# A Simple Analytical Model for the Energy-Efficient Activation of Access Points in Dense WLANs *

Marco Ajmone Marsan
1) Dipartimento di Elettronica
Politecnico di Torino - Italy
2) IMDEA Networks - Spain
marco.ajmone@polito.it

Luca Chiaraviglio
Dipartimento di Elettronica
Politecnico di Torino - Italy
luca.chiaraviglio@polito.it

Delia Ciullo
Dipartimento di Elettronica
Politecnico di Torino - Italy
delia.ciullo@polito.it

Michela Meo
Dipartimento di Elettronica
Politecnico di Torino - Italy
michela.meo@polito.it

## ABSTRACT
Energy efficient networks are becoming a hot research topic, and the networking community is increasingly devoting its attention to the identification of approaches to save energy in the networks of today. However, the networks of tomorrow will require built-in energy efficiency capabilities, so that new design techniques based on network models that account for energy efficiency are called for.

One of the simplest approaches to obtain energy efficiency is based on the activation of network resources on demand, thus avoiding to always power on all the resources that are necessary to serve users during peak traffic periods.

In this paper we both present a simple analytical model to determine the effectiveness of policies that activate APs (Access Points) in dense WLANs (Wireless LANs) according to the actual user demands, and quantify the performance that is achieved by such policies in terms of energy savings and QoS (Quality of Service). Numerical results show that, in the configurations that we studied, energy savings up to 87% are possible during low traffic periods, with hardly any sacrifice in QoS.

## Categories and Subject Descriptors
C.2 [**Computer Communication Networks**]: Network Architecture and Design, Network Operations

## 1. INTRODUCTION
The number of wireless accesses to the Internet is growing at a rate close to 100% a year, much faster than the number of wired accesses, so that some forecasts predict that in the future the majority of accesses to the Internet will be wireless. Today, the most common technologies for wireless Internet access are 3G cellular networks, and IEEE 802.11 WLANs (Wireless Local Area Networks), commonly termed Wi-Fi. The numbers relating to cellular networks are staggering, with about 4.5 billions users worldwide, served by over 4 millions cells. Similarly impressive are the numbers of Wi-Fi hotspots, with some large corporate campuses totalling thousands or even tens of thousands of access points (APs).

The energy consumed and emitted by such a huge number of wireless access points has raised the concern of researchers, specially because such energy is largely wasted in the periods of time when the number of users served by the wireless access network is low. Several groups recently started investigating approaches to improve the energy efficiency of wireless access networks. We previously studied approaches for the energy-efficient management of cellular access networks [1, 2, 3, 4], while the issues relating to the reduction of the energy consumption in WLANs were investigated by Papagiannaki et al. [5, 6]. The key idea behind those papers is that wireless access networks are used at capacity only for limited periods of time, and significant energy savings can be achieved by switching off a fraction of the access points during periods of low traffic.

In particular, in [6], the authors propose a resource-on-demand (RoD) policy to dynamically power on and off WLAN APs, based on the volume and location of user demand. Through an experimental setup, they show that huge energy savings (up to 54%) are possible in the examined configurations.

In this paper, we consider the same context of [6], and we develop a very simple analytical model that allows the investigation of the effectiveness of different policies for the AP switch-off, under different traffic loads. To the best of our knowledge, this is *the first analytical model for the study of the energy efficiency of large WLANs*. Our analytical model is suited to dense, centrally managed WLANs, where the number of APs is large (so that *clusters* of APs providing coverage over the same area can be identified), and where the centralization of management provides a powerful

*This work was supported by the Italian Ministry of University and Research through the PRIN project EFFICIENT.

instrument for the implementation of effective energy saving approaches. The case we consider is significant, since, although somewhat idealized, it is typical of today's large corporate WLANs. As an example, consider that at the time of writing this paper, in the office of one of the authors at Politecnico di Torino, the signals of 20 APs were received by the inSSIDer Wi-Fi scanner [7], 14 of which belonging to the institutional campus network.

The simplicity of the model we describe in this paper is a very valuable characteristic, since it will allow extensions to more complex environments, better reflecting the diversity of the topologies and of the traffic in large corporate WLANs.

The main improvement achieved by our approach with respect to [6] is that, while the experimental setups of [6] provide very credible and accurate indications of what can be obtained in the considered settings, they only offer general indications about the effectiveness of the approach in other contexts. Instead, our model is quite general, and easily adaptable to a variety of different scenarios, thus allowing the prediction of the energy saving achieved by the network operator, and of the quality of service (QoS) perceived by end users, in all WLANs that match the modeling assumptions.

The rest of this paper is organized as follows. Section 2 describes the class of WLANs that we consider, and discusses the main assumptions adopted in the model development, as well as the policies that can be used to switch off APs. Section 3 illustrates the analytical model of one cluster of APs in the WLAN operating under two different RoD policies, and shows how performance metrics are computed from the model solution. Section 4 discusses some significant numerical results for energy saving and QoS in an example AP cluster. In Section 5, to validate our approach, we compare the predictions of our analytical model to the results presented in [6]. Section 6 provides a short overview of related works. Finally, conclusions and an outlook on future extensions are presented in Section 7.

## 2. THE SYSTEM

Corporate WLANs of today are *dense*, which means that they comprise hundreds to thousands, and in some cases even tens of thousands access points (APs). These APs on the one side offer wireless connectivity to mobile terminals (notebooks, netbooks, PDAs, smartphones, etc.), and on the other side connect to a wired corporate LAN (normally Ethernet) made of fibers and UTP cables, switches and routers. In such dense WLANs, APs can be grouped into *clusters* [5, 6], where APs belonging to the same cluster provide overlapping coverage. Such overlapping APs are deployed to fulfill the end user requirements of high bandwidth and reliability. Several AP vendors offer centralized solutions for the management of such dense WLANs, simplifying the structure of individual APs, and concentrating in a central controller the management capabilities for the whole set of APs. Examples are the products of Cisco, Trapeze Networks, Meru Networks, Aruba Networks, Symbol Technologies by Motorola, 4ipnet. Other companies, such as Connect 802 or Develcon, offer software for the centralized management of multi-vendor WLANs.

