

On the Energy Efficiency of IEEE 802.11 WLANs

Pablo Serrano, Andres Garcia-Saavedra
Depto. de Ingeniería Telemática
Universidad Carlos III de Madrid
Avda. Universidad, 30
28911 Leganés, Spain
{pablo, agsaaved}@it.uc3m.es

Matthias Hollick
Secure Mobile Networking Lab
Technische Universität Darmstadt
Mornwegstr. 32
64293 Darmstadt, Germany
matthias.hollick@seemoo.tu-darmstadt.de

Albert Banchs
Depto. de Ingeniería Telemática
Universidad Carlos III de Madrid
Avda. Universidad, 30
28911 Leganés, Spain
banchs@it.uc3m.es

Abstract—Understanding the energy consumption of wireless interfaces is critical to provide guidelines for the design and implementation of new protocols or interfaces. In this work we analyze the energy performance of an IEEE 802.11 WLAN. Our contributions are as follows: *i*) we present an accurate analytical model that is able to predict the energy consumption, *ii*) we present an approximate model that sacrifices accuracy for the sake of analytical tractability, *iii*) based on this simplified analysis, we derive the optimal configuration to maximize energy efficiency of a WLAN, and *iv*) finally, we also analyze the trade-off between throughput and energy efficiency that IEEE 802.11 imposes. While most of these results consider a homogeneous WLAN scenario where all stations share the same energy features, we also discuss the case of heterogeneous environments, where different devices show different power consumption characteristics.

I. 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 consumption¹, 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.

However, 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.

In this paper we assess to which extent the (old) throughput maximization and the (new) energy-efficiency maximization objectives diverge for the case of IEEE 802.11 WLANs. Previous work has solved the configuration of WLANs for throughput maximization, starting from the static approaches of [1], [2] and including later adaptive approaches that tune the WLAN parameters to the estimated conditions [3], [4]. However, from the point of view of energy consumption, most of the research so far has addressed the analytical or experimental characterization of the energy consumption of the WLAN [5]–[7], which is typically divided in three states: transmission, reception and idle-state (see Table I for the energy consumption of selected wireless network cards).

There has been also some proposals for energy efficiency optimization (e.g. [8]–[10]), typically based on heuristics and sometimes requiring changes to the MAC layer. To the best of our knowledge, apart from our preliminary work in [11] only Bruno et al. [12] have considered the relation between throughput and energy and have discussed whether these could be both jointly maximized or not. However, their model consists of a p -persistent CSMA-based WLAN, where interfaces only consumed energy in two states (transmission and reception), instead of the three states described above. For this simplified case, it was shown that energy efficiency and throughput do not constitute different objectives and can be jointly maximized. In this paper, with the improved consumption model, we prove that this is not always the case. Furthermore, we also address the challenge of heterogeneous WLANs, i.e., scenarios where stations have different power consumption figures, to analyze which variable should be optimized.

The rest of the paper is organized as follows. In Section II we present and validate the analytical model of the power consumption of a WLAN. The analytical model consists of an accurate model and an approximate model, the latter of which sacrifices some numerical accuracy for analytical tractability. Section III introduces the key variable of energy efficiency, and derives the optimal configuration that maximizes this performance metric. In Section IV we discuss the relation between throughput and energy efficiency for both the case of homogeneous WLANs and the case of heterogeneous WLANs, while Section V summarizes the paper and states the main conclusions.

¹As reported in “SMART 2020: Enabling the Low Carbon Economy in the Information Age”, The Climate Group, available at <http://www.smart2020.org/>

TABLE I
POWER CONSUMPTION IN WATTS AS REPORTED IN [8] FOR DIFFERENT
WIRELESS INTERFACES

#	Card	ρ_t	ρ_r	ρ_i	ρ_r/ρ_i
A	Lucent WaveLan	1.650	1.400	1.150	1.28
B	SoketCom CF	0.924	0.594	0.066	9.00
C	Intel PRO 2200	1.450	0.850	0.080	10.63

