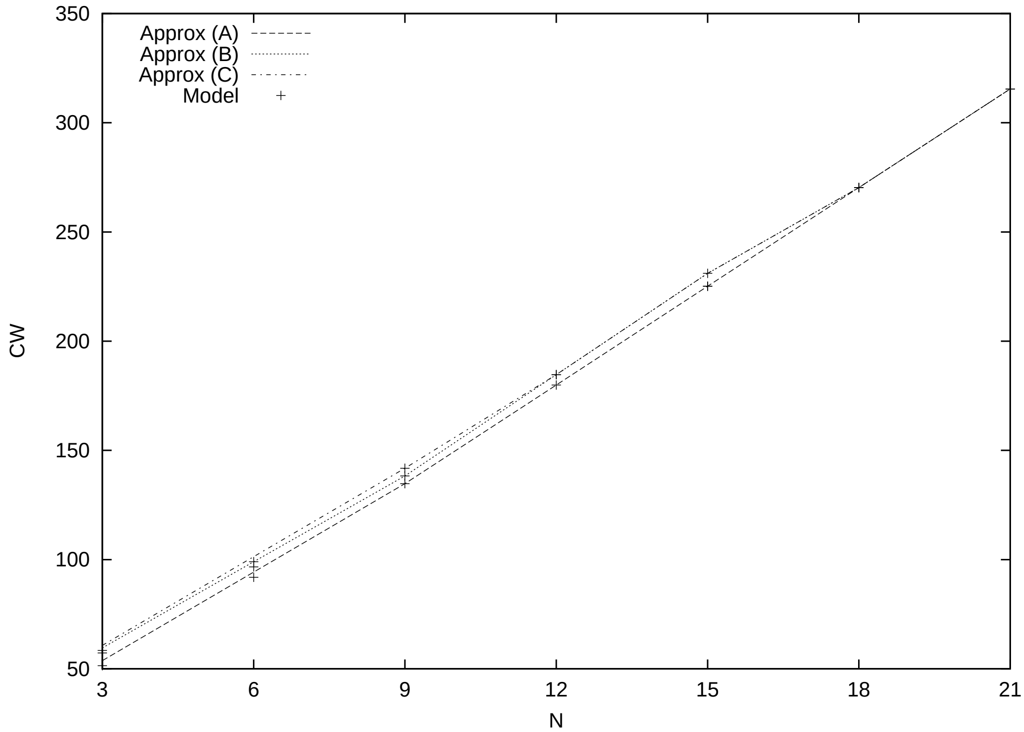

Figure 3.3: EF-optimal configuration for 802.11.

Chapter 4

EF configuration for 802.11 WLANs

Based on the energy consumption model presented in the previous chapter, we have the following expression for the energy efficiency $\hat{\eta}_i$:

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

Computing the EF-optimal configuration requires to find the τ's that maximize 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 some terms results in the following

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

that can be approximated as

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

This gives us an important first result: the τ_k's that provides the EF-optimal configuration does not depend on k, but it is the same for all stations[1]. Therefore, in order to achieve an EF-optimal configuration, stations have to fairly share the channel and thus,

$$\tau_i \approx \tau_k \ \ \forall i,k$$

[1]Note that we already saw this in previous examples.

We have then proven, that all stations have to use the same configuration. The remaining challenge is to derive this configuration. In order to tackle this, we will reformulate the criterion to take advantage of the logarithm's properties. Hence,

$$\max \sum_i \log \eta_i \Longleftrightarrow \max \prod_i \eta_i$$

Under the assumption $\tau \ll 1$ and $\beta_i < 1$, 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

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

the EF-optimal configuration can be computed by maximizing

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

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

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

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

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

$$\tau^* \approx \frac{1}{N}\sqrt{2\left(\frac{N}{\sum \alpha_i} - 1\right)} \approx \frac{1}{N}\sqrt{2\frac{T_e}{T_s}\left(\frac{1}{N}\sum \frac{\rho_i^{id}}{\rho_i^{rx}}\right)} \qquad (4.1)$$

Therefore, assuming we know the ρ parameters of all N stations in the WLAN could compute the CW that provides the optimal energy-fair configuration as follows:

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

This way, (4.1) constitutes, to the best of our knowledge, the most general expression to compute the configuration of a WLAN with N stations. In fact, it also covers the case of homogeneity that we addressed in [9] where

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

One of the major disadvantages of the use of 4.1 is that it requires obtaining the $\{\rho^{id}, \rho^{rx}\}$ parameters of all WLAN stations. Indeed, this would require not only a communication protocol to convey this information, but also that all stations are aware of their power consumption values, two requirements not straightforward to fulfill. In order to tackle this inconvenience, we can make the following *coarse* approximation (see Table 2.1)

$$\sqrt{\frac{\rho^{id}}{\rho^{rx}}} \approx 1$$

which results in the following approximate expression for the optimal τ

$$\tau^* \approx \frac{1}{N}\sqrt{\frac{2T_e}{T_s}} \tag{4.2}$$

Chapter 5

Performance evaluation

In this chapter we will assess the performance of a wireless network using the configuration proposed in (4.2). Besides, we will compare the results with the performance achieved using exhaustive searches in the CW's space to find the optimum configuration, and against the performance achieved while using the approximate expression of (4.2).