The combination of centralized management and clusters of APs suggests a very natural approach to the reduction of the WLAN energy consumption. Indeed, consider that every AP consumes about 10 W in the *on* mode, almost 90 kWh a year, which for a WLAN with 10 thousand APs [8] means almost 1 GWh a year, for a cost of the order of 150 thousand euros or dollars with current tariffs, while only a minimal amount of energy is needed by the APs in the *off* (or *sleep*) mode. Since the bandwidth and reliability of all APs in a cluster is necessary only at peak usage times, the conclusion that some APs should be turned off in periods of low traffic, so that energy can be saved, is quite natural. This amounts to what in [6] is called a RoD (resource on demand) policy. Note that, when an AP is switched off, some of the end user terminals which were connected to it may need to be transferred to another AP, which remains on, through a reassociation procedure, in a seamless way. Re-associating users is possible today by disconnecting a user, which must then identify a new AP, but in the future, explicit reassociation commands will be possible using the upcoming IEEE 802.11v standard [9], which foresees that APs can explicitly ask users to re-associate to another AP. Also, if necessary, the transmitted power of some of the APs that remain on can be increased, to provide full coverage over the area that was served by the APs that were turned off. Furthermore, users within the cluster can maintain the same data-rates, regardless the AP of the cluster they are connected to.

However, if each AP manages itself individually, it is rather difficult to distribute among APs the information about the present number of active users and traffic load, so as to decide with a distributed algorithm if some APs can be switched off, which ones, and which terminals must re-associate with which AP. On the contrary, the presence of a single WLAN controller that manages all APs implies that the information about the global state of the WLAN is available at the controller, which can then implement appropriate algorithms to decide which APs are actually needed to achieve the desired QoS, and switch off the other APs, ordering the reassociation of those users that were connected to the APs which are going off.

RoD policies are thus quite a natural solution, and they can lead to huge energy savings. However, RoD policies must be designed with great care, not to degrade the end user experience. The three laws for an effective RoD policy can be spelled out as follows:

1. the WLAN coverage must not be reduced
2. the QoS offered to end users must not be degraded
3. the WLAN operations must be stable

The first law implies that APs can be switched off according to the actual WLAN utilization, as long as at least one AP remains active in all the area served by the WLAN, to provide service to possible new incoming users. The second law guarantees that end users receive service with a quality comparable to the one which is provided by a WLAN in which all APs are active. The third law has the effect of discouraging frequent reassociations of end user terminals to

APs. A more or less equivalent definition of the desirable qualities of a RoD policy can be found in [6].

In this paper, we develop an analytical model of the behavior of one cluster of APs, which operate according to specified RoD policies. The cluster is formed by a number of APs which are in close proximity of each other, so that the coverage they offer is practically equivalent. This means that any one of the APs can be switched on or off, according to the activity of end users.

We consider two simple policies for the AP switch-off and switch-on, which differ for the state of end users that triggers the policy. In the first case, the *association-based* policy, we base our policy on the number of users associated with APs in the cluster. In the second case, the *traffic-based* policy, we consider users that are not only associated, but are in addition generating traffic. Associations of end users are very simple to track by APs. Traffic activity is less trivial to determine, but could be for example tracked by means of the presence of open TCP connections, or by using a timeout since the last packet transmission. Of course, the stability of the association-based policy is higher than the traffic-based policy.

We first describe the association-based policy. Denote with $M$ the maximum number of users that can be associated with an AP, and with $T_h \leq M$ a threshold. The basic idea of the policy is that when the number of users associated with APs in the cluster is above $kT_h$, the number of active APs must be $k + 1$. In a dynamic scenario, where users frequently enter and leave the cluster, and associate with the cluster APs, this switch-off policy can lead to frequent AP switch-off and switch-on, and to frequent reassociations of customers, thus violating the third law of RoD. To avoid this problem, we introduce a hysteresis of amplitude $T_l$ in the switch-off procedure. In practice, the $(k + 1)$-st AP is switched on when the number of users in the cluster exceeds $kT_h$, but the number of active APs decreases to $k$ again, only when the number of users in the cluster becomes $\leq kT_h - T_l$.

As an example, Fig. 1 shows how the hysteresis works in our policy, with $T_h$ =3 and $T_l = 1$. At the beginning (left part of the figure), 2 APs in the cluster are on, with 3 users each; when a new user associates with the cluster, a third AP is switched on. The users then decrease to 6 again, but no AP is switched off due to the hysteresis; the number of APs that are powered on remains 3 as long as the number of associated users remains between 6 and 9 (extremes included); an AP is switched off only when the number of associated users decreases to 5.

The traffic-based policy works in a similar way, but considers only the users that generate traffic among those which are associated with APs in the cluster. Two thresholds, $C_h$ and $C_l$, are used also in this case. When the number of traffic-generating users associated with APs in the cluster is above $kC_h$, the number of active APs must be at least $k + 1$. On the contrary, when the number of traffic-generating users in the cluster decreases below $kC_h - C_l$, $k$ active APs are sufficient to provide service.

# 3. THE CLUSTER MODEL

To evaluate and compare the performance of the system under the two proposed RoD policies, we develop a continuous-time Markov chain (CTMC) model of a cluster of APs. Since the models for the two considered policies differ for only few modifications, we first describe in detail the model of the association-based policy; then, we just describe the modifications to be introduced in the model to adapt it to the traffic-based policy.

## 3.1 Model of the association-based policy

For the model development, we introduce the following assumptions.

- The number of users associated with an AP of the cluster can vary: users associate according to a (time varying) Poisson process with rate $\lambda_s$ (a reasonable assumption, according to the experimental data of [10]); they leave the cluster after an exponentially distributed time with mean $1/\mu_s$. Thanks to the centralized WLAN management, users evenly distribute over the APs of the cluster that are powered on.
- Associated users can be *idle*, when they do not generate traffic, or *active*, when they are generating traffic. An idle user becomes active after a time whose probability density function (pdf) is exponential with mean $1/\lambda_c$. The amount of traffic generated by an active user is distributed according to an exponential pdf with mean $1/\mu_c$ (although measured distributions can be heavy-tailed [11], we expect an exponential distribution to be sufficient for a preliminary assessment of the relative merits of the considered RoD policies).
- The AP bandwidth $B$ is evenly shared among all users associated with the AP. For the sake of simplicity, we neglect the fact that users at higher distance from the AP might perceive a larger packet loss probability or, by adapting their bit rate, might transmit at lower rate. We also neglect the fact that the total AP throughput slightly decreases with the number of active users. While these assumptions may be too simplistic in short time intervals, their adoption over longer time periods seems reasonable, in light of the fact that we are interested in a comparison of the effectiveness of different RoD policies, not in the exact determination of the performance perceived by each user.
- The total number of APs in the cluster is equal to $A$; no more than $M$ users can be associated with each AP, so that no more than $AM$ users are admitted in the cluster.

