Computer Networks 56 (2012) 2522–2537

Contents lists available at SciVerse ScienceDirect

Computer Networks

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

Energy-performance trade-off in dense WLANs: A queuing study

Ana Paula Couto da Silva [a], Michela Meo [b,*], Marco Ajmone Marsan [b,c]

[a] Computer Science Department, Federal University of Juiz de Fora, Brazil
[b] Dip. Elettronica e Telecomunicazioni, Politecnico di Torino, Italy
[c] Institute IMDEA Networks, Leganes, Madrid, Spain

A R T I C L E I N F O

Article history:
Received 30 September 2011
Received in revised form 7 March 2012
Accepted 20 March 2012
Available online 30 March 2012

Keywords:
Energy efficiency
Wireless LANs
Queuing models

A B S T R A C T

The increasing concern about the energy consumption of communication networks is driving the research community to identify approaches to save energy in the networks of today. For instance, considering wireless local area networks (WLANs), the activation of network resources can be driven by the user demand, avoiding having to always power on all Access Points (APs). In this paper, we consider a portion of a dense WLAN system, where many APs are deployed to provide sufficient capacity to serve a large number of active users during peak traffic hours. To provide large capacity, a number of APs are colocated in the same position 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 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 several 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. The main focus of our study is the investigation of the type of algorithm that should be used for the association of active users with APs in order to increase the amount of saved energy in dense WLANs. Results show that when some system state information is available, such as the number of users associated with each AP, the energy consumption can decrease up to 20%. Furthermore, our study gives comprehensive insights on the trade-off between the opposite needs to save energy and provide quality of service to the users.

© 2012 Elsevier B.V. All rights reserved.

1. Introduction

Dense, centrally managed wireless local area networks (WLANs), where a very large number of Access Points (APs) is installed to provide users with a bandwidth comparable to that of wired Internet access, are becoming common in large enterprises (some industrial campuses report over 10,000 APs). Dense WLANs can also be identified in academic campus networks, although density is usually lower in this case. 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 [1], 14 of which belonging to the institutional campus network. In these cases, the WLAN deployment is often designed making use of *clusters* of APs that provide coverage and capacity over the same area. Centralized management of the clusters provides a powerful instrument for the implementation of schemes that handle interference, trying to minimize its effect while optimizing coverage and performance. For example, when traffic is low and relatively little capacity is needed, these management algorithms might decide to switch off some channels, i.e., APs, to reduce interference. Recently, the growing awareness of the need to save energy [2], with the twofold objective of reducing cost and being friendly with the environment, has introduced the use of these centralized management schemes as means to improve energy-efficiency in dense WLANs [3]. Thus, these management schemes try to

* Corresponding author.
E-mail address: michela.meo@polito.it (M. Meo).

1389-1286/$ - see front matter © 2012 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.comnet.2012.03.017

find a good trade-off among coverage, energy consumption, performance.

In this paper, we consider a portion of a dense WLAN system, where many APs are deployed to provide sufficient capacity to serve a large number of active users during peak traffic hours. To provide large capacity, a number of APs are colocated in the same position 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 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 than one group. 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. The main focus of our study is the investigation of the type of algorithm that should be used for the association of active users with APs in order to increase the amount of saved energy in dense WLANs.

Previous simulation and experimental studies in this context considered real but small WLAN setups [3], while analytical studies considered a portion of the WLAN area in which several APs exist, all providing *exactly* the same coverage [4]. In this paper, we extend the analysis of [4] to a portion of the WLAN where two groups of APs offer service over a given service area, that is the union of the two areas resulting from the coverage of the two groups. This is quite a significant extension with respect to [4], since in the case of [4] the choice of the AP to which an active user should connect is trivial. On the contrary, in the new setting, the algorithm that the active users in the area covered by both groups of APs adopt for the selection of the AP to which they should connect, influences the energy consumption of the WLAN.

To investigate the energy-performance trade-off, we develop a model based on two coupled queues: each queue represents a group of colocated APs, and customers arriving at the system might be served by one or both queues depending on their position, that is probabilistically represented in the model. We also propose extensions of the model to the case of more that two groups of APs. Due to the growth of the state space with the number of groups, the extensions are based on accurate approximations that we discuss in the final part of this paper.

Queuing models have been for many years one of the most popular tools for the quantitative analysis and design of computer and telecommunication systems, from telephony to packet networks. The types of metrics traditionally derived from queuing models are extremely varied, 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 little interest was paid in the past to metrics related to the system energy consumption. This is not due to a limitation in the modeling power of queues; rather, it reflects the lack of attention that in the past was paid to the energy characteristics of computer and telecommunication systems. Now, energy consumption has become an important element of the design space, and models are being developed with the objective of characterizing the energy properties of ICT systems of different types. As a result, energy-related metrics have started to appear in the context of queuing models. Examples of recent papers where queues are used to assess the energy properties of systems are [5], where an M/GI/1 queue with processor sharing service is used to investigate the energy properties of dynamic speed scaling in processor sharing systems, and [6,7], where single-server and multi-server queuing models of server farms are investigated, with the objective of identifying server management algorithms that optimize the Energy-Response time product.

In the context of green networking, solutions based on the use of sleep modes have been proposed in many different scenarios and contexts as approaches that, while being effective, are also easily applicable to already available technologies. Many papers consider the case of cellular access networks. In [8], the authors show that, in urban areas, it is possible to put some base stations in sleep mode during low-traffic periods, while still guaranteeing quality of service (QoS) constraints in terms of blocking probability and electromagnetic exposure limits. The impact on the power consumption of deployment strategies based on layouts featuring varying numbers of micro base stations per cell in addition to conventional macro sites has been investigated in [9]. For scenarios with full traffic load, the use of micro base stations has a rather moderate effect on the area power consumption of a cellular network. In [10], the authors evaluate the energy saving that can be achieved with the energy-aware cooperative management of the cellular access networks of more than one operator offering service over the same area. In [11], the authors propose a component-level deceleration technique in base stations operation, called Speed-Scaling, that is more conservative than entirely shutting down BSs; it dynamically saves power during periods of low load while ensuring full coverage at all times. The proposed distributed iterative algorithm, SpeedBalance, addresses the tradeoff between delay and energy in both networking and processing components of BSs. Extensive simulations show that SpeedBalance can yield to significant energy savings of about 30–45%. In [12], the authors show how cooperation increases the covered area of a base station. They propose a model to compute the power needed to cover an area and analyze how cooperation reduces the power consumption. Based on the performed analysis, they also show that using a small number of base stations to cooperate on the transmission to mobile devices is sufficient to save significant amounts of power.