5.1 Homogeneous scenario

The work performed in [9] assesses the performance evaluation in homogeneous scenarios. We will include some of the results for the sake of the completeness of the thesis. As we have seen in Chapter 4, this is actually a particular case of the general expression in (4.1). Here all the stations possess equal energy consumption properties and, therefore, the EF-optimal configuration can be computed using the following expression:

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

It is important to notice that the configuration will be different from that which gives us maximum overall throughput as they pursue different objectives. We can see this in figure 5.1, where we plot the throughput and energy efficiency achieved for the homogeneous case with N = 10, for the three interfaces of 2.1, and for different values of τ's.

Finally, in order to emphasize the trade-off between throughput-optimal and energy efficiency-optimal configurations we plot figure 5.2. Here, we display for the case of $N = 10$ the energy efficiency (in the x-axis) and the throughput performance (in the y-axis) for all possible configurations of CW for all the interfaces in Table 2.1. In the figure, we use star symbol to mark the point of maximum throughput performance, and a circle to point out the maximum energy efficiency configuration. The conclusions that come up from these results are the following:

- The interface A can be given a configuration that jointly maximizes both throughput performance and energy efficiency, that is, the point of operation that maximizes the throughput performance is the same one that maximizes energy efficiency. This is explained due to the similarity of the energy parameters values ρ^{id}, ρ^{rx}.

- For the cases of interfaces B and C, the larger the ρ^{rx}/ρ^{id} ratio, the more separate the optimum values are and, therefore, the higher the price to pay in throughput when optimizing energy (and vice-versa).

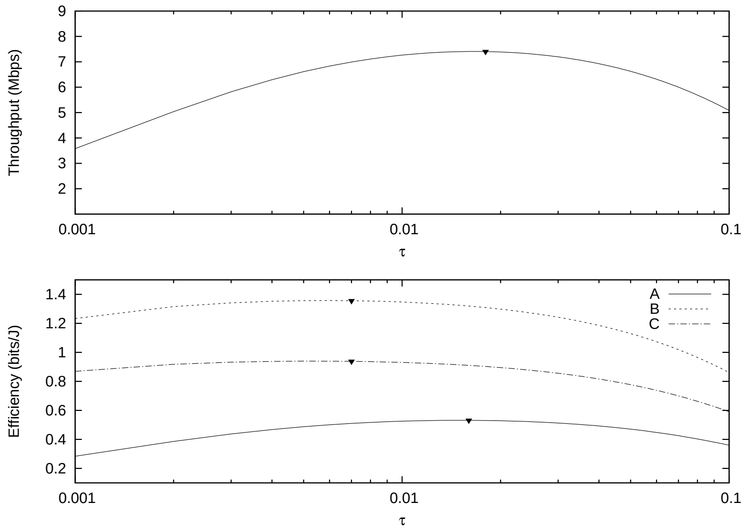

Figure 5.1: Energy efficiency and throughput vs. τ.

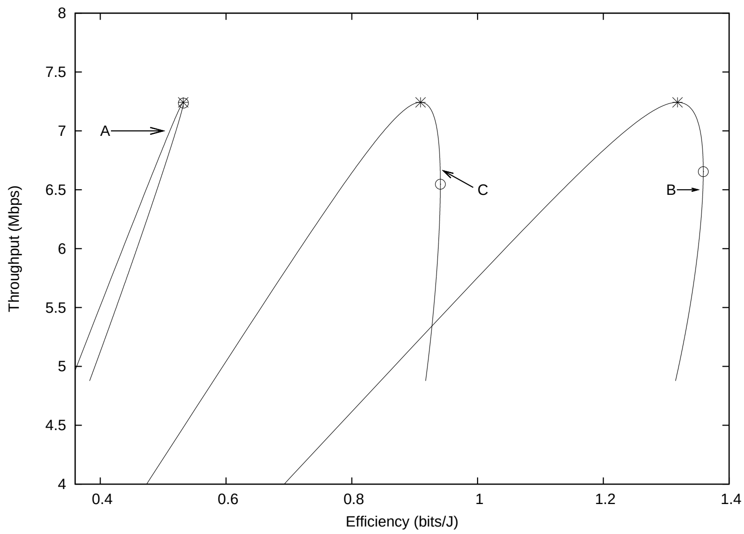

Figure 5.2: Trade-off imposed by the energy features of the WLAN interfaces

5.2 Heterogeneous scenario

In this section, we assess and compare the performance of heterogeneous wireless networks that are configured upon four different figures:

