Throughput and Energy Efficiency in IEEE 802.11 WLANs: Friends or Foes?

Pablo Serrano[1], Albert Banchs[1], Luca Vollero[2], and Matthias Hollick[3]

[1] Universidad Carlos III de Madrid, 28911, Leganés, SPAIN, `pablo@it.uc3m.es`
`banchs@it.uc3m.es`
[2] Università Campus Bio-Medico di Roma, 00128, Roma, ITALY,
`l.vollero@unicampus.it`
[3] Technische Universität Darmstadt, 64283, Darmstadt, GERMANY,
`matthias.hollick@cased.de`

Abstract. Understanding and optimizing the energy consumption of wireless devices is critical to maximize network lifetime and to provide guidelines for the design of new protocols and interfaces. In this work we first provide an accurate analysis of the energy performance of an IEEE 802.11 WLAN, and then we derive the configuration to maximize it. We also analyze the impact of the energy configuration of the device on the throughput performance, and discuss in which circumstances throughput and energy efficiency can be both maximized and where they constitute different challenges.

Key words: Energy efficiency, energy optimization, throughput optimization, IEEE 802.11

1 Introduction

ICT technologies hold one of the keys to the reduction of greenhouse gases produced worldwide. The importance of "greening of the Internet" is thus recognized as a primary design goal of future global network infrastructures. It is estimated that, today, the Internet already accounts for about 2% of total world energy consumption, 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 fuelled 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 clear that this comes not for free, and there is a price to pay when developing sustainable architectures.

In this paper we assess to which extent the (old) throughput-maximization and the (new) efficiency-maximization objectives diverge, for the case of 802.11

WLANs. Previous work has solved the configuration of WLANs for throughput maximization, starting from the statical approaches of [2, 10] and including later adaptive approaches to maximize the bits per second sent [7]. 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 [9, 5, 6], which is typically divided in three states: transmission, reception and idle-state (see Table 1 for the energy consumption of selected wireless network cards). There has been also some proposals for efficiency optimization (e.g. [1, 4, 8]), typically based on heuristic and sometimes requiring changes to the MAC layer. To the best of our knowledge, only Bruno et al. [3] have considered the relation between throughput and energy and have discussed whether they could be both jointly maximized or not. In their model, consisting of a p-persistent CSMA-based WLAN where interfaces only consumed energy in two states (transmission and reception), the answer was yes. In this paper, where we improve the accuracy of the consumption model, we prove that this is not always the case.

The rest of the paper is organized as follows. In Section 2 we present and validate an analytical model of the energy consumption of a WLAN. We further introduce a new *approximate* model that trades off accuracy for the sake of simplicity (nevertheless, as shown in the validation part, this reduction of accuracy is negligible). Section 3 presents the two approaches for performance maximization: the throughput-based approach of Bianchi, and our energy-based approach that builds upon the approximate analysis to derive a closed-form expression for the optimal transmission probability. In Section 4 we compare the resulting configuration and performance from each approach, while Section 5 concludes the paper.

2 Energy Consumption Analysis

Our analytical model for the consumption of a WLAN requires the following input parameters: N, the number of stations in the WLAN. W, defined as the minimum contention window stations use on their first attempt, and $\{\rho_t, \rho_r, \rho_i\}$, defined as the power consumed by the wireless interfaces when transmitting, receiving or idling. We assume all stations have always a packet of fixed lenght L ready for transmission, i.e., the network operates under saturation conditions, and that the sole reason for frame loss is a collision (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.

2.1 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 [2]. Then, the probability that a station operating

Card	ρ_t	ρ_r	ρ_i
Lucent WaveLan (A)	1.650	1.400	1.150
SoketCom Compact Flash (B)	0.924	0.594	0.066
Intel PRO 2200 (C)	1.450	0.850	0.080

Table 1. Power consumption in Watts for different wireless interfaces (as reported in [1])

under saturation conditions transmits upon a backoff counter decrement can be computed by means of the following equation given by [2]

$$\tau = \frac{2}{1 + W + pW \sum_{i=0}^{m-1}(2p)^i}$$

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. With this, 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:

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

where Θ is the set of events that can take place in a single timeslot, while $E(j)$ and $p(j)$ are the energy consumed in case of event j given its probability, respectively. The set Θ contains the following events, along with their probabilities:

– 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, $p_{c,\neg i}$

This way we can expand (1) with these probabilities and the energy consumed per event can be derived as follows:

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

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. [2]).

