Computer Networks 66 (2014) 82–93

Contents lists available at ScienceDirect

Computer Networks

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

Queueing systems to study the energy consumption of a campus WLAN

Marco Ajmone Marsan [a,b,*], Michela Meo [a]

[a] Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy
[b] IMDEA Networks Institute, Avenida del Mart Mediterraneo, 22, 28918 Leganes (Madrid), Spain

ARTICLE INFO

Keywords:
Energy efficiency
Wireless LANs
Queueing models

ABSTRACT

In this paper we exploit simple approximate queueing models to assess the effectiveness of the approaches that have been proposed to save energy in dense wireless local area networks (WLANs), based on the activation of access points (APs) according to the user demand. In particular, we look at a portion of a dense WLAN, where several APs are deployed to provide sufficient capacity to serve a large number of active users during peak traffic hours. To increase capacity, some APs are colocated and provide identical coverage; we say that these APs belong to the same *group*, and they serve users in the same area. The areas covered by different AP groups only partially overlap, so that some active users can only be served by a group of APs, but a fraction of active users can be served by more groups. Due to daily variations of the number of active users accessing the WLAN, some APs can be switched off to save energy when not all the capacity is needed. A real example of this setting is provided by a floor of one building of Politecnico di Torino in Italy, where a student library is located. The approximate analytical models indicate that the energy saving achievable with the proposed approaches is quite substantial, over 40% if at least one AP for each group is always kept on, even with no traffic, to be ready to accept incoming users, and it grows to almost 60% if all APs can be switched off at night, using a separate technology to activate an AP when the first user requests association in the morning.

1. Introduction

It first happened to our phones: initially we were accustomed to walk to them to answer a call, and to stay there for the entire conversation. We even found it natural not to be reachable while driving. Now, being tied to a cable for just a short call, makes us feel like dogs on leash, and we use our phones when reaching a friend's home, rather than ringing the door bell. Then it happened to our connections to the Internet: we are now addicted to reading email wherever we are: on a plane that just landed, in the elevator, in bed, stopping at a red light ... using our smartphones, tablets, laptops, phablets, and whatnot. This mutation from wired to wireless happened in spite of the unavoidable loss in performance, inherent in the poor characteristics of radio transmission, just because of convenience. We use our mobile(s) at work, even if the fixed phone is on the desk, and we prefer our laptop(s) to our desktop, even if downloads are slower. To mitigate this performance loss, large organizations have been consistently increasing the available bandwidth for mobile data users, by deploying more and more access points (APs) in their wireless local area networks (WLANs). This attitude has led to today's dense, centrally managed WLANs, where very large numbers of APs are installed, comparable to the

* Corresponding author at: Department of Electronics and Telecommunications, Politecnico di Torino, Corso Duca degli Abruzzi 24, 10129 Torino, Italy. Tel.: +39 011 0904032.
E-mail address: ajmone@polito.it (M. Ajmone Marsan).

http://dx.doi.org/10.1016/j.comnet.2014.03.012
1389-1286/

maximum number of simultaneously active users (some industrial campuses report over 10 thousand APs).

Even if the individual power consumption of APs is small, the fact that their number is very high implies a remarkable energy consumption and cost (10 W per AP, times 10 thousand APs, times 8760 h in a year, results in 876 MWh a year, with a cost of around 200 thousand euro, using the – admittedly high – kWh price in Italy). Most disturbing is the fact that the majority of this enegy is wasted, since the WLAN capacity that APs provide is not necessary 24/7; rather, capacity should be modulated according to the number of active users, or the traffic they generate, both of which vary widely over hours, days, and weeks; actually, the maximum WLAN capacity is necessary only at peak usage periods, for a limited time.

These considerations, coupled with the growing attention of the networking research community to energy efficiency, has led several research groups to investigate the energy saving which is possible with a smart management of WLAN resources, in particular by switching on and off APs according to the required capacity.