The above assumptions allow the description of the cluster dynamics with a simple continuous-time Markov chain model. Simplicity is an important and valuable characteristic of the model, since it will allow extensions, to consider more complex AP setups, and more realistic user behaviors.

The state of Markov chain $\{X(t), t \geq 0\}$ describing the cluster is given by vector $\bar{s}$ composed of three values, $\bar{s} =$

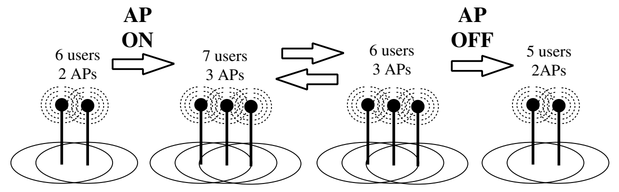


**Figure 1: Example of a hysteresis cycle, with $T_h$ =3 users per AP, and $T_l = 1$ user**

$(p, u, c)$, where $p$ is the number of APs that are on, $u$ is the number of users that are associated with any AP of the cluster, and $c$ is the number of associated users that are generating traffic.

As described in the previous section, the idea behind the association-based policy is that the number of powered on APs depends on the number of associated users: the smallest number of APs is kept on, constrained on all active APs having less than $T_h$ user associations. Users are evenly distributed among the APs in the cluster (thanks to the central management of the WLAN), so that, when $p$ APs are on and the number of users is equal to $pT_h$, $T_h$ users are associated with each AP. As soon as a new user associates with an AP in the cluster, a new AP is switched on. To avoid that the switch-on and switch-off frequency is too high, when $p$ APs are on, one of them can be switched off only when the number of users becomes $\leq pT_h - T_l$.

The reachable states in the MC compose the state space $\mathcal{S}^{(\mathrm{X})}$ and jointly satisfy the following relations:

$$\begin{cases} 1 \leq p \leq A \\ 0 \leq u \leq MA \\ 0 \leq c \leq u \\ (p-1)T_h - T_l < u \leq pT_h \end{cases} \tag{1}$$

The third constraint refers to the assumption that each user can be either idle, or active in generating traffic, while the last constraint is the core of the RoD policy.

Table 1 reports the possible CTMC transitions from the generic state $(p, u, c)$, and, for each transition, it indicates the destination state, the transition rate, and the condition over $(p, u, c)$ for the transition to be possible. When a user associates, it might trigger the switch-on of an AP; similarly, when a user leaves the cluster, an AP might be switched off (notice that a user leaving the cluster might also interrupt an active connection). Connections are started by associated users that are not active yet. The rate of a connection termination is proportional to $\mu_c$ according to a factor $B(\overline{s})$, which represents the amount of bandwidth that is available for the connection. By assuming that the bandwidth is fairly shared among the active users associated with an AP, and by neglecting the fact that the number of active users associated with an AP is discrete, we may say that in state $(p, u, c)$ the average number of connections per AP is $c/p$, and the bandwidth per connection is:

$$B(\overline{s}) = \min(B, Bp/c) \tag{2}$$

**Table 1: Association-based policy: CTMC transitions from state $(p, u, c)$**

| Action | Destination | Rate | Condition |
|---|---|---|---|
| **User associates** | | | |
| AP ON | $(p+1, u+1, c)$ | $\lambda_s$ | $u = T_h p \wedge p < A$ |
| AP – | $(p, u+1, c)$ | $\lambda_s$ | $u \neq T_h p \wedge u < AM$ |
| **Connection starts** | | | |
| AP – | $(p, u, c+1)$ | $(u-c)\lambda_c$ | $c < u$ |
| **Connection ends** | | | |
| AP – | $(p, u, c-1)$ | $cB(\overline{s})\mu_c$ | $c > 0$ |
| **Active user leaves** | | | |
| AP – | $(p, u-1, c-1)$ | $c\mu_s$ | $c>0 \wedge u \neq T_h(p-1)-T_l+1$ |
| AP OFF | $(p-1, u-1, c-1)$ | $c\mu_s$ | $c>0 \wedge u=T_h(p-1)-T_l+1 \wedge p>1$ |
| **Non-active user leaves** | | | |
| AP – | $(p, u-1, c)$ | $(u-c)\mu_s$ | $u>0 \wedge u \neq T_h(p-1)-T_l+1$ |
| AP OFF | $(p-1, u-1, c)$ | $(u-c)\mu_s$ | $u>0 \wedge u=T_h(p-1)-T_l+1 \wedge p>1$ |

**Table 2: Traffic-based policy: CTMC transitions from state $(p, u, c)$**

| Action | Destination | Rate | Condition |
|---|---|---|---|
| **User associates** | | | |
| AP – | $(p, u+1, c)$ | $\lambda_s$ | $u < AM$ |
| **Connection starts** | | | |
| AP – | $(p, u, c+1)$ | $(u-c)\lambda_c$ | $c < u \wedge c \neq pC_h$ |
| AP ON | $(p+1, u, c+1)$ | $(u-c)\lambda_c$ | $c < u \wedge c = pC_h \wedge p < A$ |
| **Connection ends** | | | |
| AP – | $(p, u, c-1)$ | $cB(\overline{s})\mu_c$ | $c>0 \wedge c \neq C_h(p-1)-C_l+1$ |
| AP OFF | $(p-1, u, c-1)$ | $cB(\overline{s})\mu_c$ | $c>0 \wedge c=C_h(p-1)-C_l+1 \wedge p>1$ |
| **Active user leaves** | | | |
| AP – | $(p, u-1, c-1)$ | $u\mu_s$ | $c > 0$ |
| **Non-active user leaves** | | | |
| AP OFF | $(p-1, u-1, c)$ | $(u-c)\mu_s$ | $u > 0 \wedge p > 1$ |

## 3.2 Model of the traffic-based policy

Under the traffic-based policy, APs are switched on and off, based on the amount of traffic that is carried by the cluster. When the number of connections per AP grows above a threshold $C_h$, a new AP is switched on; when the number of connections decreases below $pC_h - C_l$, an AP is switched off.

The model presented in the previous section can be easily modified to describe this policy. Let $\{Y(t), t \geq 0\}$ denote the CTMC of the traffic-based policy. The state is defined as before, but the actions that determine the AP switch-on and off are now decided by the variable $c$, that represents the number of active connections. The reachable states are different with respect to the previous case. In particular, the last constraint in (1) is replaced by:

$$(p-1)C_h - C_l < c \leq pC_h \tag{3}$$

Let the state space for this model be denoted by $\mathcal{S}^{(\mathrm{Y})}$.