This paper is organized as follows: Section 2 contains the problem statement, with the description of the considered portion of the dense WLAN, and its queuing model, together with the definition of performance metrics and customer management algorithms. Section 3 presents and discusses numerical results obtained with the detailed model of the system. Section 4 describes the approximate models and discusses their accuracy, as well as their use in more general settings. Finally, Section 5 concludes the paper.

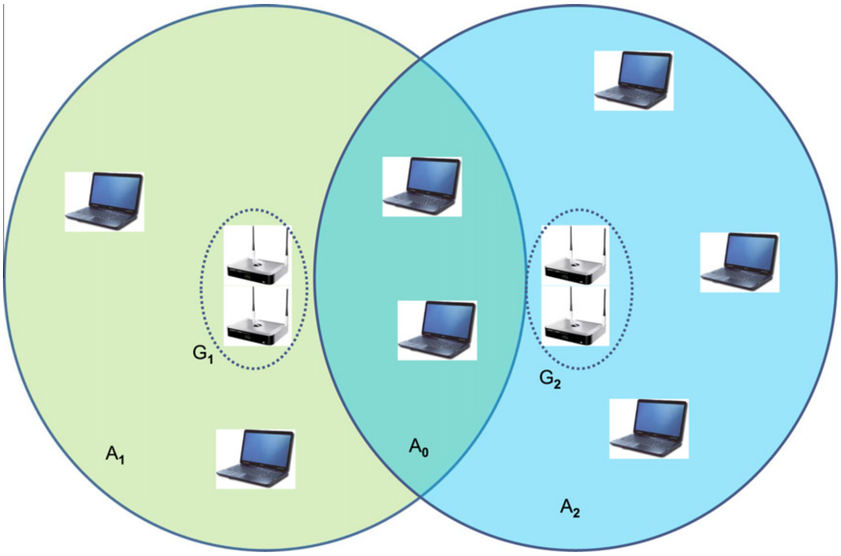

Fig. 1. Sketch of the considered portion of a dense WLAN.

2. Problem statement

In this section, we describe the system we focus on, we present the analytical model we propose to study the system and we discuss possible customer management (CMN) algorithms.

The most important terms of the notation used throughout the paper are summarized in Table 1.

2.1. The Wireless LAN

We consider a portion of a dense WLAN in which an area $A = A_1 \bigcup A_2$ is covered by several access points (AP) in a manner such that a group G_1 of n_1 APs covers area A_1, and another group G_2 of n_2 APs covers area A_2, see Fig. 1 for a sketch of the system. Each AP can allocate up to K mobile terminals (MTs); this means that no more than K MTs can be associated at the same time with the AP.

Table 1
Notation.

Symbol	Description
n_i	Number of APs in group i
K	Maximum number of MTs simultaneously associated with an AP
λ	Rate of association request
$1/\mu$	Mean association time
N_i	Maximum number of MTs associated with group i
α_i	Probability that a request is originated from a MT in area A_i
ρ	System load
η	Energy consumption of an active AP [W]
ϵ	Energy needed for an AP activation/deactivation [J]
u_i	Number of unconstrained MTs in A_i
c_i	Number of constrained MTs in A_i
S	State space of the proposed model
γ_i	Approximate customer arrival rate to A_i
δ	Approximation relative error

Access points can be switched on and off according to the number of MTs to be served. However, at least one AP in each group must be on at any given time, in order to provide coverage over the whole area A. A fraction Ψ_1 of the total area A corresponds to $A_1 \cap \overline{A}_2$, so that the APs in group G_2 cannot serve MTs in this area. Similarly, a fraction Ψ_2 of the total area A corresponds to $A_2 \cap \overline{A}_1$, that cannot be served by the APs in group G_1. Finally, MTs located in area $A_1 \cap A_2 = A_0$, which corresponds to a fraction $1 - \Psi_1 - \Psi_2 = \Psi_0$ of the total area, can be served by any AP. APs can be activated and deactivated based on some algorithms (which are the main focus on this paper) triggered by MT arrivals and departures. When an AP is activated triggered by the arrival of a MT, the switching time might translate into an increase of the time needed by the MT to find and associate with a network. While we neglect the effect of this time, let us notice here that some solutions are possible to reduce it. For example, the MT can be temporary associated with other already active APs, or an AP can be activated before it is strictly needed. In this paper, we do not discuss these techniques to engineer the proposed solutions, we just estimate the possible achievable saving.

MTs are uniformly distributed over the area, and they collectively request access according to a Poisson process with rate λ. The MTs remain associated with the AP for a time described by an exponentially distributed random variable with rate μ. Since the time a MT is associated with an AP depends on the user behavior and only very marginally on the system performance, the rate μ is independent on the number of MTs in the system. Indeed, the time a MT is associated with an AP is equivalent to the time a user is connected to a network and it is of the order of minutes or hours; thus, the possible effect of network performance on user's behavior can be considered negligible and μ can be made independent on the number of MTs connected to the AP.

2.2. The queuing model

We model the portion of the dense WLAN with the queuing system shown in Fig. 2. The system is composed of two multiserver stations with no waiting room. Each station models a group of APs. The service corresponds to the MT association with an AP. The two stations comprise N_1 and N_2 servers, respectively, where $N_i = Kn_i$ corresponds to the maximum number of MTs that can be associated with group i, that is composed of n_i APs, each one limited to host K MTs. As an example, Fig. 3 shows the case of station 1 with $n_1 = 2$ APs.