II. POWER CONSUMPTION MODEL

Our analytical model for the energy consumption of a WLAN requires the following input parameters: N , the number of stations in the WLAN; CW_{min} , defined as the minimum contention window stations use on their first transmission attempt; and $\{\rho_t, \rho_r, \rho_i\}$, defined as the power consumed by the wireless interfaces when transmitting, receiving and idling, respectively. We assume all stations have always a packet of fixed length L ready for transmission², i.e., the network operates under saturation conditions, and that the sole reason for frame loss is a collision (i.e., where two or more stations transmit simultaneously). We further assume that each station randomly selects the destination for each frame out of the other $N - 1$ stations.

A. Model

With the assumption that each transmission attempt collides with a constant and independent probability, we can model the behavior of a station with the same Markov chain used in [1]. Then, the probability that a station operating under saturation conditions transmits upon a backoff counter decrement τ can be computed by means of the following equation given by [1]

$$\tau = \frac{2}{1 + CW_{min} + pCW_{min} \sum_{i=0}^{m-1} (2p)^i}$$

where p is the probability that a transmission attempt of a station collides. This probability can be computed as

$$p = 1 - (1 - \tau)^{N-1}$$

The above constitutes a system of two non-linear equations that can be solved numerically, giving the value for τ . However, for the case of $CW_{min} = CW_{max}$, the transmission probability can be computed as

$$\tau = \frac{2}{CW + 1}$$

Once the value of τ is computed, we next proceed to compute the energy per slot consumed by a station, which we denote by e .

We compute e by applying the total probability theorem as follows:

$$e = \sum_{j \in \Theta} E(j)p(j) \quad (1)$$

²Note that, following our analysis in [13], the model could be extended to account for variable packet sizes.

where Θ is the set of events that can take place in a single timeslot³, while $E(j)$ and $p(j)$ are the energy consumed in case of event j and its probability, respectively. The set Θ of events, as well as 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, $p_{s,-i}$
- There is a collision and the considered station is involved, $p_{c,i}$
- There is a collision but the considered station is not involved, $p_{c,-i}$

The probability of each event can be easily computed based on the probability of a transmission τ as follows

$$\begin{aligned} p_e &= (1 - \tau)^N \\ p_s &= N\tau(1 - \tau)^{N-1} \\ p_{s,i} &= \tau(1 - \tau)^{N-1} \\ p_{s,-i} &= p_s - p_{s,i} \\ p_c &= 1 - p_e - p_s \\ p_{c,i} &= \tau(1 - (1 - \tau)^{N-1}) \\ p_{c,-i} &= p_c - p_{c,i} \end{aligned}$$

We proceed similarly to [6] to compute the energy consumed by a particular station depending on each event j . With this, we can expand (1) as follows:

$$\begin{aligned} e &= p_e \rho_i T_e + \\ &+ p_{s,i} (\rho_t T_s + \rho_r T_{ack} + \rho_i (SIFS + DIFS)) + \\ &+ p_{s,-i} \left[\rho_r T_s + \frac{1}{N-1} (\rho_t T_{ack}) + \right. \\ &+ \left. \frac{N-2}{N-1} (\rho_r T_{ack}) + \rho_i (SIFS + DIFS) \right] + \\ &+ p_{c,i} (\rho_t T_s + \rho_i EIFS) + p_{c,-i} (\rho_r T_s + \rho_i EIFS) \end{aligned}$$

where T_e , T_s , and T_{ack} are the durations of an empty slot, a successful transmission and the transmission of an acknowledgment, while $SIFS$, $DIFS$, and $EIFS$ are physical constants (for the computation of these values, see e.g. [1]).