Table 2 reports the possible CTMC transitions from the generic state $(p, u, c)$, considering this policy.

## 3.3 Performance indices

To investigate and compare the performance of our RoD policies, we evaluate a set of parameters related to the power consumption of the APs, and to the Quality of Service (QoS) perceived by end users, based on the steady-state probabilities of the CTMC models. Let the steady-state probability of state $\overline{s}$ be denoted by $\pi_s^{(X)}$, and $\pi_s^{(Y)}$, respectively, for the models of the association-based and traffic-based policies.

Let us first focus on the *switching rate*, i.e., the average number of times an AP is switched on (or off) in the time unit. At steady state, the average number of on/off and off/on transitions per unit time must be equal; let $R^{(X)}$ and $R^{(Y)}$ denote this value for the two policies. We can derive $R^{(X)}$ from:

$$R^{(X)} = \sum_{\overline{s} \in S_{on}^{(X)}} \pi_s^{(X)} \lambda_s \quad (4)$$

with $S_{on}^{(X)}$ the set of states from which an AP might be switched on,

$$S_{on}^{(X)} = \{(p, u, c) | u = T_h p \wedge p < A\} \subset \mathcal{S}^{(X)} \quad (5)$$

Similarly, for the traffic-based policy we have:

$$R^{(Y)} = \sum_{\overline{s} \in S_{on}^{(Y)}} \pi_s^{(Y)} (u - c) \lambda_c \quad (6)$$

with $S_{on}^{(Y)}$ given by:

$$S_{on}^{(Y)} = \{(p, u, c) | c = C_h p \wedge c < u \wedge p < A\} \subset \mathcal{S}^{(Y)} \quad (7)$$

We now compute the *average bandwidth per connection*. This metric is directly related to the QoS perceived by the end users. For both policies, it is computed as:

$$\bar{B} = \frac{\sum_{\overline{s} \in \mathcal{S}'} B(\overline{s}) \pi_s}{\sum_{\overline{s} \in \mathcal{S}'} \pi_s} \quad (8)$$

where $S'$ is the subset of $S$ where $c > 0$.

Let $P_A$ be the power consumption of a cluster that always keeps on all its APs (the always-on policy):

$$P_A = AG \quad (9)$$

where $G$ is the power consumption of an AP, and $A$ is the total number of APs in the cluster. The power consumption of our RoD policies is given by:

$$P = G \sum_{\overline{s} \in \mathcal{S}} p \pi_s \quad (10)$$

that can be specialized for both the policies.

We thus define the *percentage power saving* as:

$$PS = 100 \frac{P_A - P}{P_A} \quad (11)$$

# 4. NUMERICAL RESULTS

In this section we discuss numerical results, computed with the models presented in the previous section. We use realistic model parameters, and we show that the power savings achieved by the proposed RoD policies can be huge in a dense WLANs scenario typical of today's corporate environments.

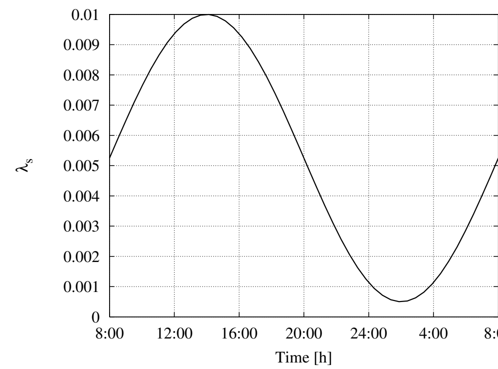


**Figure 2: Association rate of users during a 24 h period**

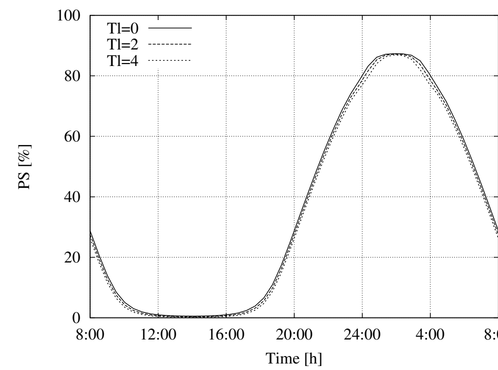


**Figure 3: Association-based policy: Power saving for different values of the hysteresis $T_l$ during a 24 h period**

For both policies, our model is applied to a 1-day period, assuming that the behavior of the association rate $\lambda_s$ versus time follows a typical periodic daily behavior, which we represent with a simple sinusoidal curve, as shown in Fig. 2 (which is not far from the observations of [11]). The periodicity of associations is quite intuitive, since users mostly come to office in the morning and turn on their terminals, and leave in the evening. We take the ratio between the maximum and the minimum of the association rate to be about 20. Indeed, the maximum value for $\lambda_s$ is 0.01 associations/s, i.e., one association every 100 s, while the minimum value is close to one association every half an hour.

We consider one cluster composed by $A = 8$ APs. Each AP can serve up to $M = 10$ users, therefore, the maximum number of associated users within the cluster is $MA = 80$. We set the threshold for the association-based policy to the maximum possible value: $T_h = M = 10$. For the traffic-based policy, we set the threshold to $C_h = 4$. The other system parameters were set to: $\mu_s = 0.0001\ s^{-1}$, $\mu_c = 0.005\ s^{-1}$, $\lambda_c = 0.0008\ s^{-1}$.

Note that, to obtain time-dependent results, we generate the steady-state solution of our model for parameters evaluated every 30 minutes intervals.

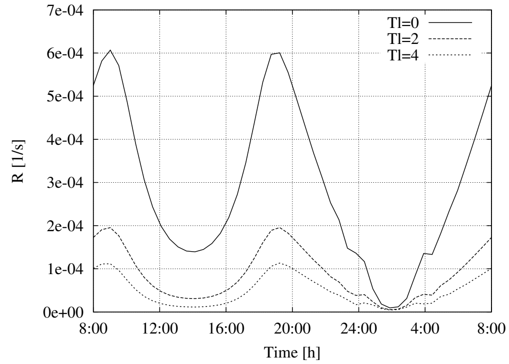


**Figure 4: Association-based policy: AP switch-off rate for different values of the hysteresis $T_l$ during a 24 h period**

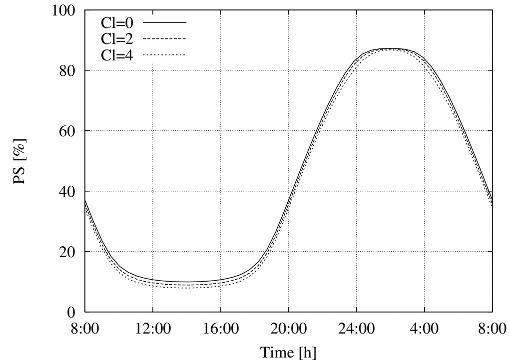


**Figure 5: Traffic-based policy: Power saving for different values of the hysteresis $C_l$ during a 24 h period**

## 4.1 Association-based policy

We start by considering the association-based policy, and we evaluate performance indexes for different values of the hysteresis $T_l$.