MTs request service, i.e., request to associate, according to a Poisson process with rate λ. This is the customer arrival process in the queueing model. With probability equal to α_1 an arriving customer is in area $A_1 \cap \overline{A}_2$ and it is, thus, constrained to proceed to station 1. By assuming uniform arrival rate in the area, the probability of arriving in A_1 is proportional to the size of the area, and it is thus $\alpha_1 = \Psi_1$. if the customer cannot be served by any AP in group 1, because the maximum number of MTs has already been reached (i.e., it cannot find an idle server), the customer is lost. The same occurs to group 2 with probability $\alpha_2 = \Psi_2$ (where $\alpha_1 + \alpha_2 < 1$). With probability $\alpha_0 = \Psi_0 = 1 - \alpha_1 - \alpha_2$, an arriving customer is unconstrained, i.e., it can be associated with either one of the two groups, according to the system CMN algorithm, which we discuss later, and the customer is lost only when no idle server exists at either station.

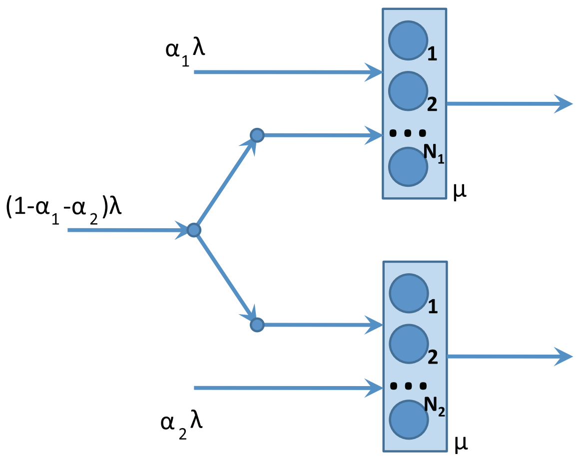

Fig. 2. Queuing model of the system.

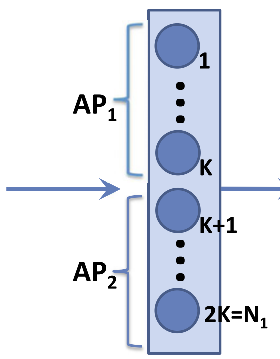

Fig. 3. Example of the queuing system modeling a group of APs.

Service times are i.i.d. random variables with exponential distribution with rate μ. At the end of a service, unconstrained customers that remain in the system may change station, according to the CMN algorithm. The system load can be defined as $\rho = \lambda/(\sum_i n_i K\mu)$.

Each AP (and the corresponding K servers in the model) can be either *active* or *dormant*. Active servers can be either busy, i.e., serving a customer, or idle, i.e., ready to provide service to an incoming customer. In the model, the transition of an AP from dormant to active and vice versa corresponds to groups of K servers that, at a time, become dormant or active. We assume that as soon as K active servers are idle, they are deactivated, and that if all active servers are busy when a customer arrives, K dormant servers are activated, if available. We assume that the time needed to complete a transition is negligible with respect to the time-scale of request arrivals and departures.

Each active AP (group of K servers) consumes energy at a rate η W, while a group of dormant servers consumes negligible energy. The energy consumption may be a function of load, or of the number of active servers in the group; however, in this paper we always assume η to be constant, which is consistent with the characteristics of today's hardware. The activation or deactivation of an AP requires an energy equal to ϵ J.

The state of the queuing system is defined by a quadruple $\bar{s} = (u_1, c_1, u_2, c_2)$, with u_i the number of unconstrained customers at station i and c_i the number of constrained customers at station i. Considering that $u_i + c_i \leqslant N_i$, the cardinality of the state space S is given by:

$$|S| = \left(\frac{(N_1+1)(N_1+2)}{2}\right)\left(\frac{(N_2+1)(N_2+2)}{2}\right) \quad (1)$$

that is of the order of $(N_1N_2)^2$.

2.3. The performance metrics

For the WLAN and the queuing system we introduce the following performance metrics:

- customer loss probability (probability that a MT cannot associate with an AP)
- average number of customers at each station (average number of MTs associated with each AP group)
- average power consumption (average power consumption of APs)
- average energy consumed to serve one customer (average energy consumed by APs to serve one MT).

The customer loss probability can be computed from the queuing model steady-state distribution as:

$$P[loss] = \alpha_1 \sum_{\bar{s}\in S_1} \pi(\bar{s}) + \alpha_2 \sum_{\bar{s}\in S_2} \pi(\bar{s}) + (1-\alpha_1-\alpha_2) \sum_{\bar{s}\in S_3} \pi(\bar{s}), \quad (2)$$

where $S_1 = \{\bar{s}|u_1 + c_1 = N_1\}$, $S_2 = \{\bar{s}|u_2 + c_2 = N_2\}$, and $S_3 = \{\bar{s}|u_1 + c_1 = N_1 \& u_2 + c_2 = N_2\}$. The customer loss probability is computed as the fraction of customers that,

Table 2
Customer management algorithms.

Algorithm	Working principle	Triggering action
R	Request is randomly allocated to G_1 or G_2	MT arrival on A_0
F1	Request is preferentially associated with G_1	MT arrival on A_0
F1-R	Request is preferentially associated with G_1	MT arrival on A_0 or MT de-association
AB	Request is associated with the group with the smallest number of MTs	MT arrival on A_0
EA	Request is preferentially associated with G_1 provided a new AP is not turned on	MT arrival on A_0
EA-AB	Request is preferentially associated with G_1 provided a new AP is not turned on and the number of MTs in G_1 is smaller than in G_2	MT arrival on A_0
EA-AB-R	Request is preferentially associated with G_1 provided a new AP is not turn on and the number of MTs in G_1 is smaller than in G_2	MT arrival on A_0 or MT de-association

when arriving at the system, find it full. Thus, it is derived as the sum of the probabilities of all cases in which a MT arrives to area A_i, and the system is unable to accommodate it.

