Computer Communications 000 (2017) 1–11

Contents lists available at ScienceDirect

Computer Communications

journal homepage: www.elsevier.com/locate/comcom

On the energy efficiency of rate and transmission power control in 802.11

Iñaki Ucar[a,*], Carlos Donato[b], Pablo Serrano[a], Andres Garcia-Saavedra[c], Arturo Azcorra[a,b], Albert Banchs[a,b]

[a] Universidad Carlos III de Madrid, Leganés, 28911, Spain
[b] IMDEA Networks Institute, Leganés, 28918, Spain
[c] NEC Labs Europe, Heidelberg, 69115, Germany

ARTICLE INFO

Article history:
Available online xxx

Keywords:
WLAN
802.11
Rate adaptation
Transmission power control
Energy efficiency

ABSTRACT

Rate adaptation and transmission power control in 802.11 WLANs have received a lot of attention from the research community, with most of the proposals aiming at maximising throughput based on network conditions. Considering energy consumption, an implicit assumption is that optimality in throughput implies optimality in energy efficiency, but this assumption has been recently put into question. In this paper, we address via analysis, simulation and experimentation the relation between throughput performance and energy efficiency in multi-rate 802.11 scenarios. We demonstrate the trade-off between these performance figures, confirming that they may not be simultaneously optimised, and analyse their sensitivity towards the energy consumption parameters of the device. We analyse this trade-off in existing rate adaptation with transmission power control algorithms, and discuss how to design novel schemes taking energy consumption into account.

1. Introduction

In recent years, along with the growth in mobile data applications and the corresponding traffic volume demand, we have witnessed an increased attention towards “green operation” of networks, which is required to support a sustainable growth of the communication infrastructures. For the case of wireless communications, there is the added motivation of a limited energy supply (i.e., batteries), which has triggered a relatively large amount of work on energy efficiency [1]. It turns out, though, that energy efficiency and performance do not necessarily come hand in hand, as some previous research has pointed out [2,3], and that a criterion may be required to set a proper balance between them.

This paper is devoted to the problem of rate adaptation (RA) and transmission power control (TPC) in 802.11 WLANs from the energy consumption’s perspective. RA algorithms are responsible for selecting the most appropriate modulation and coding scheme (MCS) to use, given an estimation of the link conditions, and have received a vast amount of attention from the research community (see e.g. [4,5] and references therein). In general, the challenge lies in distinguishing between those loses due to collisions and those due to poor radio conditions, because they should trigger different reactions. In addition, the performance figure to optimise is commonly the throughput or a related one such as, e.g., the time required to deliver a frame.

On the other hand, network densification is becoming a common tool to provide better coverage and capacity. However, densification brings new problems, especially for 802.11, given the limited amount of orthogonal channels available, which leads to performance and reliability issues due to RF interference. In consequence, some RA schemes also incorporate TPC, which tries to minimise the transmission power (TXP) with the purpose of reducing interference between nearby networks. As in the case of “vanilla” RA, the main performance figure to optimise is also throughput.

It is generally assumed that optimality in terms of throughput also implies optimality in terms of energy efficiency. However, some previous work [6,7] has shown that throughput maximisation does not result in energy efficiency maximisation, at least for 802.11n. However, we still lack a proper understanding of the causes behind this “non-duality”, as it may be caused by the specific design of the algorithms studied, the extra consumption caused by the complexity of MIMO techniques, or any other reason. In fact, it could be an inherent trade-off given by the power

* Corresponding author.
E-mail addresses: inaki.ucar@uc3m.es (I. Ucar), carlos.donato@imdea.org (C. Donato), pablo@it.uc3m.es (P. Serrano), andres.garcia.saavedra@neclab.eu (A. Garcia-Saavedra), azcorra@it.uc3m.es (A. Azcorra), banchs@it.uc3m.es (A. Banchs).

http://dx.doi.org/10.1016/j.comcom.2017.07.002
0140-3664/

consumption characteristics of 802.11 interfaces, and, if so, RA-TPC techniques should not be agnostic to this case.

This work tackles the latter question from a formal standpoint. A question which, to the best of the authors' knowledge, has never been addressed in the literature. For this purpose, and with the aim of isolating the variables of interest, we present a joint goodput (i.e., the throughtput delivered on top of 802.11) and energy consumption model for single 802.11 spatial streams in the absence of interfering traffic. Packet losses occur due to poor channel conditions and RA-TPC can tune only two variables: MCS and TXP.

Building on this model, we provide the following contributions: (*i*) we demonstrate through an extensive numerical evaluation that energy consumption and throughput performance are different optimisation objectives in 802.11, and not only an effect of MIMO or certain algorithms' suboptimalities; (*ii*) we analyse the relative impact of each energy consumption component on the resulting performance of RA-TPC, which serves to identify the critical factors to consider for the design of RA-TPC algorithms; (*iii*) we experimentally validate our numerical results; and (*iv*) we assess the performance of several representative RA-TPC algorithms from the energy consumption's perspective.

The rest of this paper is organised as follows. In Section 2, we develop the theoretical framework: a joint goodput-energy model built around separate previous models. In Section 3, we provide a detailed analysis of the trade-off between energy efficiency and maximum goodput, including a discussion of the role of the different energy parameters involved. We support our numerical analysis with experimental results in Section 4. Section 5 explores the performance of RA-TPC algorithms from the energy consumption's perspective. Finally, Section 6 summarises the paper.

2. Joint goodput-energy model

In this section, we develop a joint goodput-energy model for a single 802.11 spatial stream and the absence of interfering traffic. It is based on previous studies about goodput and energy consumption of wireless devices. As stated in the introduction, the aim of this model is the isolation of the relevant variables (MCS and TXP) to let us delve in the relationship between goodput and energy consumption optimality in the absence of other effects such as collisions or MIMO.