We compute the power saving of our RoD policy with respect to an always-on policy. Similar to [6], we consider that each AP consumes about $G$ =10 W. The power saving is reported in Fig. 3 for different values of hysteresis, specifically $T_l \in \{0, 2, 4\}$. The power saving reaches 87% during the low traffic period, and it is almost independent on $T_l$ (the decrease with $T_l$ is negligible). The value 87% corresponds to periods in which only one AP is powered on in the whole cluster, so that a fraction equal to 7/8 of the energy can be saved. On the contrary, no energy is saved during peak traffic periods (as expected), since all APs must be on at those times.

To dig further into the impact of the hysteresis, we evaluate the AP switch-off rate, $R$. Since $R$ is equal to both the on/off and off/on transition rates (that are, on average, the same, because of equilibrium), the total switching rate is actually $2R$. Fig. 4 shows $R$ during the whole day, for different values of $T_l$. The values of $R$ are small when traffic is high: in this case the number of associated users is so high that all the APs in the cluster are needed, and are almost continuously powered on. Consequently, no energy is saved at peak traffic, coherently with what we observed in Fig. 3. When traffic is low, transitions are also rare, since very few users request association to the APs in the cluster, and one AP only is sufficient to provide them with connectivity, and to guarantee coverage, as required by the 1st law of RoD policies. The intermediate values between the maximum and the minimum association rate correspond to the highest values of $R$: in the worst case, a switching rate of 0.0006 corresponds to about one switch-off every 27 minutes, on average, that, considering both the switch-on and -off transitions, corresponds to a transition every quarter of an hour. The hysteresis can drastically reduce the switching rate: $T_l = 2$ is already sufficient to reduce the maximum value of $R$ by a factor 3, and $T_l = 4$ reduces the maximum value by a factor close to 6, i.e., a switching transition every 3 hours. This value is surely adequate to satisfy the 3rd law of RoD policies.

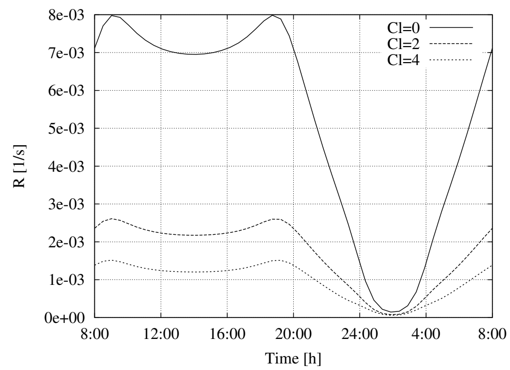


**Figure 6: Traffic-based policy: AP switch-off rate for different values of the hysteresis $C_l$ during a 24 h period**

## 4.2 Traffic-based policy

For the analysis of the traffic-based RoD policy, we use the same scenarios considered before, and set the threshold of the policy to $C_h = 4$, so that an additional AP has to be powered on when the number of active users $c$ is greater than $C_h \cdot p$, where $p$ is the number of APs in the cluster that are powered on.

Also in this case, we compute the power saving of our policy with respect to an always-on policy, for different values of the hysteresis $C_l$. Fig. 5 shows the results. With the traffic profile we have used (see Fig. 2), the minimum power saving, that occurs at the peak traffic hour, is about 10%. Notice that with the association-based RoD policy, no saving could be achieved at the peak traffic hour; this shows that, as expected, the traffic-based policy better adapts the number of active APs to the actual traffic demand. The cost that must be paid for this better performance in terms of power saving is a higher switch-off (switch-on) rate $R$. As shown in Fig. 6, that reports the values of $R$ for different values of $C_l$, the rate is one order of magnitude larger in this case than with the association-based RoD policy: in the worst case, the system performs, on average, a transition per minute, which might translate into an excessive system instability, thus violating the 3rd law of RoD policies. Again, the hysteresis can significantly improve the situation.

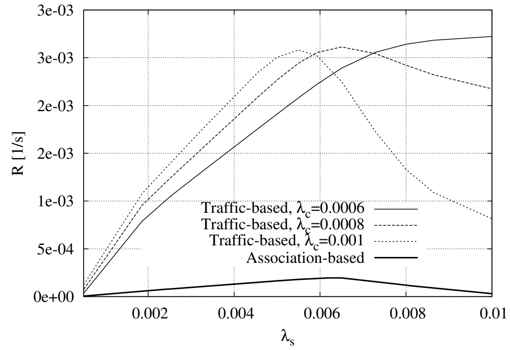


**Figure 7: AP switch-off rate versus the user association rate $\lambda_s$ for different values of the connection generation rate $\lambda_c$; $T_l = 2$, $C_l = 2$**

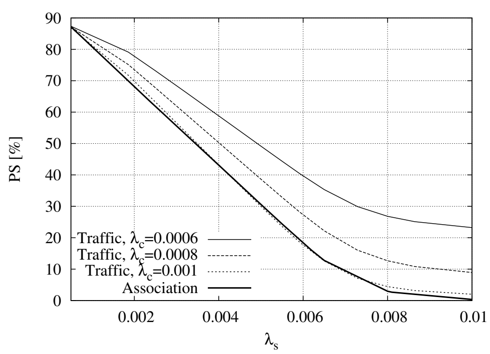


**Figure 8: Power Saving versus the association rate $\lambda_s$ for different values of the connection generation rate $\lambda_c$ ($C_l = 2$)**

## 4.3 Comparison of the two RoD policies

To compare the two policies, we first investigate the impact of traffic parameters on the system energy consumption and QoS. We consider different scenarios, obtained by varying the association rate $\lambda_s$ (which is equivalent to considering different times during the day) and the connection generation rate $\lambda_c$.