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

$$p_e = (1 - \tau)^N$$

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

$$p_{s,i} = \tau(1 - \tau)^{N-1}$$

$$p_{s,\neg i} = p_s - p_{s,i}$$

$$p_c = 1 - p_e - p_s$$

$$p_{c,i} = \tau(1 - (1 - \tau)^{N-1})$$

$$p_{c,\neg i} = p_c - p_{c,i}$$

However, note that the full expression of (1) consists of a sum of several terms that non-linearly depends on τ. In order to derive the value of τ that provides the best energy performance, we introduce the following simplified expression for e

$$\hat{e} = (1 - \tau)^N \rho_e T_e + \tau \rho_t T_s + (1 - \tau)\left(1 - (1 - \tau)^N\right) \rho_r T_s$$

This way, we have simplified the set Θ of events by considering only three cases: i) nobody transmits, ii) the station transmits (without the distinction if there is a collision or a success), and iii) someone else transmits (again, no matter if there is a success of a collision).

The above can be expressed as:

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

where $E = \rho_e T_e$, $T = \rho_t T_s$, and $R = \rho_r T_s$. We further write $T' = T - R$ and $R' = R - E$, therefore:

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

(Note that in the following section we assess the accuracy obtained both via (1) and (2).) Finally, we define the energy efficiency η as the ratio between the bits transmitted and the energy consumed in a timeslot:
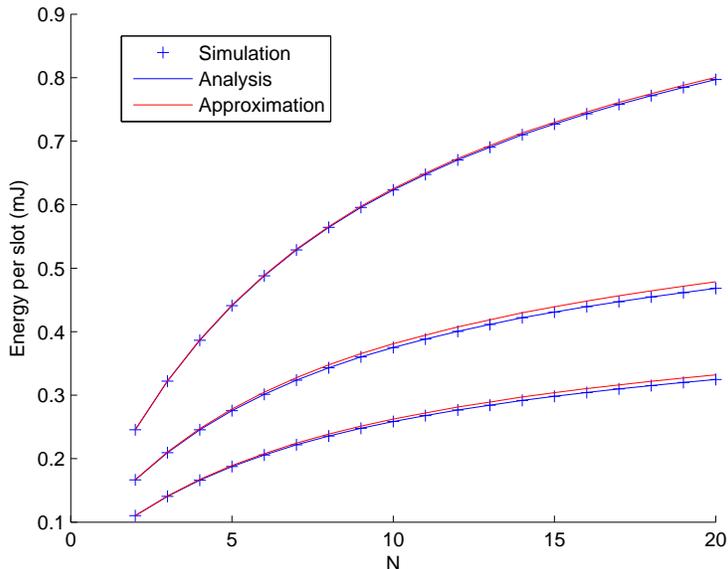
$$\eta = \frac{p_{s,i} L}{e} \tag{3}$$

Fig. 1. Energy per slot-time consumed for different interfaces and number of stations. The arrays of curves from top to bottom show the results for energy profiles (A), (C), and (B)

2.2 Validation

We first compare the accuracy of the exact and approximate models for e and $\hat{e}$ versus results obtained via simulation. To this end, we compare the energy consumed per timeslot for the three selected power consumption sets listed in Table 1 for different values of N and the default DCF configuration. Results are shown in Fig. 1.

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 $\hat{e}$ follows quite closely the behavior of the WLAN but slightly overestimating the energy consumed for large values of N.

We take advantage of the accurate analytical model to further explore the energy consumption of the WLAN, identifying where is the energy consumed. To this aim, we account for the relative amount of energy spent on successful transmissions, collisions and idling, with the results of Fig. 2 for the case of $N = 10$ and the interface A of Table 1.