However, note that the full expression of (1) consists of a sum of several terms that non-linearly depends on τ . We next introduce an simplified expression for e , denoted by \hat{e} , that will be used in Section III-A to derive the τ value that provides optimal energy performance, because of its analytical tractability.

$$\begin{aligned} \hat{e} &= (1 - \tau)^N \rho_i T_e + \tau (\rho_t T_s + \rho_r T_{ack} \\ &+ \rho_i (SIFS + DIFS)) + (1 - \tau) (1 - (1 - \tau)^{N-1}) \cdot \\ &(\rho_r T_s + \rho_t T_{ack} + \rho_i (SIFS + DIFS)) \end{aligned}$$

This way, we have simplified the set Θ of events by considering only three cases:

³A timeslot is defined as the amount of time between two backoff counter decrements of a station, see [1].

- Nobody transmits, and therefore the station remains in state idle.
- The considered station transmits, and we assume it is a successful transmission.
- Some other station transmits, and we assume again it is a successful transmission.

The above can be expressed as:

$$\hat{e} = R + \tau(T - R) - (1 - \tau)^N(R - E)$$

where

$$E = \rho_i T_e$$

$$T = \rho_t T_s + \rho_r T_{ack} + \rho_i(SIFS + DIFS)$$

$$R = \rho_r T_s + \rho_t T_{ack} + \rho_i(SIFS + DIFS)$$

We further write $T' = T - R$ and $R' = R - E$, therefore:

$$\hat{e} = R + \tau T' - (1 - \tau)^N R' \quad (2)$$

Given the average energy consumption of a station per timeslot, the power consumption of a wireless interface is computed as

$$\pi = \frac{e}{T_{slot}} \quad (3)$$

where T_{slot} is the average slot duration, computed as

$$T_{slot} = p_e T_e + (1 - p_s - p_c) T_s \quad (4)$$

B. Validation

We first compare the accuracy of the exact and approximate energy consumption models against results obtained via simulation, using the physical layer parameters of the IEEE 802.11b standard. To this end, we first compare the energy consumed per second π for the three power consumption profiles of Table I and different values of N , using the default DCF configuration. We plot in Fig. 1 the resulting numbers for three cases: the “complete” analytical model, the approximated model, and the use of simulations. From the results, it is clear that the detailed analytical model e provides values that almost coincide with those derived from simulations, while the approximate model \hat{e} reasonably follows the behavior of the WLAN⁴.

To gain further the energy behavior of the WLAN, we set the number of stations to $N = 10$ and obtain the energy consumed for different values of the contention window used. The results are given in Fig. 2 and can be explained as follows: for small values of CW , stations are most of the time attempting transmission (although most of these result in collision), and therefore the power consumed approaches ρ_t . On the other hand, for large CW all stations spend most of the time in the idle state, and therefore the power consumed tends to ρ_i (note that actually reaching ρ_i would require the use of an *extremely large* CW value, to set $\tau \approx 0$).

⁴Note that the approximate model \hat{e} overestimates the energy consumed because all transmissions are assumed to be successful, this resulting in an increased energy consumption because of the acknowledgement.




Fig. 1. Power consumption for the energy profiles of Table I with the standard DCF configuration.




Fig. 2. Power consumption for the energy profiles of Table I for different CW_{min} configurations, $N = 10$.

The previous analysis presents a model to capture the power consumption of a WLAN. However, it does not provide any information about the way the energy is consumed, e.g., it could be that large values of power consumption of Fig. 2 corresponds to states where large volumes of data are being sent in the WLAN, while small values of π corresponds to a completely idle WLAN and therefore all energy is being wasted. We take advantage of the accurate analytical model to further explore the energy consumption of the WLAN, identifying where the energy consumption is rooted. To this aim, for the same case of $N = 10$ stations we obtain the relative amount of energy wasted during empty slots or collisions for different values of the contention window, with the results of Fig. 3 for interfaces A and C of Table I.