Fig. 7 shows the switch-off rate $R$ versus $\lambda_s$, for different values of $\lambda_c$, for both the considered RoD policies. Notice that the association-based policy is insensitive to $\lambda_c$; thus, one curve only is plotted for this policy. Initially, the switch-off rate increases with $\lambda_s$, i.e. the higher the user arrival rate is, the more dynamic the policy is. Then, after reaching a maximum value, the number of transitions decreases again, corresponding to high traffic conditions in which all the APs are powered on, and only a few transitions occur. This happens for both policies, but the switch-off rate of the association-based policy is one order of magnitude smaller than the one of the traffic-based policy. For the traffic-based policy, the higher the connection rate is, the more visible this phenomenon is.

The percentage power saving PS versus $\lambda_s$ and $\lambda_c$ is reported in Fig. 8, that shows that the traffic-based policy

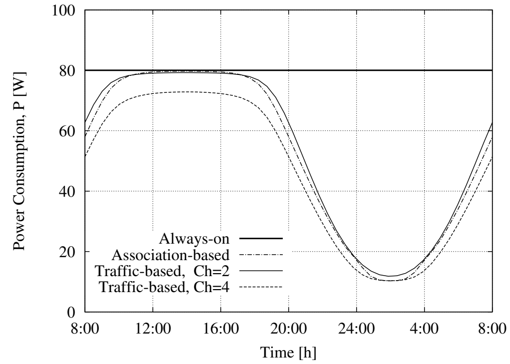


**Figure 9: Comparison of the RoD policies: Power consumption during a 24 h period**

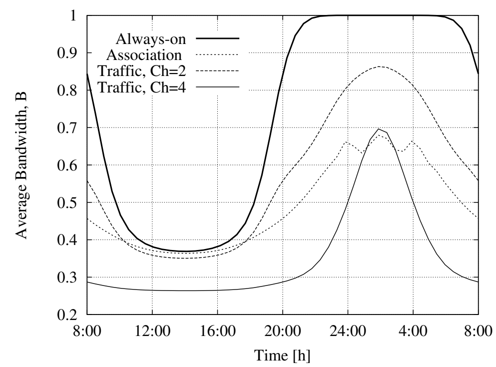


**Figure 10: Comparison of the RoD policies: Average bandwidth per active user during a 24 h period**

allows larger energy savings, especially when the traffic per associated user is low. In this case, the traffic-based policy can better follow the actual demand of capacity, as expected, at the price of a lower stability in the WLAN operation.

The power consumptions of the association-based, traffic-based, and always-on policies are reported in Fig. 9 for the usual sinusoidal pattern of the association rate (see Fig. 2); for the traffic-based policy we consider two values of $C_h$, $C_h = \{2, 4\}$. As previously mentioned, the power consumption of the always-on policy is constant, and equal to 80 W, since all the 8 APs of the cluster are always powered on, and consume $G = 10$ W each. Both our RoD policies lead to limited or no energy saving at peak traffic times; the saving for the traffic-based policy with $C_h = 4$ is achieved by exploiting fluctuations of traffic that allow the switch-off of one AP. A substantial saving (87%, as we already observed, since one AP only is powered on, rather than 8), during low traffic periods can be achieved by both the proposed RoD policies. Again, the traffic-based policy with $C_h = 4$ achieves the largest energy saving.

The always-on policy consumes 1.92 kWh per day, therefore, the energy consumption in one year is: $E_{year}^{TOT} \simeq 701$ kWh. Table 3 reports the power consumption and the consequent saving, for the two RoD policies, considering three different values of hysteresis $T_l$ and $C_l$ and, again, $T_h = 10$ and $C_h =$

Table 3: Energy consumption and energy saving per cluster per year

| Policy | $E_{year}^{SM}$ [kWh/year] | $ES$ [%] |
|---|---|---|
| Association | | |
| $T_l = 0$ | 447 | 36.2 |
| $T_l = 2$ | 453 | 35.3 |
| $T_l = 4$ | 460 | 34.4 |
| Traffic | | |
| $C_l = 0$ | 398 | 43.2 |
| $C_l = 2$ | 405 | 42.2 |
| $C_l = 4$ | 414 | 40.9 |

4. The saving is computed with respect to the always-on policy. The values in Table 3 indicate that the energy saving can be quite substantial, largely over 30% in all cases.

The average bandwidth per user is shown, for the three policies, in Fig. 10. Note that, for simplicity, the bandwidth $B$ provided by each AP is normalized to one. The curves show that the better energy performance of the traffic-based policy with $C_h = 4$ is paid not only in terms of a higher switch-off rate, but also in terms of a reduced bandwidth per connection. Indeed, by reducing, on average, the number of APs that are powered on, hence the available capacity, a higher energy saving is possible, but lower QoS (in terms of bandwidth) is provided to users. The bandwidth reduction that we observe in Fig. 10 for the traffic-based policy with $C_h = 4$ is not particularly critical during the low traffic period, when each user still succeeds in obtaining about 70% of the bandwidth of the only AP that is powered on, but it becomes critical during the peak traffic period, when the bandwidth per user is reduced of about 30% with respect to the always-on policy, thus violating the 2nd law of RoD policies (in exchange for an energy saving of just 10%).

The behavior of the association-based policy is very similar to that of the traffic-based policy with $C_h = 4$ when traffic is lowest, but at peak traffic becomes almost identical to that of the always-on policy, exhibiting a behavior that seems compatible with the 2nd law of RoD policies.

However, it is quite interesting to observe that, if the threshold $C_h$ for the traffic-based policy is reduced to 2, the bandwidth per user during the high traffic period becomes very close to that of the always-on policy, and during the low traffic period it is around 85% of the one of the always-on policy, much more than with the previous two policies. This is quite satisfactory, specially considering that the power consumption of this policy, reported in Fig. 9, is very close to that of the association-based policy.