As can be seen from the figure, it is clear that for relatively small values of CW_{min} there is a lot of energy wasted in collisions, while the energy spent idling is quite small. Then, with increasing CW_{min} values the energy wasted in collisions decreases rapidly, while there is a slower increase in the part corresponding to idling. This behavior is intuitively explained as follows. Increasing
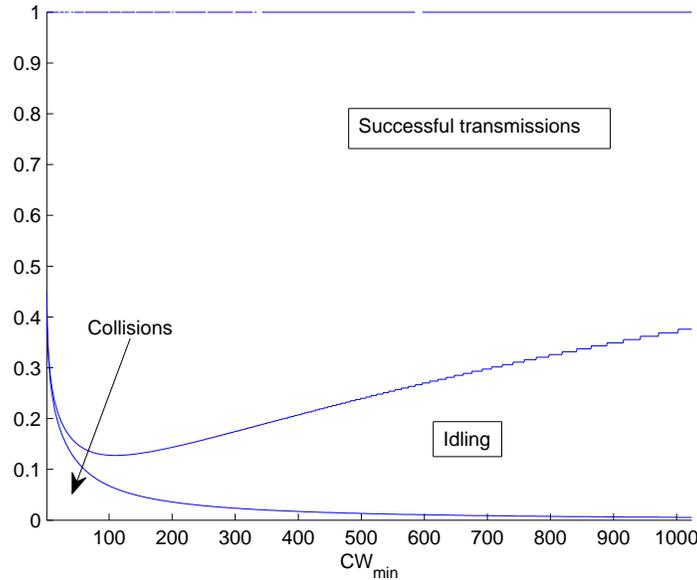
Fig. 2. Relative energy devoted to successful transmissions, collisions and idling

CW_{min} results in a smaller collision probability and larger probability of empty timeslots. However, the savings in energy due to the absence of collisions are "multiplied" by the power consumption when receiving ρ_r or transmitting ρ_t as well as, approximately, the length of a successful transmission T_s. On the other hand, the increase of energy consumption because of the larger number of empty timeslots is weighted by ρ_i and T_e, both being smaller than their counterparts.

Another result from the figure is that there exists a maximum for the energy devoted to successful transmission (in the scenario considered, for $CW_{min} \approx$ 100). This optimun value sits in the tradeoff between the decrease of the energy devoted to collisions and the increase in the energy spent when idling, and its computation is derived in Section 3.2

Finally, we compare the efficiency η for three different WLAN scenarios (one for each of the interfaces of Table 1) and $N = 10$. We compare the numerical values given by simulations against the ones provided with our simplified analytical model, i.e., using (3) but substituting e with $\hat{e}$. We can see that the model is quite accurate, in particular in the relatively "flat" region where the efficiency is maximum, and that the optimal value of CW is different for each of the WLAN scenarios—a result we analyze next.
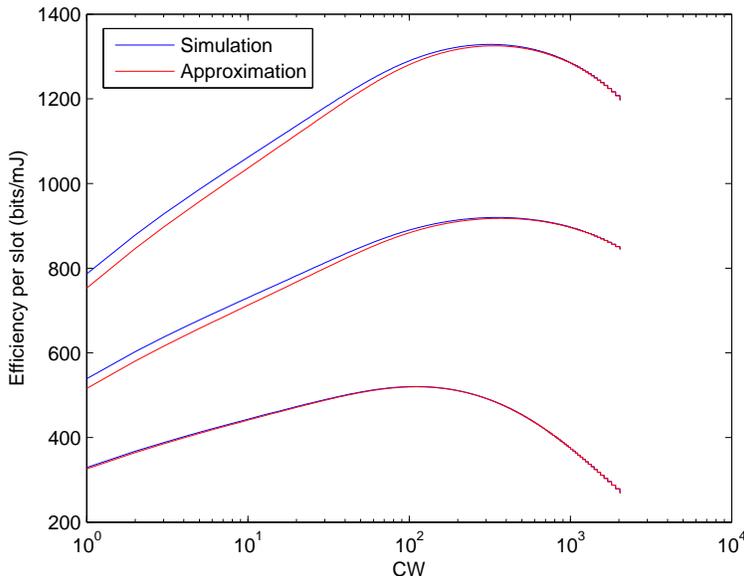
Fig. 3. Impact of the CW_{min} used on the efficiency. The arrays of curves from top to bottom show the results for energy profiles (A), (C), and (B)

3 Configuration of 802.11

We provide in this section closed-form expressions for the optimal transmission probability τ, depending on the optimization objective throughput maximization in Section 3.1, and energy optimization in Section 3.2. Note that if we set $CW_{min} = CW_{max}$, the transmission probability τ is easily related to the CW to use as follows

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

3.1 Throughput maximization

When optimizing throughput, it is well known that CSMA/CA algorithms have an optimal transmission probability that depends on the network load, in terms of traffic generated and number of contending stations. For the case of saturated 802.11 WLANs, Bianchi [2] analytical 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, τ_t

$$\tau_t \approx \frac{1}{N}\sqrt{\frac{2T_e}{T_s}} \qquad (4)$$

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. Indeed, this is the motivation behind some adaptive algorithms (e.g. Idle Sense [7]) that equalize the amount of time wasted in collisions with the amount of time waiting in backoff decrements.