This figure is particularly relevant, as it helps to further gain insight into the observed behavior of Fig. 2. Indeed, although in the previous figure the power each device consumes does not seem to significantly change with the CW_{min} used, in this figure we observe that the energy is indeed spent in different

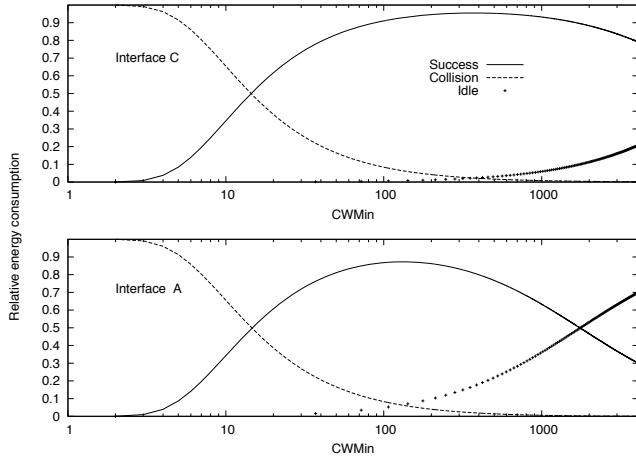


Fig. 3. Relative energy devoted to collisions and idling for the selected energy profiles A and C from Table I

events: the relative amount of energy devoted to collisions decreases with CW_{min} , while the energy spend during idling increases. Furthermore, it is approximately at the CW_{min} value that makes the energy devoted to collisions equal to the energy devoted to idling where the energy devoted to successful transmission is maximized. This would be, intuitively, the optimal configuration to use in a “green” WLAN. How to compute this configuration is addressed in the next section, with the introduction of the energy efficiency.

Remark: Note that the previous throughput-optimization approach of [3] achieved its objective by adaptively forcing the time spent in collisions to be equal to the time spent in idle slots. For the case of energy optimization, we have a similar behavior, although in this case the *cost* of a timeslot that contains a collision or is empty is “multiplied” by the power consumed when transmitting/receiving or idling, respectively.

III. ENERGY EFFICIENCY

We have just seen in Fig. 3 that depending on the CW used, the energy can be devoted to successful transmissions, idling or collisions, and that there seems to be a CW value that maximizes the energy devoted in successful transmissions. Motivated by this finding, we define the energy efficiency η as the ratio between the bits transmitted and the energy consumed in a timeslot:

$$\eta = \frac{p_{s,i}L}{e} \quad (5)$$

Next we compare the efficiency η for three different WLAN scenarios (one for each of the interfaces of Table I) and the standard recommended configuration of CW_{min} . To this aim, we plot in Fig. 4 the value of η given by simulations against the ones provided by the accurate analytical model of (5), and the simplified model, i.e., using (5) but substituting e with \hat{e} .

From the figure it is evident that the default recommendation shows an efficiency η that rapidly decreases with the number of stations N , an expected results because the well-known throughput degradation of DCF with N (caused by collisions)




Fig. 4. Energy efficiency of the default DCF configuration vs. N , for the selected energy profiles of Table I




Fig. 5. Energy efficiency of the WLAN for $N = 10$ stations and the interfaces of Table I

is further emphasized by the relatively larger weight of collisions over idle slots (because of the energy consumption).

To analyze the impact of the CW value used on η , we analyze the relation between the efficiency and CW_{min} for the case of $N = \{10, 20\}$ in Figs. 5 and 6, respectively (note that $CW_{min} = 32$ corresponds to the standard configuration). We plot the energy efficiency of the WLAN for different values of CW , where we represent with points the results obtained from simulations and with lines the ones derived from the analytical models.

From the figures, we can see two main results: first, the approximate model follows the behavior of the actual energy consumption, in particular in the region where the maximum value of η is reached⁵; second, for each WLAN interface and number of stations this maximum value of η is obtained with different CW values. Note that this observation is aligned with the results obtained in Fig. 3, where the relative costs of idling

⁵Note that this is indeed very relevant, as we will use the approximate model to derive the most energy-efficient configuration.