## 5. MODEL VALIDATION

In this section we validate our analytical model by comparing its predictions against the experimental results of [6]. We consider as input the same traces [12], from which we extract the input parameters of our model (the association rate $\lambda_s$, the connection generation rate $\lambda_c$, the mean association time $1/\mu_s$, and the mean connection duration $1/\mu_c$), and we study a small cluster composed of 3 APs, each capable of serving up to 3 users, like in [6]. However, while in [6] the authors evaluate the energy savings only for low

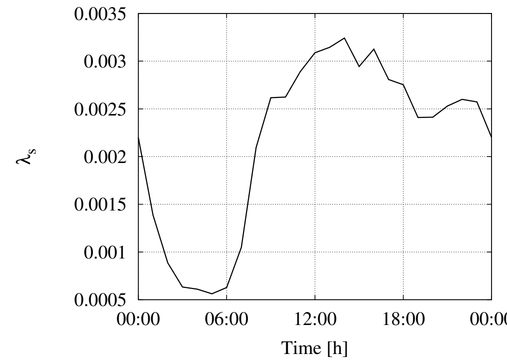


Figure 11: Association rate during a 24 h period derived from traces

Table 4: System parameters derived from traces

| time | $\lambda_s[1/s]$ | $\lambda_c[1/s]$ | $1/\mu_s[s]$ | $1/\mu_c[s]$ |
|---|---|---|---|---|
| on peak | $3.24 \cdot 10^{-3}$ | $3.27 \cdot 10^{-3}$ | 2081.68 | 2235.69 |
| off peak | $5.62 \cdot 10^{-4}$ | $3.24 \cdot 10^{-3}$ | 870.28 | 1855.28 |

and high traffic periods, our model allows to perform the analysis during the whole day.

We consider portions of the traces of length equal to one hour; for each portion, we select the 9 users which exchange the largest amount of traffic. From this group of users we extract $\lambda_c$ and $\mu_c$. Measurements in the traces are performed every 5 minutes. Let $B_j(i)$ denote the total amount of information that user $j$ exchanges at time interval $i$. We assume that a new connection is started at time $i$ if there exists a $j$ such that $B_j(i) > 0$ and $B_j(i-1) = 0$. Similarly, we assume that the connection of user $j$ ends if $B_j(i) = 0$ and $B_j(i-1) > 0$. Finally, we compute $\lambda_c$ as the total number of connections divided by the total time, and $\mu_c$ as the inverse of the average duration of a connection.

The association parameters $\lambda_s$ and $\mu_s$ are instead obtained directly from the traces. In particular, at 5 minutes intervals, the AP records the number of associated users. From this value, we derive $\lambda_s$ and the average number of associated users, $E[N_s]$. The parameter $\mu_s$ is then set to $\mu_s = \lambda_s / E[N_s]$. Finally, the parameters are averaged over all the APs.

Fig.11 shows the association rate $\lambda_s$ extracted from the traces, considering averaged values from five working days. As expected, $\lambda_s$ follows a typical day-night pattern, with off peak rate equal to 17% of the peak rate. Table 4 reports the parameter values at times of maximum and minimum traffic.

Fig.12 shows the variation of the average number of *on* APs throughout the day, considering both the association and the traffic-based policies. The results are comparable with those presented in [6]. In particular, during low traffic periods, only one AP is powered on, while during the rest of the day at least two APs need to be *on*, to support the user traffic.

The good agreement between our analytical results and the

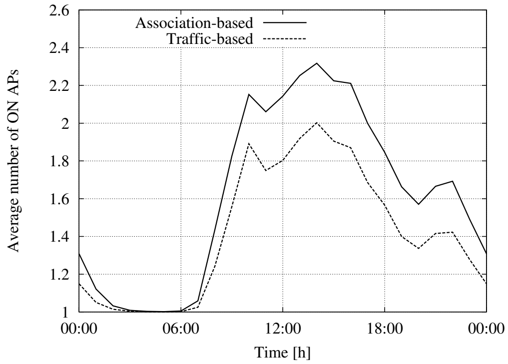


**Figure 12: Average number of on APs during a 24 h period**

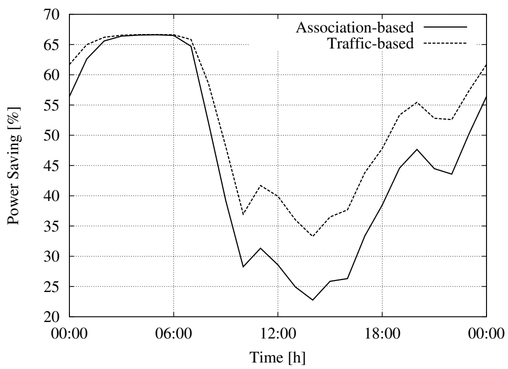


**Figure 13: Power saving achievable throughout the day**

experimental results of [6] indicates that the simplifying assumptions introduced in the model are not critical, at least in this simple setting.

To obtain more insight, Fig.13 details the power saving throughout the day. As expected, during the night the RoD policies allow saving 66% of the cluster power.

# 6. RELATED WORK

Several approaches have been recently proposed to reduce the power consumption of networks, and of the Internet in particular, starting from the seminal work of Gupta [13]. Researchers have shown that RoD policies can be applied to core [14], edge [15], mobile [4] and data-center networks [16]. In this paper we have proposed an energy-efficient RoD policy for corporate WLANs, in which the dense distribution of APs is exploited to power off unused devices. Several recent works [17, 18] proposed approaches to reduce the energy consumption in WLANs, but they mainly focused on the user side, in order to preserve the battery lifetime. These approaches can be easily integrated with energy-aware policies, that efficiently control the AP power consumption.

Looking at the internal architecture of the APs normally deployed in WLANs, researchers have observed that the largest amount of power is due to base components, rather than transmission circuits, so that an AP consumes approximately the same amount of power, independently from the traffic that is flowing through it [19]. This confirms that powering off unused APs can be a viable solution to save energy.

Industries are currently studying how to power on and switch off APs (and, in general, network devices) [20], using existing techniques, such as Power-Over-Ethernet, to remotely perform these actions [21]. Our approach is naturally integrated into such environments, allowing the network manager to effectively decide *when* devices should enter sleep mode and when they must be waken up. As reported in [22], introducing power-off mechanisms in APs could save millions of dollars worldwide, or, equivalently, many tons of $CO_2$.

Finally, renewable sources can be easily adopted to supply energy for APs. This idea has been investigated in the last years by researchers [23, 24], and commercial solutions using photovoltaic systems already exist [25]. Our approach can be of benefit also in this case, by limiting the amount of power needed during low traffic periods, specially at night, when the energy of the sun cannot be exploited.

# 7. CONCLUSIONS

In this paper we have presented a first simple analytical model for the assessment of the effectiveness of resource-on-demand policies for the energy-efficient operation of dense WLANs with centralized management, typical of today's corporate environments. The model simplicity stems from several simplifying assumptions, that will permit extensions to incorporate more complex network topologies and user characteristics.