- The optimum configuration found by an exhaustive search in the CW's space and denoted by "Exhaustive". Note that this value is used only as reference, given its practical unfeasibility given the required computational time.
- The EF-optimal configuration given by 4.1 and denoted by "EF-config.".
- The throughput-optimal configuration given by 4.2 and denoted by "Approx".
- The DCF standard configuration.

In order to generate heterogeneous scenarios we consider different mixtures of interfaces. We denote with N_A , N_B and N_C the number of WLAN stations with the power properties of interfaces A, B and C from Table 2.1, respectively. Firstly, we plot in 5.3, 5.4 and 5.5 three different figures that represent three different scenarios with $N = 20$ as total number of stations:

- Figure 5.3 represents a scenario with stations that have interfaces A and B. The x-axis shows the number of stations with interface A.
- Figure 5.4 represents a scenario with stations that have interfaces A and C. The x-axis shows the number of stations with interface A.
- Figure 5.5 represents a scenario with stations that have interfaces B and C. The x-axis shows the number of stations with interface B.

The obtained EF values for the four considered configurations can lead the following results:

- The performance of the default standard configuration rapidly decreases with the number of stations, as most of the resources are wasted in energy-consuming collisions.
- Our configuration provides EF values very close to the ones achievable by means of the exhaustive search. Indeed, as results show, the differences between the "EF-config." and the "Exhaustive" lines are almost negligible, this way providing the ability of 4.1 to derive the WLAN to the EF-optimal point of operation.
- When the energy consumption information is not available, a WLAN configured according to the "Approx." approach of 4.2 provides performance values that, although smaller than the maximum achievable ones, significantly outperforms the values derived from the use of the standard configuration. Note that, when N_A is relatively large the EF values for the "Approx." and for the "EF-config." approaches are very similar, being this explained because of the similarity between its ρ^{id} and ρ^{rx} parameters.

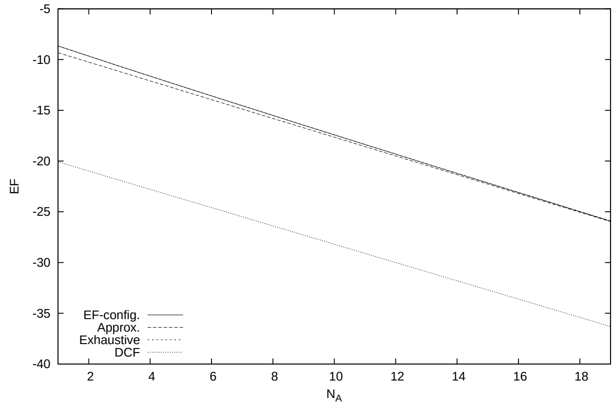

Figure 5.3: EF performance for a heterogeneous WLAN with two types of interfaces: A+B

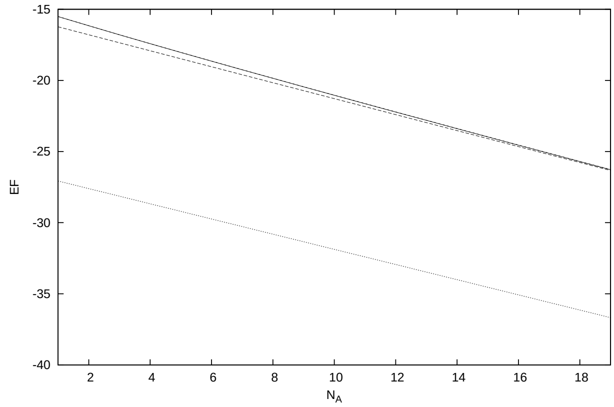

Figure 5.4: EF performance for a heterogeneous WLAN with two types of interfaces: A+C