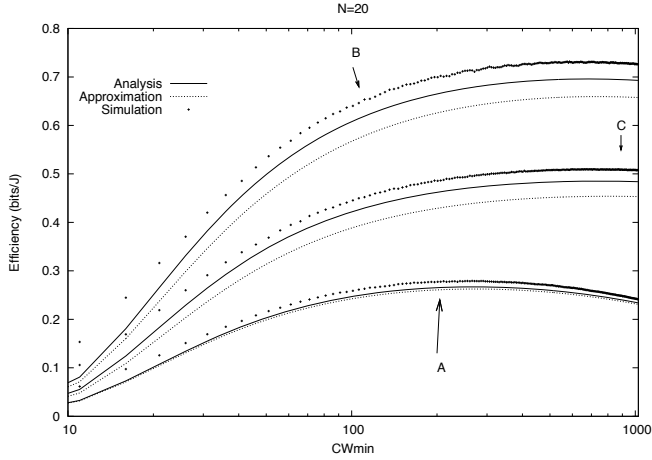


Fig. 6. Energy efficiency of the WLAN for $N = 20$ stations and the interfaces of Table I

and colliding are multiplied by different values depending on the interface. This way, the CW value that achieves best performance is larger when the ρ_i value is relatively smaller, as the cost of a collision is relatively larger. In the next section we analyze how to obtain this maximum value, i.e., the relation between the CW that optimizes energy efficiency and the power consumption parameters of the WLAN card.

A. Maximum Efficiency Configuration

We provide in this section closed-form expressions for the optimal transmission probability τ . Note that, to derive a configuration rule, if we set $CW_{min} = CW_{max}$, the transmission probability τ is related to the size of the contention window CW to be used as follows

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

To compute the transmission probability that optimizes the consumption of energy τ_e we start from the expression of η with the approximation for \hat{e}

$$\eta = \frac{\tau(1-\tau)^{N-1}L}{R + \tau T' - (1-\tau)^N R'}$$

And then compute the τ value that maximizes the above by

$$\frac{d\eta}{d\tau} = 0$$

which leads to

$$(N-1)\tau^2 T' + (1-\tau)^N R' + N\tau R - R = 0$$

By the following Taylor expansion of $(1-\tau)^N$

$$(1-\tau)^N \approx 1 - N\tau + \frac{1}{2}N(N-1)\tau^2$$

We have the following equation

$$a\tau^2 + b\tau + c = 0$$

where

$$a = (N-1)T' + \frac{1}{2}N(N-1)R'$$




Fig. 7. Energy efficiency of the WLAN for $N = \{10, 20\}$ stations and the interfaces of Table I

$$b = NE$$

$$c = -R$$

If we now define α and β as follows

$$\alpha = \frac{T'}{E}, \quad \beta = \frac{R'}{E}$$

Then we have the following for the computation of τ_e :

$$\tau_e = \frac{-N + \sqrt{N^2 + 4(N-1)\alpha + 2N(N-1)\beta}}{2(N-1)\alpha + N(N-1)\beta}$$

Which can be approximated as follows

$$\tau_e \approx \frac{1}{N} \sqrt{\frac{2}{\beta}} \approx \frac{1}{N} \sqrt{\frac{2\rho_i T_e}{\rho_r T_s}} \quad (6)$$

B. Validation

We validate the previous expression for the optimal configuration (in terms of energy efficiency) of a WLAN for the interfaces of Table I and different values of N . To this aim, we plot in Fig. 7 the resulting η from an exhaustive search on the CW space, denoted by η_{ex} , and the one derived from the use of (6), denoted by η_{conf} . For all scenarios considered, the resulting η values almost coincide in all cases, a result caused by the “flatness” of η in the region close to the maximum value (as seen in, e.g., Fig. 4). We conclude, therefore, that indeed (6) provides the most energy-efficient configuration to use in a WLAN.