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. 3 that the optimal CW value indeed depends on the energy consumption of the WLAN interfaces. This relationship is what we investigate in the next section.

3.2 Energy optimization

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

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

This leads to the following

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

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

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

$$b = nE$$

$$c = -R$$

If we now define α and β as follows

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

Then we have the following for the computation of τ_e:

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

That can be approximated as follows

$$\tau_e \approx \frac{1}{n}\sqrt{\frac{2}{\beta}} \approx \frac{1}{n}\sqrt{\frac{2\rho_e T_e}{\rho_r T_s}} \tag{5}$$

Note that, if we divide (4) by (5), we have that 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.,

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

a relation that we analyze in the next section.

4 Energy Efficiency vs. Throughput Maximization

We first compare the resulting configuration obtained when maximizing throughput and when maximizing energy efficiency. To this end, in Fig. 4 we show the resulting CW configuration for each maximization variable, for the three considered interfaces of Table 1 and an increasing number of stations N. From the figure is obvious to see that, while the throughput maximization provides the same CW for a given number of stations, the optimal CW for energy efficiency depends quite noticeably on the power characteristics of the WLAN interface. It can be seen that, the larger the ρ_r/ρ_e ratio, the larger the CW. This could be expected from the results of Fig. 5, as collisions have a larger cost and therefore it is more efficient to spend more time on the backoff, instead of taking the risk of transmitting and suffering from a no-success but energy-consuming collision.

We next compare the performance of both approaches, both in terms of energy efficiency and in terms of throughput, to gain further in the behavior of the WLAN under the different criteria. Results for each approach, as well as for the standard recommended values (DCF), are provided in Figs. 5 and 6, and can be summarized as follows:
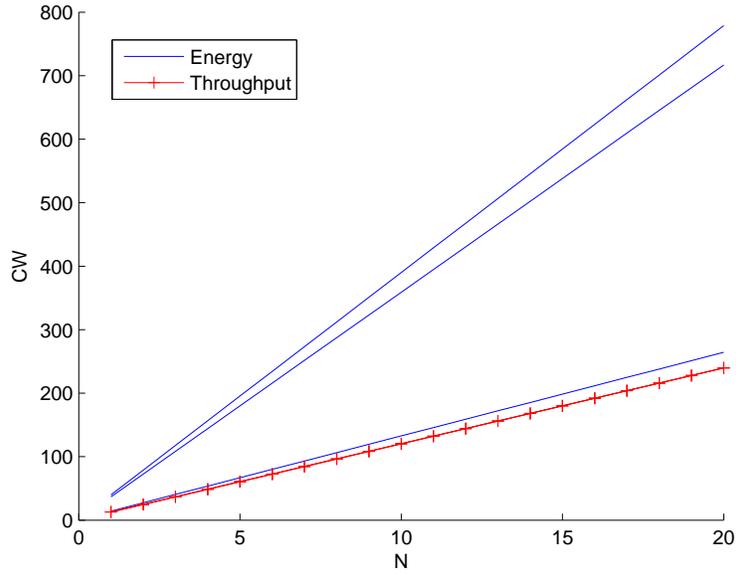
Fig. 4. Resulting *CW* configuration from each approach. The energy curves from top to bottom show the results for energy profile (C), (B), and (A)
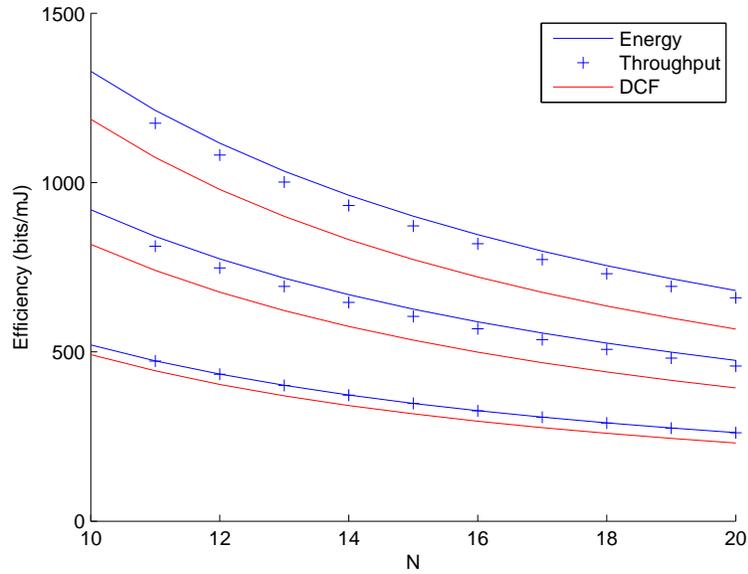