The average number of customers at station i can be computed as:

$$E[q_i] = \sum_{\bar{s}\in S}(u_i + c_i)\pi(\bar{s}). \tag{3}$$

The average power consumption can be computed as:

$$E[P] = E[P_s] + E[P_{act}] + E[P_{de}], \tag{4}$$

where $E[P_s]$ is the power consumed by active server groups, and $E[P_{act}]$, $E[P_{de}]$ are the powers consumed for server group activation and deactivation, respectively. We have:

$$E[P_s] = \eta\left\{2\pi_{(0,0,0,0)} + \sum_{j=1}^{n_2}\sum_{\bar{s}\in S_{0j}}(j+1)\pi(\bar{s}) + \sum_{i=1}^{n_1}\sum_{\bar{s}\in S_{i0}}(i+1)\pi(\bar{s}) + \right.$$
$$\left. + \sum_{i=1}^{n_1}\sum_{j=1}^{n_2}\sum_{\bar{s}\in S_{ij}}(i+j)\pi(\bar{s})\right\}, \tag{5}$$

where $S_{0j} = \{\bar{s}|u_1 + c_1 = 0\&(j-1)K + 1 \leqslant u_2 + c_2 \leqslant jK\}$, $S_{i0} = \{\bar{s}|(i-1)K + 1 \leqslant u_1 + c_1 \leqslant iK\&u_2 + c_2 = 0\}$, $S_{ij} = \{\bar{s}|(i-1)K + 1 \leqslant u_1 + c_1 \leqslant iK\&(j-1)K + 1 \leqslant u_2 + c_2 \leqslant jK\}$. The first term of the sum accounts for the consumption when no customer is present in the system, in this case one AP per group must be kept on. The second term accounts for the case in which there is nobody in the first group (but one AP must be active) and j APs are active in the second group; the third term is the dual for the case in which there is nobody in the second group. And, finally, the last term corresponds to the case in which i and j APs are active in the first and second group, respectively.

$E[P_{act}]$ can be written as:

$$E[P_{act}] = \lambda\epsilon\left\{\alpha_1\sum_{\bar{s}\in S_{act1}}\pi(\bar{s}) + \alpha_2\sum_{\bar{s}\in S_{act2}}\pi(\bar{s}) + (1-\alpha_1-\alpha_2)\sum_{\bar{s}\in S_{act0}}\pi(\bar{s})\right\} \tag{6}$$

$$E[P_{de}] = \mu\epsilon\left\{\sum_{\bar{s}\in S_{de1}}(u_1 + c_1)\pi(\bar{s}) + \sum_{\bar{s}\in S_{de2}}(u_2 + c_2)\pi(\bar{s})\right\}, \tag{7}$$

The sets S_{act1}, S_{act2}, and S_{act0} include those states where an arrival produces the activation of a group of K servers. In the cases of S_{act1} and S_{act2}, a constrained customer arrives on areas A_1 and A_2, respectively, and a new group of servers is activated at queues 1 and 2, respectively. In the case of S_{act0}, an unconstrained customer arrives on area A_0, and a new group of servers is activated at either queue 1 or queue 2. The sets S_{de1} and S_{de2} include those states where an end of service at queues 1 and 2, respectively, produces the deactivation of a group of K servers.

The average energy consumed to serve one customer can be computed as:

$$E[E_c] = \frac{E[P]E[T]}{E[q_1] + E[q_2]} = \frac{E[P]}{\lambda(1 - P[loss])}, \tag{8}$$

where $E[T]$ is the average time spent by a customer in the queuing system, which is computed by Little's Law as:

$$E[T] = \frac{E[q_1] + E[q_2]}{\lambda(1 - P[loss])}. \tag{9}$$

2.4. The customer management algorithms

We consider several alternatives as regards the algorithm used to manage unconstrained customers, that correspond to MTs in A_0. The algorithms mainly differ in two aspects. The first one is the state information that is used to take decisions about MT association with the APs: in the various algorithms presented below, state information goes from the simple binary information about the system being full or not, to complete state including, for all the APs, the state and number of associated MTs. The second aspect is the kind of event that can trigger the algorithm: in some cases the algorithm is triggered by the request of a new association (a customer arrival, in the queuing system model), in other cases both arrivals and departures, i.e., MT associations and de-associations, trigger the algorithm.

Table 2 shows the main characteristics of the proposed CMN algorithms, that are described below in detail:

1. **Random (R)** – Each MT in A_0 that requests access chooses the access point group G_1 or G_2 with equal probability. If the chosen group cannot accommodate the MT request, the MT is associated with the other group, if possible. If both groups are full with customers, the request is refused.
   This algorithm does not try to optimize energy consumption, allowing MTs to randomly choose the AP group. The only information required at MTs concerns the full/not full state of AP groups at arrival instants.