IV. ON THE RELATION BETWEEN THROUGHPUT AND ENERGY EFFICIENCY

A. Homogeneous WLANs

For the case of saturated IEEE 802.11 WLANs, Bianchi [1] analytically derived the optimal transmission probability τ by maximizing the following expression for throughput

$$R = \frac{p_s L}{T_{slot}}$$

where T_{slot} is the average slot duration, given by

$$T_{slot} = (1 - \tau)^N T_e + (1 - (1 - \tau)^N) T_s$$

This optimization is done by deriving the above with respect to τ , and solving a second-grade equation resulting from the approximation $\tau \ll 1$. This results in the following approximate value for the optimal transmission probability that maximizes throughput, which we denote by τ_t

$$\tau_t \approx \frac{1}{N} \sqrt{\frac{2T_e}{T_s}} \quad (7)$$

Note that this optimal value of τ depends on the number of stations N , but also on the relative size of an empty timeslot T_e as compared to a timeslot that contains a transmission T_s . This way, apart from the number of stations, the ratio between the timeslot lengths sets the optimal tradeoff between the *cost of a collision* and the *cost of idling*, as we already saw at the end of Section II.

However, because τ_t does not take into account energy consumption, for similar scenarios with different WLAN interfaces it will provide the same configuration for CW , while we have seen in Fig. 4 that the optimal CW value indeed depends on the energy consumption of the WLAN interfaces. Actually, if we divide (7) by (6), the relation between τ_t and τ_e is given by the ratio of the power consumption of the interface when receiving a frame over the power consumption when idling, i.e.,

$$\frac{\tau_t}{\tau_e} = \sqrt{\rho_r / \rho_i}$$

Therefore, the optimal value of τ for the case of energy efficiency depends not only on the relative size of the timeslots, like in the case of throughput maximization, but also on the relative power consumed when receiving or idling. This way, for *old* interfaces where idling and receiving consumes approximately the same value (e.g., interface A of Table I) the formula of Bianchi still holds, while for *new* interfaces where the cost of idling is smaller (e.g., interface C), the CW_{min} value to use is larger than the one obtained with (7).

Throughput maximization and energy efficiency optimization, then, pursue different objectives. To illustrate this, we plot in Fig. 8 the throughput and energy efficiency achieved for the case of $N = 10$, with the three interfaces of Table I, and using different values of the transmission probability τ (for the case of throughput there is a single line). Let us focus on the case of interface C: the τ value that achieves optimal energy efficiency performance is approximately 0.007, while the value that maximizes throughput is around 0.02. If we divide these, the ratio is approximately the one obtained with the square root of the relative power consumption given in Table I, i.e., $\sqrt{\rho_r / \rho_i} = \sqrt{10} \approx 3$.

We have, therefore, that there is a different configuration for CW depending on the variable (i.e. throughput or energy) to optimize. We next analyze this tradeoff, i.e., the price to pay in terms of energy efficiency (throughput) when the variable to maximize is the throughput (energy efficiency). We first analyze the reduction in energy efficiency when the variable to




Fig. 8. Energy efficiency and Throughput vs τ




Fig. 9. Relative energy efficiency loss for the throughput-optimized configuration

maximize is the throughput. To this aim, we obtain the relative throughput loss by computing the ratio between the total WLAN throughput obtained with the τ that maximizes energy efficiency over the throughput obtained when using Bianchi's expression. The results for different values of N are given in Fig. 9 and show that the use of τ_t inflicts performance loss on η that, as expected, depends on the considered interface. For the case of interface A the difference is negligible, while for the other two cases the price to pay when using the maximum throughput configuration is around 5% of energy efficiency.

We next analyze the reduction in throughput for the case of energy optimization with a similar approach. This way, we first compute the configuration that maximizes throughput, and then obtain the relative loss of performance per each wireless interface when using the configuration that maximizes the energy efficiency. The results are given in Fig. 10, showing a similar behavior to the previous case. However, in this case the price to pay in throughput to achieve the most energy efficient configuration is slightly larger, as it requires a throughput degradation of 8–10% for interfaces B and C.