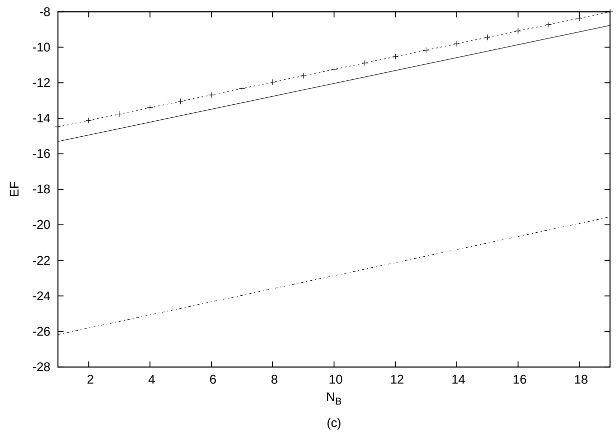

Figure 5.5: EF performance for a heterogeneous WLAN with two types of interfaces: B+C

To conclude, let us set up several scenarios with the aim of evaluating the performance for more cases. The table 5.1 summarizes the results for eight different topologies giving the results for the four defined figures. Results show, like in the previous case, that the DCF performance is very poor for all scenarios, worsening when the total number of stations increases. On the other hand, the proposed configuration is always very close to the maximum achievable values, while the use of the approximate expression of 4.2 results in a small performance decrease.

Table 5.1: Performance of the four considered configuration approaches for a (overhearing) heterogeneous WLAN scenario.

Scenario			Performance			
N_A	N_B	N_C	DCF	Approx.	EF-Config	Exhaustive
5	5	5	-5.99	-0.57	-0.29	-0.27
	5	10	-5.27	5.04	5.46	5.55
	10	5	-17.33	-6.31	-5.85	-5.82
	10	10	-19.62	-1.67	-1.05	-0.96
10	5	5	-22.14	-11.47	-11.23	-11.20
	5	10	-24.48	-6.90	-6.53	-6.45
	10	5	-36.86	-18.61	-18.20	-18.18
	10	10	-42.13	-14.82	-14.26	-14.19

5.3 Non-overhearing scenario

Some modern Wireless Network Interface Cards (Wireless NICs) are able to not overhear all the transmissions. That is, instead of listening to all transmissions regardless of their destination (and therefore spending most of the time in the reception state), stations only listen to the preamble T_{PLCP} and the header H, and, in case the transmission is intended to a different station, remain the rest of the transmission time in the idle state.

In order to assess the performance of our EF-optimal configuration in a non-overhearing scenario we will model a heterogeneous set of stations where, again, the three different interfaces in table 2.1 are evenly used in the topology. The results, as seen in figure 5.6, show that, though achieving smaller EF values than the optimum possible, our configuration still performs better than the one optimized for overall throughput. However, a more detailed analysis of this scenario constitutes part of our future research.

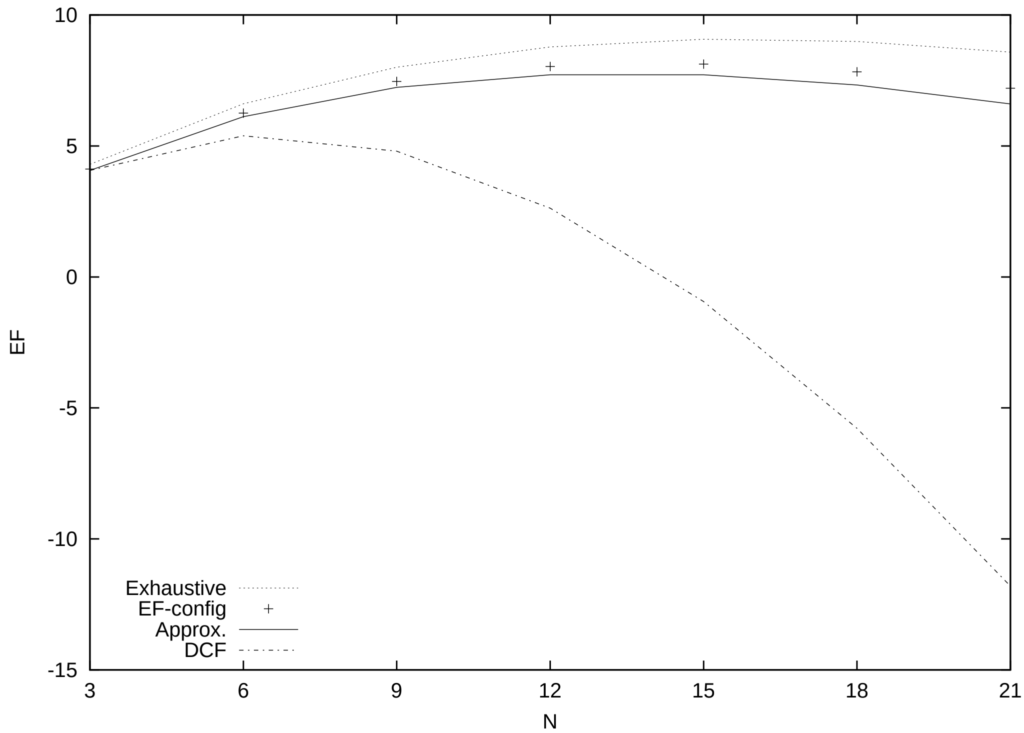

Figure 5.6: Performance of the four considered configuration approaches for a (non-overhearing) heterogeneous WLAN scenario.