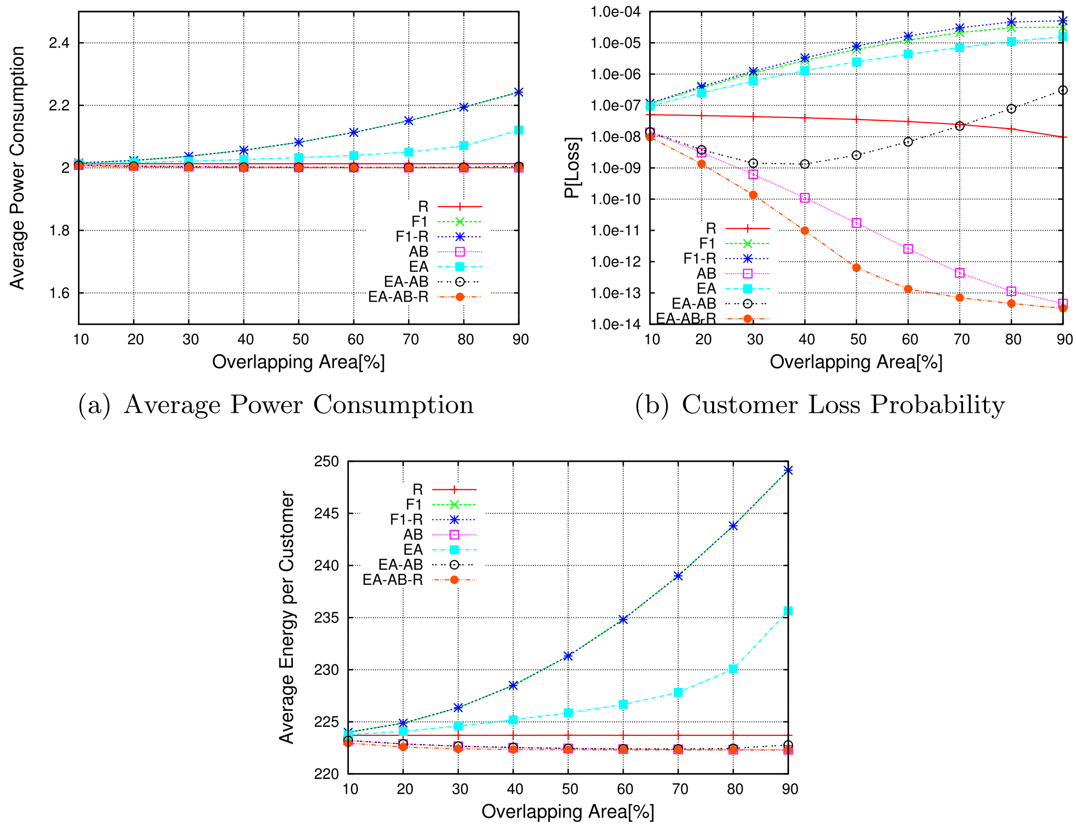

(a) Average Power Consumption

(b) Customer Loss Probability

(c) Average Energy per Customer

Fig. 4. Set of metrics for the equal load case with $\rho = 0.2$.

2. **Fill G_1 first (F1)** – MTs in A_0 that request access are preferentially associated with G_1, insofar as possible. If G_1 is full, MTs are associated with G_2, but if also G_2 is full, the request is refused.
   This algorithm tries to concentrate the load on G_1, thus reducing the need for active APs in G_2. The only information required at MTs concerns the full/not full state of AP groups at arrival instants.
3. **Fill G_1 first with Reassociation (F1-R)** – This strategy is similar to the previous one, F1, but the algorithm can also be triggered by a MT de-association. Thus, the end of a service of a MT in A_0 associated with G_1 triggers a new control over the active MTs. The control, in its turn, can make a MT in A_0 associated with G_2 be transferred to G_1.
   This algorithm tries to concentrate the load on G_1, even more than F1. The information required at MTs concerns the full/not full state of AP groups at both arrival and departure instants.
4. **Association Balancing (AB)** – MTs in A_0 that request access are associated with the group of access points that is serving the smallest number of MTs (considering both constrained and unconstrained MTs). If the two groups are serving the same number of MTs, the MT is associated with G_1. If both groups are full, the request is refused.
   This algorithm tries to optimize energy consumption through load balancing over the two AP groups. The information required at MTs concerns the exact state of AP groups at arrival instants.
5. **Energy-Aware (EA)** – MTs in A_0 that request access are associated with G_1, provided this does not require switching on a new access point. Otherwise, the same check is made on G_2. Furthermore, a new access point is preferentially turned on in G_1.
   This algorithm tries to optimize energy consumption by reducing the number of active APs in the two groups. The information required at MTs concerns the exact state of AP groups at arrival instants.
6. **Energy-Aware with Association Balancing (EA-AB)** – MTs in A_0 that request access are associated with G_1, provided this does not require switching on a new access point and the number of MTs served by G_1 is smaller than for G_2. Otherwise, the MT is associated with G_2, provided this does not require switching on a new access point. A new access point is preferentially turned on in G_1.
   This algorithm tries to optimize energy consumption by reducing the number of active APs in the two groups, and by load balancing. The information required at MTs concerns the exact state of AP groups at arrival instants.