Fig. 10. Relative throughput loss for the energy-optimized configuration




Fig. 11. Trade-off imposed by the energy features of the WLAN interfaces

Therefore, these results confirm that there is a tradeoff between energy and throughput maximization that depends on the characteristics of the WLAN interface. Indeed, for some ratios of power consumption we have the same result of [12], that both throughput and energy efficiency can be simultaneously maximized. However, our results show also that, for existing WLAN interfaces, this is not always the case. These findings are summarized in Fig. 11. In this figure, we plot for the case of $N = 10$ the resulting values of the energy efficiency (in the x-axis) and throughput (in the y-axis) for all possible configurations of the CW parameter and all the interfaces of Table I. In the figure, we mark with a star the point of maximum throughput performance, and with a circle the point of maximum energy efficiency. This figure provides valuable insights on the observed behavior:

- For the case of interface A, given its relatively similar values of power consumption ρ_t , ρ_r and ρ_i , both energy and throughput can be jointly maximized, given the “linear” shape of the resulting curve.
- For the cases of interfaces B and C, the larger the ρ_r/ρ_i

ratio, the more separate the optimum values are and therefore the higher the price to pay in throughput when optimizing energy (and vice-versa).

- Lastly, all curves share the same maximum throughput value, because of the use of the physical layer parameters of 802.11b⁶. However, their maximum efficiency value is not tied only to the ρ_i parameter or the ρ_r/ρ_i ratio, but instead it depends on the whole set of ρ parameters.

This finding backs our above analysis and demonstrates that, depending on the power characteristics of the WLAN interfaces, gains w.r.t. energy are achievable if a energy-optimized configuration is chosen over a throughput-optimized one.

B. Heterogeneous WLANs

In the previous sections we have only addressed *homogeneous* scenarios, i.e., deployments where all stations’ interfaces have the same power consumption behavior. However, in real-life scenarios different users may have different interfaces and, therefore, it could be that, e.g., for some users the *cost* of idling is relatively much larger than for other users. While for the case of throughput maximization there is no difference between heterogeneous and homogeneous scenarios, for the case of energy efficiency optimization the heterogeneous scenario constitutes a very different challenge: given each interface in the WLAN may have a different set of $\{\rho_t, \rho_r, \rho_i\}$ parameters, it is unclear what is the variable to maximize.

Let us consider the overall energy efficiency, i.e., the total number of bits transmitted over the total energy consumed. If this variable is to be maximized, the most immediate solution would be to *choke* those stations with larger power consumption figure (configuring an extremely large CW), while optimally configuring the least consuming interfaces with the use of (7). This solution would indeed maximize the overall efficiency, at the cost of an extreme unfairness in throughput. On the other hand, an energy-unaware configuration would provide all interfaces with the same CW value, what can be seen as sacrificing energy efficiency for the sake of fairness.

To assess this trade-off we run the following experiment. We consider a WLAN scenario with two groups of 10 stations each. We model the power consumption of one of the groups according to Interface A of Table I, while the other group follows the consumption of Interface C. We denote with $CW[1]$ the CW used by the first group, and with $CW[2]$ the CW used by the second group. Then, we set a given ratio $CW[1]/CW[2]$ and we search on the $CW[2]$ space the configuration that maximizes the overall η , computing also the throughput each group experiences ($R1$ and $R2$, respectively). The results from this experiment are depicted in Fig. 12.

As expected, the throughput ratio and the energy efficiency are strongly coupled: the larger the asymmetry in throughput towards the less consuming interface, the higher the efficiency in the WLAN, while the fair share of the wireless

⁶Therefore, the use of, e.g., 802.11g parameters would have shifted the curves to the 54 Mbps range.




Fig. 12. Trade-off between WLAN energy efficiency and throughput share

throughput ($R1/R2 = 1$) corresponds to the least efficient scenario. Therefore, to compute the “optimal” configuration in heterogeneous scenarios, a careful definition of the variable to maximize is required. This constitutes part of our current work.

V. SUMMARY AND CONCLUSIONS

While the energy consumption was only considered a key performance figure in very specific environments (e.g., sensor networks), greening the communication protocols is nowadays recognized as a primary design goal of future global network infrastructures. This design goal requires switching from “information per unit of time” to “information per unit of energy”. However, for the case of 802.11 WLANs, energy optimization has been typically addressed through the use of ad-hoc sleeping heuristics, and not through the proper tuning of the CSMA/CA access mechanism.