Chapter 6

Experimental measurements

Current manufacturers of wireless chipsets provide basic power information regarding their interface cards. However, it is unclear the accuracy of this information. In order to analyze and understand the energy consumption of WLAN devices we have run a series of experiments that will allow us to determine the power characteristics of a generic wireless Network Interface Card (NIC).

6.1 Set-up

We set up a scenario like the one displayed in Fig. 6.1. The station 1 (STA 1) is an ASUS WL-500G P with a wireless NIC Alfa Network 11a/b/g MiniPCI Card AWPCI085S installed and associated to an AP which uses a D-Link DWL-AG660 card. The power analyzer is a PCE-PA 6000 power analyzer from PCE-iberica[1] and the wireless monitor is a laptop with a serial connection to the power analyzer and a wireless card set to monitoring the channel. The traffic will be generated using Iperf and we will use 802.11a (to lessen the impact from other WLANs) with a transmission rate of 6Mbps and a transmission power of 15dBm.

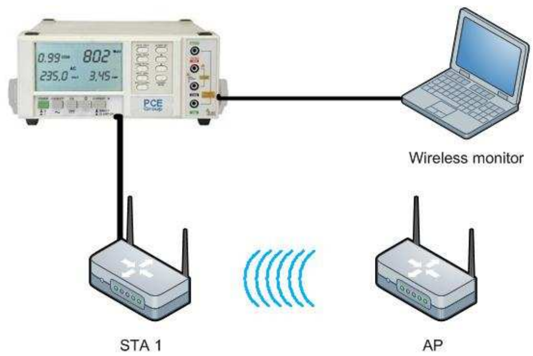

Figure 6.1: Scenario for the experiments

While running controlled experiments we will collect the following information:

[1] http://pce-iberica.es/

- Sniffer traces: The wireless monitor will capture in a trace file all the frames exchanged during the experiment. This will allow us later processing of this information. For this we will use tshark[2] over a monitoring wireless interface.

- Power consumption: The power analyzer will measure the energy consumed by STA 1 for the same period of time.

In order to measure the power parameters of a wireless NIC we will have to perform the numerical analysis on transmissions traces. We will use the three power parameters from Chapter 3 to characterize the interface's behavior:

- $\rho_i^{tx} \equiv$ Power consumption of interface i while transmitting.
- $\rho_i^{rx} \equiv$ Power consumption of interface i while receiving.
- $\rho_i^{id} \equiv$ Power consumption of interface i while idling.

As one interface can be in three different states: transmitting, receiving or idling, we can calculate the total energy consumed with the following expression:

$$e = \rho^{id} \cdot T_{id} + \rho^{tx} \cdot T_{tx} + \rho^{rx} \cdot T_{rx} \tag{6.1}$$

where

$T_{tx} \equiv$ Time spent in the transmitting state.
$T_{rx} \equiv$ Time spent in the receiving state.
$T_{id} \equiv$ Time spent in the idling state.

In order to compute these values we run controlled experiments where we saturate the channel and ensure a negligible number of retry frames (i.e., using an non-populated channel). The computerized processing of the trace files will allow us to count the number of frames and their length to calculate the value of T_{tx} and T_{rx}. Assuming saturation and a long transmission time T_T, T_{id} can be calculated using:

$$T_{id} \approx (N_{tx} + N_{rx}) \cdot (DIFS + \frac{CW_{min}}{2} \cdot T_e + SIFS) \approx T_T - T_{tx} - T_{rx}$$

where

$N_{tx} \equiv$ Number of frames transmitted.
$N_{rx} \equiv$ Number of frames received.
$DIFS, SIFS, T_e \equiv$ Constants defined by the standard 802.11.

[2] http://www.wireshark.org/

6.2 Measurements

Once we have described the set up, we will proceed to explain the results obtained. Each experiment will consist in a series of four different measurements, each of one performed for a time $T_T = 300$ seconds:

1. STA 1 with the NIC uninstalled. This way measuring the base power consumption of the device.

2. STA 1 with the wireless interface up. Note that though T_{id} should be relatively high, it still can receive/transmit frames alien to the measurement, i.e beacon frames.