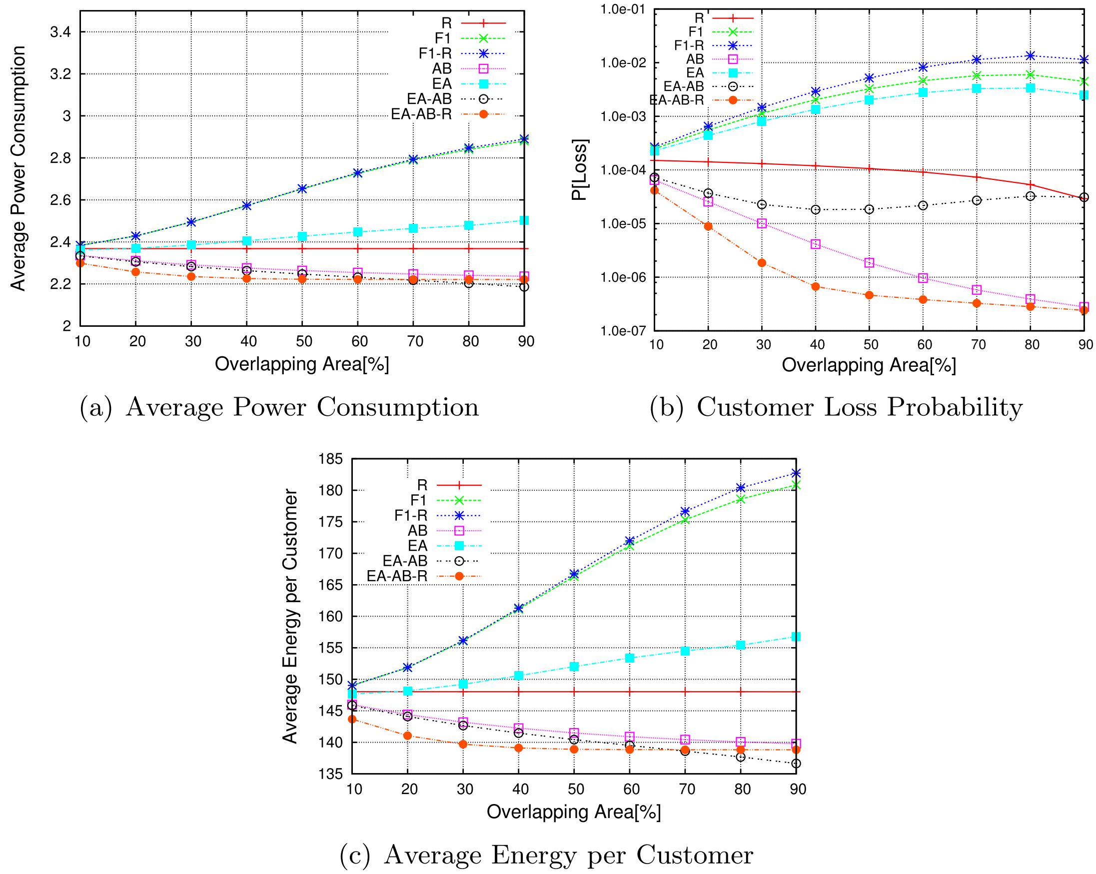

(a) Average Power Consumption (b) Customer Loss Probability

(c) Average Energy per Customer

Fig. 5. Set of metrics for the equal load case with $\rho = 0.4$.

7. **Energy-Aware with Association Balancing and Reassociation (EA-AB-R)** – The strategy is the same as the previous one, but it is also run at each user de-association.
   This algorithm tries to optimize energy consumption by reducing the number of active APs in the two groups, and by load balancing. The information required at the MTs includes the exact state of AP groups at arrival and departure instants.

3. Numerical results

In this section we discuss numerical results derived from the solution of the queuing model with the different CMN algorithms. We present in particular a comparative analysis of the behavior of the different CMN algorithms, trying to gain insight in the approaches to be followed for energy efficiency.

We consider scenarios that differ in the system load, as well as in the fraction of the service area that can be served by both groups of APs. We fix the time that each MT remains associated with an AP to be exponentially distributed with mean equal to 1000 s (i.e., $\mu = 0.001\ s^{-1}$, adopting the same values used in [4]), and we vary λ, as well as α_1 and α_2.

The energy consumption per time unit for each access point is assumed to be constant, and is set to $\eta = 1$ W. Initially, the energy spent for an AP activation or deactivation is neglected, i.e., we set $\epsilon = 0$ J.

3.1. Basic scenario

We start by considering the case in which the load generated in the areas $A_1 \cap \overline{A_2}$ and $A_2 \cap \overline{A_1}$ are equal. The fraction of the load generated in area $A_0 = A_1 \cap A_2$ varies from 10% to 90% of the total load, meaning that the overlap between the areas A_1 and A_2 in Fig. 1 varies. Since the distribution of MTs over the service area is assumed to be uniform, the load in a given area is proportional to its size. The maximum number of MTs that can associate with an AP equals $K = 10$, and group i comprises $n_i = 2$ APs, with $i = 1, 2$, so that the maximum number of MTs allowed in each AP group equals $N_i = 20$. The results are shown for a total load on each area A_i which is $\rho_i = \lambda/(n_i K \mu) \in \{0.2, 0.4, 0.6, 0.8\}$.

Figs. 4–7 show numerical results for: (a) average power consumption, (b) customer loss probability, (c) average energy per customer, versus the fraction of the load generated in area A_0, for variable system load.

Several interesting observations derive from the results in Figs. 4–7:

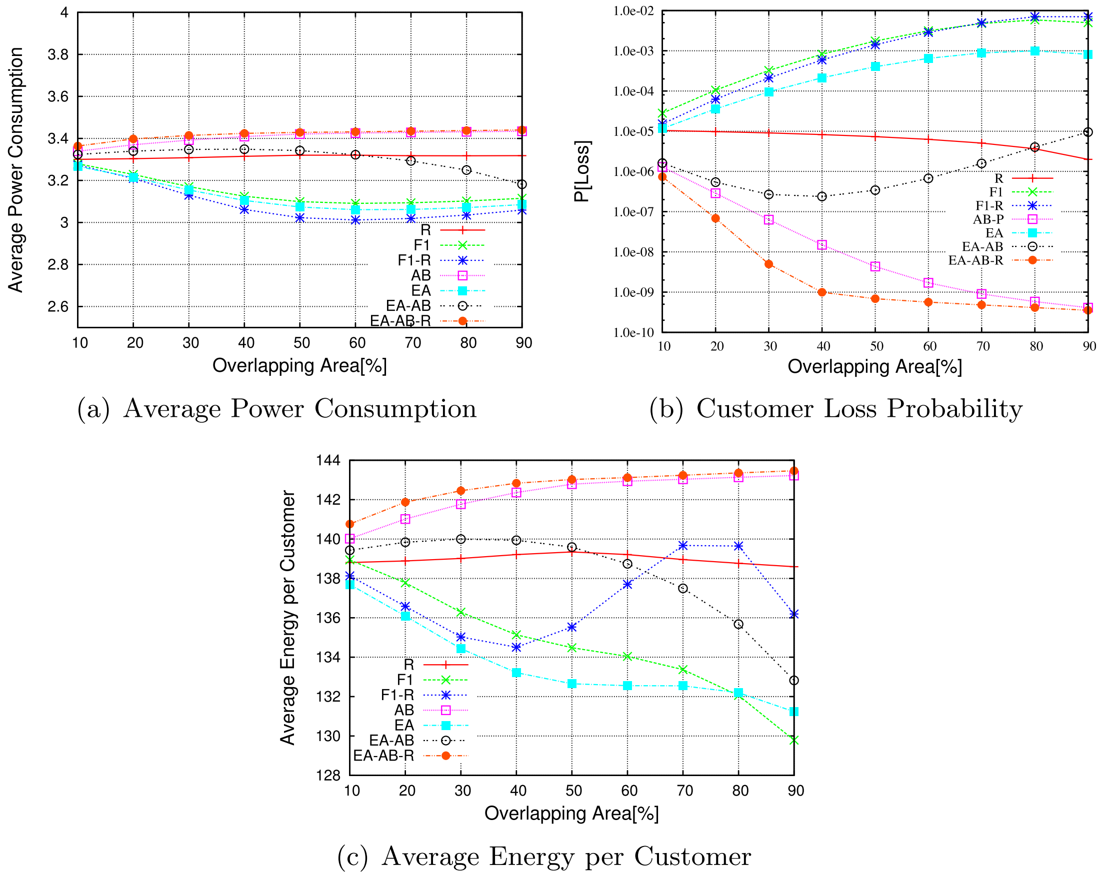

Fig. 6. Set of metrics for the equal load case with $\rho = 0.6$.