We have used the proposed model to study two simple resource-on-demand policies, that, based on the instantaneous WLAN parameters, select the appropriate number of APs to activate, thus trying to avoid to waste energy on underutilized APs. The first policy is based on the number of users that are *associated* with each AP (and is thus called association-based), the second is based on the number of *active* users, i.e., users that are generating traffic (and is thus called traffic-based). Both policies can be tailored, by defining a value of threshold for the activation of additional APs, and a value of hysteresis to increase the operation stability.

In both cases, our results indicate that ample room exists for energy saving. For the considered setups, we identified the possibility to save up to 87% of the power requested to operate the WLAN APs, during low traffic periods, with both policies.

The traffic-based policy with threshold 4 also achieves a limited saving during peak traffic periods, but it exhibits lower stability than the association-based policy, and limits the per-user bandwidth in high traffic conditions. However, by reducing the threshold to 2, the amount of per-user bandwidth improves drastically. Thus, our results indicate that the association-based policy, which is also simpler to implement, may be a wise choice for the preliminary implementation of a simple and effective approach to reduce the energy consumption of corporate WLANs. However, more dynamic policies, like the traffic-based, if carefully tuned, may be able to provide better overall performance.

A much more careful and extensive study is necessary to completely characterize the behavior of the two considered policies in different WLAN scenarios. This study will be made possible by the simple analytical model that was introduced in this paper, and will be the subject of further research, together with the definition of more elaborate policies to simultaneously achieve large energy savings, and a QoS very close to the one obtained when all APs are powered on.

# 8. REFERENCES


[1] L. Chiaraviglio, D. Ciullo, M. Meo, M. Ajmone Marsan, "Energy-Aware UMTS Access Networks," *First International Workshop on Green Wireless (W-GREEN 2008)*, Lapland, September 2008.

[2] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, "Optimal Energy Savings in Cellular Access Networks," *First International Workshop on Green Communications (GreenComm 2009)*, Dresden, Germany, June 2009.

[3] M. Ajmone Marsan, M. Meo, "Energy Efficient Management of two Cellular Access Networks," *GreenMetrics 2009*, Seattle, WA, June 2009.

[4] L. Chiaraviglio, D. Ciullo, M. Meo, M. Ajmone Marsan, "Energy-Efficient Management of UMTS Access Networks," *21st International Teletraffic Congress (ITC 2009)*, Paris, France, September 2009.

[5] A. Jardosh, G. Iannaccone, K. Papagiannaki, B. Vinnakota, "Towards an Energy-Star WLAN Infrastructure," *11th ACM International Workshop on Mobile Computing Systems and Applications (HotMobile)*, Tucson, AR, USA, February, 2007.

[6] A. Jardosh, K. Papagiannaki, E. Belding, K. Almeroth, G. Iannaccone, and B. Vinnakota, *Green WLANs: On-Demand WLAN Infrastructure*, Mobile Networks and Applications (MONET), special issue on Recent Advances in WLANs, April 2009.

[7] http://www.metageek.net/products/inssider.

[8] *Aruba WLAN at Microsoft Exceeds 11,000 Access Points*, http://www.wi-fiplanet.com/news/article.php/3753466.

[9] IEEE 802.11v: Wireless Network Management. http://grouper.ieee.org/groups/802/11/Reports/ tgv update.htm.

[10] M. Papadopouli, H. Shen, and M. Spanakis, "Modeling Client Arrivals at Access Points in Wireless Campus-Wide Networks," *14th IEEE Workshop on Local and Metropolitan Area Networks (LANMAN 2005)*, Crete, Greece, September 2005.

[11] A. Mahanti, C. Williamson, and M. Arlitt, *Remote Analysis of a Distributed WLAN Using Passive Wireless-Side Measurement*, Performance Evaluation, Vol. 64, n. 9-12, pp. 909-932, 2007.

[12] D. Kotz, T. Henderson, and I. Abyzov, CRAWDAD trace set. Dartmouth/campus/snmp (v. 2004-11-09), http://crawdad.cs.dartmouth.edu/dartmouth/ campus/snmp, November 2004.

[13] M. Gupta, S. Singh, "Greening of the Internet," *ACM SIGCOMM 2003*, Karlsruhe, Germany, August 2003.

[14] P. Barford, J. Chabarek, C. Estan, J. Sommers, D. Tsiang, S. Wright, "Power Awareness in Network Design and Routing," *IEEE INFOCOM 2008*, Phoenix, USA, April 2008.

[15] P. Tsiaflakis, Y. Yi, M. Chiang, and M. Moonen, "Green DSL: Energy-Efficient DSM," *IEEE ICC'09*, Dresden, Germany, June 2009.

[16] A. Qureshi, R. Weber, H. Balakrishnan, J. Guttag and B. Maggs, "Cutting the Electric Bill for Internet-Scale Systems," *ACM SIGCOMM 2009*, Barcelona, Spain, August 2009.

[17] Y. Agarwal, R. Chandra, A. Wolman, P. Bahl, K. Chin, and R. Gupta, "Wireless Wakeups Revisited: Energy mManagement for VoIP over Wi-Fi Smartphones," *ACM International Conference on Mobile Systems, Applications, and Services (MobiSys) 2007*, San Juan, Puerto Rico, June 2007.

[18] W. Ye, J. Heidemann, and D. Estrin, *Medium Access Control with Coordinated Adaptive Sleeping for Wireless Sensor Networks*, IEEE/ACM Transactions on Networking (ToN), Vol. 12, n. 3, pp. 493-506, June 2004.

[19] *A Comparison of Efficiency, Throughput, and Energy Requirements of Wireless Access Points*, http://www.iol.unh.edu/services/testing/wireless.

[20] *Cisco Energy Wise*, http://www.cisco.com/.

[21] *Reducing Power consumption: Improving the Heat and power Efficiency of switching and Telephony equipment* Alcatel White Paper, http://www.alcatel-lucentbusinessportal.com.

[22] *Sustainable Mobility: Cisco Strategies for Green Wireless*, www.cisco.com/web/go/nextgen-wireless.

[23] A.A. Sayegh, T.D. Todd, "Energy Management in Solar Powered WLAN Mesh Nodes Using Online Meteorological Data," *IEEE ICC'07*, Glasgow, UK, June 2007.

[24] F. Zhang, T.D. Todd, D. Zhao, V. Kezys, *Power Saving Access Points for IEEE 802.11 Wireless Network Infrastructure*, IEEE Transactions on Mobile Computing, Vol. 5, n. 2, pp. 144-156, February 2006.

[25] http://meraki.com/news/2007/06/03/meraki-introduces-first-solar-powered-outdoor-wi-fi-access-kit/.