In this paper we have revisited previous 802.11 performance analyses to derive an approximate model, analytically tractable, of the energy consumption of a 802.11 WLAN. Based on this model, we have derived the optimal CW configuration for homogeneous scenarios to maximize energy efficiency. We have also discussed for these scenarios whether throughput and energy can be jointly maximized or not, and we have found that for currently available WLAN interfaces, there is a reduction of up to 10% of one variable when maximizing the other. Finally, we have analyzed heterogeneous scenarios, where we have shown that throughput fairness and energy efficiency do constitute significantly different optimization objectives, and therefore optimization requires the definition of a proper maximization criterion.

VI. ACKNOWLEDGMENTS

The authors would like to thank the anonymous reviewers for their valuable comments. The research leading to these results has received funding from the European Community’s Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 214994. The views and conclusions contained here are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either express of the CARMEN project or the European Commission. It was also partly funded by the Ministry of Science and Innovation of Spain, under the QUARTET project (TIN2009-13992-C02-01).

REFERENCES

- [1] G. Bianchi, “Performance analysis of the IEEE 802.11 distributed coordination function,” *Selected Areas in Communications, IEEE Journal on*, vol. 18, no. 3, pp. 535–547, Mar 2000.
- [2] H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma, “Performance of reliable transport protocol over IEEE 802.11 wireless LAN: Analysis and enhancement,” in *Proceedings of IEEE INFOCOM’02*, June 2002.
- [3] M. Heusse, F. Rousseau, R. Guillier, and A. Duda, “Idle sense: an optimal access method for high throughput and fairness in rate diverse wireless lans,” *SIGCOMM Comput. Commun. Rev.*, vol. 35, no. 4, pp. 121–132, 2005.
- [4] F. Cali, M. Conti, and E. Gregori, “Dynamic tuning of the IEEE 802.11 protocol to achieve a theoretical throughput limit,” *IEEE/ACM Trans. Netw.*, vol. 8, no. 6, pp. 785–799, 2000.
- [5] X. Wang, J. Yin, and D. P. Agrawal, “Analysis and optimization of the energy efficiency in the 802.11 dcf,” *Mob. Netw. Appl.*, vol. 11, no. 2, pp. 279–286, 2006.
- [6] M. Ergen and P. Varaiya, “Decomposition of energy consumption in IEEE 802.11,” in *Communications, 2007. ICC ’07. IEEE International Conference on*, June 2007, pp. 403–408.
- [7] L. Feeney and M. Nilsson, “Investigating the energy consumption of a wireless network interface in an ad hoc networking environment,” in *Proceedings of IEEE INFOCOM’01*, April 2001.
- [8] V. Baiamonte and C.-F. Chiasserini, “Saving energy during channel contention in 802.11 WLANs,” *Mob. Netw. Appl.*, vol. 11, no. 2, pp. 287–296, 2006.
- [9] J.-C. Chen and K.-W. Cheng, “Edca/ca: Enhancement of IEEE 802.11e edca by contention adaptation for energy efficiency,” *Wireless Communications, IEEE Transactions on*, vol. 7, no. 8, pp. 2866–2870, August 2008.
- [10] C. Wang, B. Li, and L. Li, “A new collision resolution mechanism to enhance the performance of IEEE 802.11 dcf,” *Vehicular Technology, IEEE Transactions on*, vol. 53, no. 4, pp. 1235–1246, July 2004.
- [11] A. Banchs, P. Serrano, M. Hollick, and L. Vulliamy, “Throughput and energy efficiency in IEEE 802.11 WLANs: Friends or foes?” in *ICST Qshine ’09, Las Palmas de Gran Canaria, Spain*, November 2009.
- [12] 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*, vol. 1, no. 1, pp. 10–31, Jan-Mar 2002.
- [13] P. Serrano, A. Banchs, and A. Azcorra, “A throughput and delay model for IEEE 802.11e edca under non saturation,” *Wirel. Pers. Commun.*, vol. 43, no. 2, pp. 467–479, 2007.