1. Focusing on the lower loads, $\rho = \{0.2, 0.4\}$, we see that CMN algorithms that balance the load between areas (i.e., the AB, EA-AB and EA-AB-R algorithms) yield lower average energy consumption with respect to algorithms that push the MTs to AP group 1, and the former consume less energy than the algorithm that randomly chooses an AP group for unconstrained customers, which also stochastically balances customers between APs. Remarkably, these algorithms, while yielding smaller customer loss probability, consume less energy. As a result, they provide the smallest values of energy per user. Indeed, as can be noted from the average number of customers at the two stations (not plotted for brevity), algorithms F1 and F1-R (and even EA, although at a lower extent) tend to bring more customers to AP group 1, and, thus, lead to high loss probability; however, they also consume more. The reason is that by pushing customers to the first group, the second group is quite under-utilized, but still one AP in that group is always powered on, and energy is not that effectively used, as shown also by the energy per customer plot. The largest difference can be observed at 90% of unconstrained MTs, where the EA-AB-R algorithm decreases the average energy consumption by about 25% with respect to F1.
2. Focusing on the higher loads $\rho = \{0.6, 0.8\}$, we see that the situation is reversed: the CMN algorithms that push the MTs to AP group 1 outperform the other policies in terms of energy consumption, at the price of a probably too high loss probability. Indeed, by having higher loss probability, these policies end up serving a lower number of MTs and having a higher chance to switch off some APs. This, in its turn, makes the average energy per customer smaller for these policies than for the AB, EA-AB and EA-AB-R algorithms.
3. As expected, the average number of MTs connected to the APs in group 1 is larger for the algorithms that push MTs to this group and, consequently, the average number of MTs in the APs of group 2 is smaller; this can be observed in Fig. 8 for load equal to 0.8. If the objective, for a given particular system design, is to decrease the power consumption of a specific AP group, the F1, F1-R and EA policies are more appropriate, but loss probabilities become quite high for growing loads.
4. The random algorithm is probably the simplest one in terms of implementation. Indeed, the main algorithm characteristic is that it requires no state information. It must be noted that the performance of this algorithm is quite close to optimal. Indeed, for $\rho = \{0.2, 0.4\}$ the best performing algorithm improves energy consumption

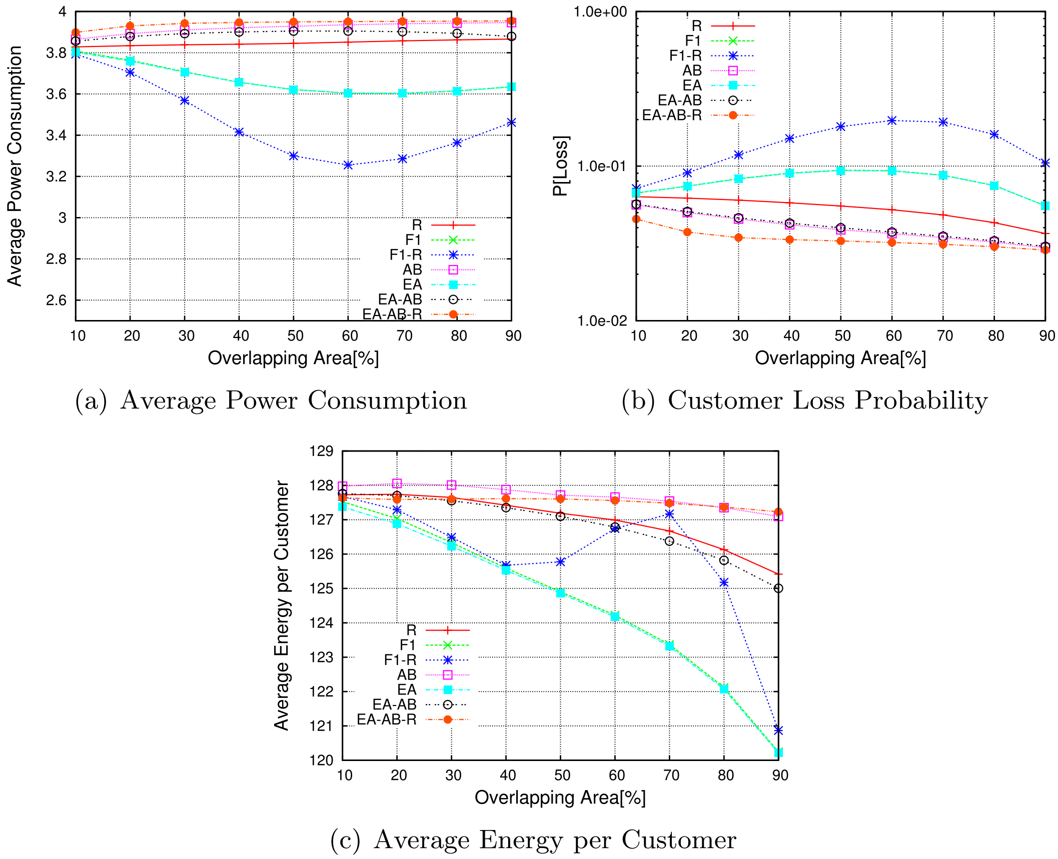

(a) Average Power Consumption (b) Customer Loss Probability

(c) Average Energy per Customer

Fig. 7. Set of metrics for the equal load case with $\rho = 0.8$.