Fig. 5. Energy efficiency of each approach. The arrays of curves from top to bottom show the results for energy profile (B), (C), and (A)
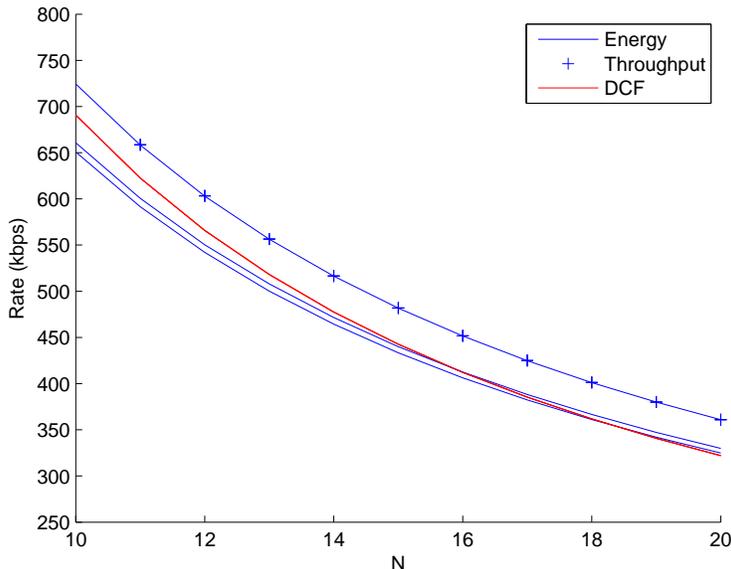
Fig. 6. Throughput performance of each approach. The energy curves from top to bottom show the results for energy profile (A), (C), and (B)

– Considering energy efficiency, despite both throughput and energy optimizing approaches substantially outperform the DCF default configuration, the maximum efficiency approach provides the larger values of bits per Joule. As expected from the results of Fig. 4, the larger the ρ_r/ρ_e ratio, the larger the differences in performance between τ_e and τ_t

– Considering throughput performance, it is clear that τ_t provides the largest values, as expected. It is quite remarkable, on the other hand, that while for one case the energy consumption provides almost the same results (this will happen as long as $\sqrt{\rho_r/\rho_e} \approx 1$), for the other two cases there is a price to pay. Indeed, the throughput for these two interfaces is smaller than the one provided by DCF for $N \leq 17$. However, this slightly smaller throughput is obtained with a different CW value that results in quite different values of energy spent in collisions and backoff counter decrements.

Therefore, 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 [3], 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, and there is a price to pay in throughput to achieve the most efficient behavior.

5 Conclusions

Greening the communication protocols is recognized as a primary design goal of future global network infrastructures. This paper presents a three-fold contribution on this field. First, it provides an approximate analytical model for the energy consumption of IEEE 802.11 LANs. Second, it defines an optimal configuration strategy that minimizes energy consumption for within such networks. Eventually, it provides a comparison of energy minimization against throughput optimization, this way assessing the price to pay. Our future work will focus on experimental analysis and measurements on the field.

6 Acknowledgments

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. 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.
3. 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.
4. J.-C. Chen and K.-W. Cheng. Edca/ca: Enhancement of ieee 802.11e edca by contention adaption for energy efficiency. *Wireless Communications, IEEE Transactions on*, 7(8):2866–2870, August 2008.
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. L. Feeney and M. Nilsson. Investigating the energy consumption of a wireless network interface in an ad hoc networking environment. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 3, pages 1548–1557 vol.3, 2001.
7. 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.*, 35(4):121–132, 2005.
8. 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*, 53(4):1235–1246, July 2004.

9. X. Wang, J. Yin, and D. P. Agrawal. Analysis and optimization of the energy efficiency in the 802.11 dcf. *Mob. Netw. Appl.*, 11(2):279–286, 2006.
10. H. Wu, Y. Peng, K. Long, S. Cheng, and J. Ma. Performance of reliable transport protocol over IEEE 802.11 wireles LAN: Analysis and enhancement. In *Proceedings of IEEE INFOCOM'02*, June 2002.