The first work in this field is [5], where the authors suggest that, in dense WLANs, APs can be clustered, based on their Euclidean distance. When the WLAN traffic (or the number of users) is low, only one AP in each cluster, the cluster-head, is switched on. When the traffic/number of users increases, additional APs can be switched on, to provide adequate capacity. Note that keeping track of the number of users associated with the APs in the cluster is easier (and more stable) than measuring the traffic load, so this quantity is preferred as a control variable in the dense WLAN central controller that turns on/off the cluster APs. The estimated energy saving can be 20–50% in less dense scenarios, whereas in more dense WLANs it can grow to 50–80%. An improvement of the AP clustering scheme was proposed in [6], by using the number and signal strength of the received beacons. This paper also suggested that a suitable metric to estimate user demand and drive the provided capacity can be the percentage of time the channel is busy due to transmission and interframe spacing.

A similar approach was proposed in [12] to reduce the number of switched-on APs, suggesting that it can be even possible to switch off all the APs in a cluster, provided the area covered by the cluster can be served by neighboring AP clusters. In this case the user demand is estimated from the number of users associated with APs, and energy savings are quantified in about 60%.

A first analytical model for the estimation of the energy gain achievable by modulating the number of switched-on APs as a function of the user demand was proposed in [2], considering just one AP cluster. Assuming as input a measured traffic trace, the analytical predictions indicate that the energy saving can be of the order of 40%.

In [10], the authors developed an ILP (Integer Linear Program) optimization model to adapt the number of switched-on APs to the number and the position of active users, achieving up to 63% power saving. In a subsequent paper [9], the same authors devised a heuristic approach which reduces the problem complexity, but also the achievable power saving.

A WLAN power saving scheme based on the maximum coverage problem was proposed in [3]. The proposed algorithm runs on a central controller, which collects the number of users associated with each AP and their data rates, and switches APs on and off dynamically, while maintaining coverage and guaranteeing user performance. Power saving of about 80% is reported at the expense of frequent user associations and significant delays.

The authors of [4] suggest that a drastic reduction of the density of APs in WLANs is possible, provided that the few APs remaining active can provide the coverage required to discover the presence of users. The key point here is that the detection of the user presence is possible with very limited AP coverage, and additional APs can then be switched on, to provide adequate capacity to users. Numerical evaluations show that up to 98% of APs can be switched off with this approach.

Other approaches exist, which are based on the presence of a secondary channel, that can be used to alert switched-off APs of the user presence. For example, the approaches proposed in [13–15] assume that all inactive APs can be switched off regardless of coverage. To react to user presence, an auxiliary low-power channel is available, that wakes the APs when users require access. The authors of [8] assumed that users are connected to a cellular network, which is capable of requesting the switch-on of a WLAN AP in the vicinity of the user. In [16] the presence of a secondary Bluetooth interface in both APs and mobile stations is assumed, and the Bluetooth interface can be used to request the activation of the WLAN APs.

In a previous paper [1] we considered a portion of a dense WLAN, where clusters of APs partially overlap in coverage, so that some active users can only be served by the APs of one cluster, but a fraction of active users can be served by APs in more than one cluster. We investigated the energy-performance trade-off with a model based on coupled queues: each queue represents a cluster of colocated APs, and customers arriving at the system might be served by one or several queues, depending on their position. We also proposed approximations based on single-queue and two-queue analysis, which provide fairly accurate performance estimates.

In this paper we expand on those single-queue approximate models, first showing that measurements of WLAN AP coverage validate the setting we consider, and that measurements of user activity justify the interest in energy saving approaches for a university campus WLAN. Then, we use simple queueing models, with and without setup times, to compare the cases in which the AP wake-up can or cannot exploit the presence of a secondary channel for the detection of user presence.

Queueing systems [7] have been for many years one of the most popular and versatile tools for the quantitative analysis and design of networks. The types of metrics traditionally derived from queueing system models are extremely diverse, ranging from very simple indicators, like the average number of waiting customers, or the average time in the queue, to more elaborate parameters, like, for example, server busy and idle period duration distributions. Of course, the metrics of interest depend on the type of system and on the objective of the quantitative analysis. Very