3. STA 1 transmitting to the AP, where T_{tx} should be relatively high.

4. AP transmitting to the STA 1, where the time devoted to receive data (T_{rx}) should be the highest.

We will perform a series of experiments to calculate the power properties of a wireless card. The processing of the traces for each measurement will give us the T_{id}, T_{rx}, and T_{tx} that will be used in the equation 6.1. The results from the power analyzer (and gathered by the wireless monitor) provide the total energy consumed. Hence, we will have four equations that can be solved using matrix computation:

$$
[\ T\ ] \times [\ \rho\ ] = [\ e\ ]
$$

$$
\begin{bmatrix} T_T & T_{id1} & T_{rx1} & T_{tx1} \\ T_T & T_{id2} & T_{rx2} & T_{tx2} \\ T_T & T_{id3} & T_{rx3} & T_{tx3} \\ T_T & T_{id4} & T_{rx4} & T_{tx4} \end{bmatrix} \times \begin{bmatrix} \rho^{base} \\ \rho^{id} \\ \rho^{rx} \\ \rho^{tx} \end{bmatrix} = \begin{bmatrix} e_1 \\ e_2 \\ e_3 \\ e_4 \end{bmatrix} \tag{6.2}
$$

Tables 6.1 and 6.2 show the results for four experiments done at different times of day. After these measurements, we use the numbers to fill up the matrix T and e and obtain a linear equations system $AX = B$ whose solution for the case 3 is given by

$$
\begin{pmatrix} \rho^{base} \\ \rho^{id} \\ \rho^{rx} \\ \rho^{tx} \end{pmatrix} = \begin{pmatrix} 4.011 \\ 0.9899 \\ 2.3035 \\ 3.2192 \end{pmatrix} [Watts]
$$

The solutions for the rest of the measurements provide very similar values. However, as part of our future work, we will construct an overdetermined system of equations that can be solved using least squares methods in order to improve robustness of the results.

Table 6.1: Power consumed (in mJ) per event for the interfaces of of Table 2.1

			Transmitted		Received		
#	Experiment	Watts	Data	ACK	Data	ACK	Beacon
1	1	4.0114	0	0	0	0	0
	2	5.0014	1	65	3	1	2921
	3	6.3257	14	131185	131174	14	2926
	4	6.9942	128869	51	13	128867	2917
2	1	4.1785	0	0	0	0	0
	2	5.0528	1	50	4	1	2919
	3	6.3385	13	131224	131228	13	2931
	4	6.9814	128924	62	10	128926	2927
3	1	4.0114	0	0	0	0	0
	2	5.0014	1	52	6	1	2916
	3	6.2357	20	131211	131214	20	2930
	4	6.9942	128880	64	13	128878	2930
4	1	4.0110	0	0	0	0	0
	2	4.9885	1	66	3	1	2920
	3	6.0300	13	131186	131188	13	2933
	4	6.9942	128856	57	8	128859	2932

Table 6.2: Power consumed (in mJ) per event for the interfaces of of Table 2.1

#	Experiment	Watts	Transmitted Time($\mu secs$)	Received Time($\mu secs$)	Idled Time($\mu secs$)
1	1	4.0114	0	0	0
	2	5.0014	4799	6249	299992280
	3	6.3257	5538736	271399594	24269956
	4	6.9942	266632103	5439311	23843746
2	1	4.1785	0	0	0
	2	5.0528	4169	8318	299987513
	3	6.3385	5538305	271511278	24279521
	4	6.9814	266746360	5435582	23853654
3	1	4.0114	0	0	0
	2	5.0014	4253	12456	299983291
	3	6.2357	5552242	271482606	24278242
	4	6.9942	266655408	5439773	23845989
4	1	4.0110	0	0	0
	2	4.9885	4841	6249	299988910
	3	6.0300	5536709	271428518	24272153
	4	6.9942	266605458	5428630	23840672

The above procedure provides the power properties in term of ρ's parameters for the network card being tested. However, we detected some unexpected behavior of the energy consumption reported. To illustrate this, we run several experiments with different devices. Fig 6.2 shows the energy consumption for a Soekris net4826 box (`www.soekris.com`) using the same network card as the previous experiments for the different states we defined before. The difference between the figure on the left and the figure on the right is that the latter was performed after a long time of prior measurements (note that both subfigures use the axis).