Beyond this primary intent, it is worth noting that these assumptions conform with real-world scenarios in the scope of recent trends in the IEEE 802.11 standard development, namely, the amendments 11ac and 11ad, where device-to-device communications (mainly through beamforming and MU-MIMO) are of paramount importance.

2.1. Goodput model

We base our study on the work by Qiao et al. [8], which develops a robust goodput model that meets the established requirements. This model analyses the IEEE 802.11a Distributed Coordination Function (DCF) over the assumption of an AWGN (Additive White Gaussian Noise) channel without interfering traffic.

Let us briefly introduce the reader to the main concepts, essential to our analysis, of the goodput model by Qiao *et al.*. Given a packet of length l ready to be sent, a frame retry limit $n_{\max}$ and a set of channel conditions $\hat{s} = \{s_1, \ldots, s_{n_{\max}}\}$ and modulations $\hat{m} = \{m_1, \ldots, m_{n_{\max}}\}$ used during the potential transmission attempts, the expected effective goodput $\mathcal{G}$ is modelled as the ratio between the expected delivered data payload and the expected transmission time as follows:

$$\mathcal{G}(l, \hat{s}, \hat{m}) = \frac{\mathbb{E}[\text{data}]}{\mathbb{E}[\mathcal{D}_{\text{data}}]} = \frac{\Pr[\text{succ} \mid l, \hat{s}, \hat{m}] \cdot l}{\mathbb{E}[\mathcal{D}_{\text{data}}]} \tag{1}$$

where $\Pr[\text{succ} \mid l, \hat{s}, \hat{m}]$ is the probability of successful transmission conditioned to $l, \hat{s}, \hat{m}$, given by Eq. (5) in [8]. This model is valid as long as the coherence time is equal or greater than a single retry, i.e., the channel condition s_i is constant.

The expected transmission time is defined as follows:

$$\mathbb{E}[\mathcal{D}_{\text{data}}] = \left(1 - \Pr[\text{succ} \mid l, \hat{s}, \hat{m}]\right) \cdot \mathcal{D}_{\text{fail}|l,\hat{s},\hat{m}} + \Pr[\text{succ} \mid l, \hat{s}, \hat{m}] \cdot \mathcal{D}_{\text{succ}|l,\hat{s},\hat{m}} \tag{2}$$

where

$$\mathcal{D}_{\text{succ}|l,\hat{s},\hat{m}} = \sum_{n=1}^{n_{\max}} \Pr[n\text{succ} \mid l, \hat{s}, \hat{m}] \cdot \Big\{ \sum_{i=2}^{n_{\max}} \Big[\overline{T}_{\text{bkoff}}(i) + T_{\text{data}}(l, m_i) + \overline{\mathcal{D}}_{\text{wait}}(i) \Big] + \overline{T}_{\text{bkoff}}(1) + T_{\text{data}}(l, m_1) + T_{\text{SIFS}} + T_{\text{ACK}}(m'_n) + T_{\text{DIFS}} \Big\} \tag{3}$$

is the average duration of a successful transmission and

$$\mathcal{D}_{\text{fail}|l,\hat{s},\hat{m}} = \sum_{i=1}^{n_{\max}} \Big[\overline{T}_{\text{bkoff}}(i) + T_{\text{data}}(l, m_i) + \overline{\mathcal{D}}_{\text{wait}}(i+1) \Big] \tag{4}$$

is the average time wasted during the $n_{\max}$ attempts when the transmission fails.

$\Pr[n\text{succ} \mid l, \hat{s}, \hat{m}]$ is the probability of successful transmission at the nth attempt conditioned to $l, \hat{s}, \hat{m}$, and $\overline{\mathcal{D}}_{\text{wait}}(i)$ is the average waiting time before the i-th attempt. Their expressions are given by Equations (7) and (8) in [8]. The transmission time (T_{data}), ACK time (T_{ACK}) and average backoff time ($\overline{T}_{\text{bkoff}}$) are given by Eq. (1)–(3) in [8]. Finally, T_{SIFS} and T_{DIFS} are 802.11a parameters, and they can be found also in Table 2 in [8].

2.2. Energy consumption model

The selected energy model is our previous work of [9], which has been further validated via ad-hoc circuitry and specialised hardware [10] and, to the best of our knowledge, stands as the most accurate energy model for 802.11 devices published so far, because it accounts not only the energy consumed by the wireless card, but the consumption of the whole device. While classical models focused on the wireless interface solely, this one demonstrates empirically that the energy consumed by the device itself cannot be neglected as a device-dependent constant contribution. Conversely, devices incur an energy cost derived from the frame processing, which may impact the relationship that we want to evaluate in this paper.

The energy model is a multilinear model articulated into three main components:

$$\overline{P}(\tau_i, \lambda_i) = \underbrace{\rho_{\text{id}} + \sum_{i \in \{\text{tx,rx}\}} \rho_i \tau_i}_{\text{classical model}} + \sum_{i \in \{\text{g,r}\}} \gamma_{\text{x}i} \lambda_i \tag{5}$$

where the first two addends correspond to the classical model and the third is the contribution described in [9]. These components are the following:

- A platform-specific baseline power consumption that accounts for the energy consumed just by the fact of being powered on, but with no network activity. This component is commonly referred to as *idle consumption*, ρ_{id}.
- A component that accounts for the energy consumed in transmission, which linearly grows with the airtime percentage τ_{tx}, i.e., $\overline{P_{\text{tx}}}(\tau_{\text{tx}}) = \rho_{\text{tx}} \tau_{\text{tx}}$. The slope ρ_{tx} depends linearly on the radio transmission parameters MCS and TXP.