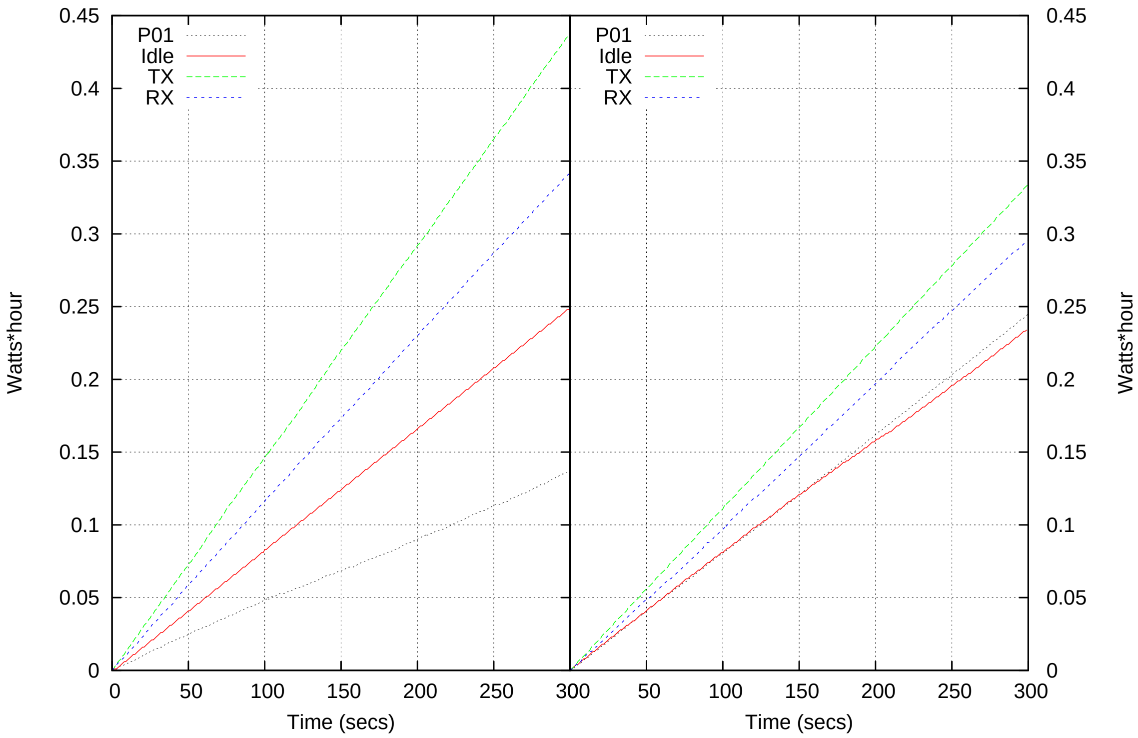

Figure 6.2: Measurements on a Soekris box for different device's uptime

This figure shows different values of energy consumed for the same period of time depending on the uptime of the devices. This way, in the second subplot (latest uptime), the energy devoted to idling is lower than the base energy! We found that one reason could be that the zero-calibration point of the measurement device had changed.

In order to understand these results, we compared the base power consumed by a Soekris box with that of a 60W lamp for a long period of time, in this case 10000 seconds. We calculated the power consumption using a time window of 50 seconds. The results, displayed in fig. 6.3, show that decrease in the reported power consumption occurs in both cases. The remaining challenge is to identify the source of this behavior and carefully assess the reliability of the measuring tools.

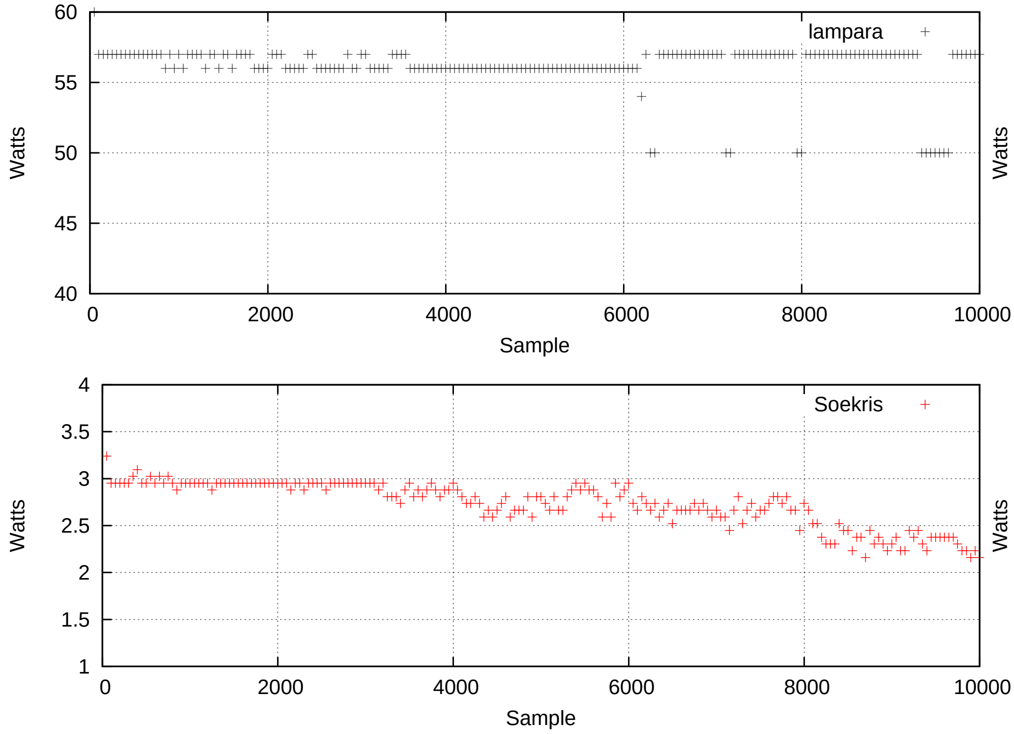

Figure 6.3: Measurements on a Soekris box for different device's uptime

Chapter 7

Conclusions and future work

The increasing concern about the energy consumption of ICTs has motivated the study of the efficiency of communication protocols, as well as new and "greener" proposals. However, most of these works are based on the unrealistic assumption of homogenous devices, i.e., all nodes having the same power consumption characteristics. While this may be the case for some particular scenarios (e.g., sensor networks), it cannot be safely assumed given the diversity of WLAN devices, and therefore this assumption has to be relaxed.

In this work we have first showned why these heterogeneous scenarios, as compared against previous homogeneous problems, constitute a different research challenge. We have identified the risk of extreme unfairness if the overall efficiency is the only variable taken into account, and have proposed the EF criterion to achieve a trade-off between energy efficiency and throughput fairness. We have then used an analytical model of the power consumption of an IEEE 802.11 WLAN to derive a closed-form expression of the configuration to use in order to achieve EF fairness. For this case, we have seen that this criterion results in a throughput-fair allocation, but the criterion itself could also be applied to other scenarios, which constitutes part of our future work. The proposed configuration has been validated through extensive simulations, and has been showned to substantially outperform the default configuration, being very close to the maximum achievables values derived from exhaustive searches on the configuration space.

We have also presented an experimental methodology to characterize the power consumption of wireless devices. We have derived some preliminary figures about the power consumption of an 802.11 interface. However, althought these numbers are consistent with the considered analytical model, we have also identified that the proposed methodology requires a more careful validation given the bias introduced by the measurement device. How to characterize this bias and the methodology to lessen its impact constitues part of our future work as well.

References

[1] V. Baiamonte and C.-F. Chiasserini. Saving energy during channel contention in 802.11 wlans. *Mob. Netw. Appl.*, 11(2):287–296, 2006.

[2] A. Banchs and P. Serrano. Revisiting 802.11e edca performance analysis. *Wirel. Pers. Commun.*, 43(4):1145–1149, 2007.

[3] G. Bianchi. Performance analysis of the ieee 802.11 distributed coordination function. *Selected Areas in Communications, IEEE Journal on*, 18(3):535–547, Mar 2000.

[4] R. Bruno, M. Conti, and E. Gregori. Optimization of efficiency and energy consumption in p-persistent csma-based wireless lans. *Mobile Computing, IEEE Transactions on*, 1(1):10–31, Jan-Mar 2002.

[5] M. Ergen and P. Varaiya. Decomposition of energy consumption in ieee 802.11. In *Communications, 2007. ICC '07. IEEE International Conference on*, pages 403–408, June 2007.

[6] IEEE 802.11e. *Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications: Medium Access Control (MAC) Enhancements for Quality of Service (QoS)*. Supplement to IEEE 802.11 Standard, November 2005.

[7] R. Jain, W. Hawe, and D. Chiu. A quantitative measure of fairness and discrimination for resource allocation in shared computer systems. In *DEC-TR-301*, September 1984.

[8] F. Kelly, A. Maulloo, and D. Tan. Rate control in communication networks: shadow prices, proportional fairness and stability. In *Journal of the Operational Research Society*, volume 49, 1998.

[9] P. Serrano, A. Garcia-Saavedra, M. Hollick, and A. Banchs. On the energy efficiency of ieee 802.11 wlans. In *European Wireless 2010, Invited Session on Greening of the Wireless Internet*, April 2010.