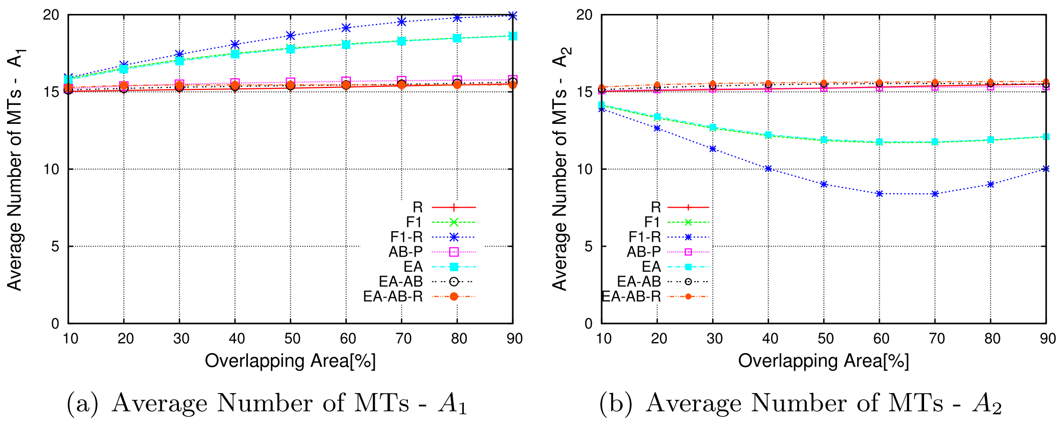

(a) Average Number of MTs - A_1 (b) Average Number of MTs - A_2

Fig. 8. Mean number of MTs associated with the two groups, $\rho = 0.8$.

with respect to the random (R) by less than 10%, at a price of some state information. For $\rho = \{0.6, 0.8\}$, the maximum improvement amounts to about 15%, but it comes with a large increase of loss probability.

5. The average energy per customer decreases for increasing load, due to a better utilization of the APs that always remain on, to guarantee coverage.

3.1.1. Lower capacity access points

In the derivation of previous results we considered groups of two APs, that can accept up to $K = 10$ MTs each. Since some of the previously described effects might be due to the quantization inherent in the small number of APs in each group, we now turn our attention to a case where quantization is lower: each AP can serve at most five

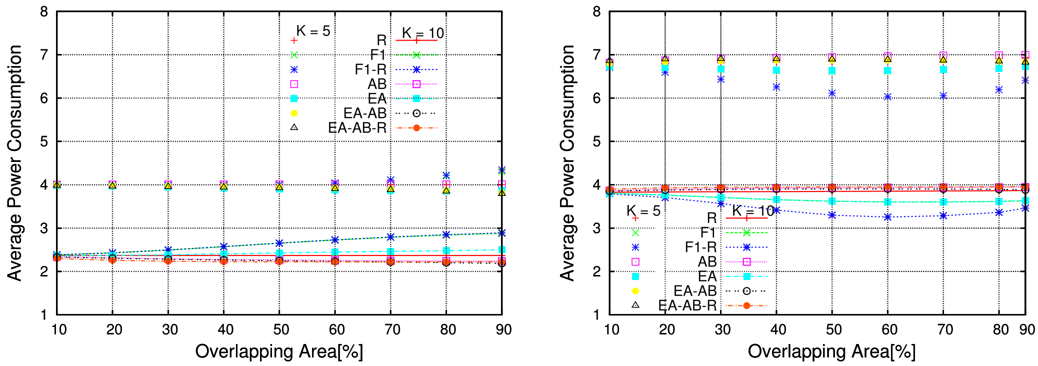

(a) $\rho = 0.4$ (b) $\rho = 0.8$

Fig. 9. Impact of parameter K on the policies' behavior.

users, and each group comprises four APs, thus leaving the maximum number of active users per AP group equal to 20, as in the previous case.

Fig. 9 shows the results for the average power consumption considering $\rho \in \{0.4, 0.8\}$ and $K \in \{5, 10\}$. Taking into account that the case $K = 5$ corresponds to doubling the number of APs with respect to the $K = 10$ case, and, thus, when all devices are powered, to double the consumption (since we assume that in both cases an AP consumes 1 W when on), no significant qualitative difference can be observed in the two cases. We can conclude that the effect of quantization on results is limited. The reduced granularity in the number of MTs served by each AP leads in fact to a small increase in efficiency, but the relative merits of the different CMN algorithms remain the same.

3.1.2. Including the energy cost of turning on an access point

The results presented so far assumed that the AP activations or deactivations do not incur energy cost. However, it is interesting to investigate whether a reasonable energy cost associated with the AP activation or deactivation may impact the comparison of the considered CMN algorithms. Here, we investigate only the association of an energy cost with the AP switch-on, neglecting the switch-off cost, which can be assumed to be significantly smaller.

To compute the energy cost of switching on an AP, we first derive the *switch-on rate*, i.e, the average number of times an AP is switched on in the time unit. Let us denote by R the *switch-on rate*. We can derive R from:

$$R = \sum_{\bar{s} \in S_{on}} \pi_{\bar{s}} \lambda_{\bar{s}},$$

with $S_{on} = \{\bar{s} | [\frac{u_1+c_1}{K} < n_1 \& u_1 + c_1 < N_1] \, || \, [\frac{u_2+c_2}{K} < n_2 \, \& u_2 + c_2 < N_2]\}$. The energy cost of switching on an AP is then derived as the product of R by the energy needed for each switch-on; the cost is, thus, equal to $R\epsilon$ and it is measured in Watts.

Initially, we only present results for the switch-on energy cost, without adding the power consumption cost value attached to the MTs connected to each access point. Fig. 10 shows the results for $\rho \in \{0.4, 0.8\}$ and $\epsilon = 10$ J. Assuming $\epsilon = 10$ J implies that the switch-on energy cost equals the energy cost of 10 s of normal operation. This may be reasonable, assuming that the AP requires a time of the order of 10 s to switch on. We can note that the switch-on energy cost is at most of the order of few

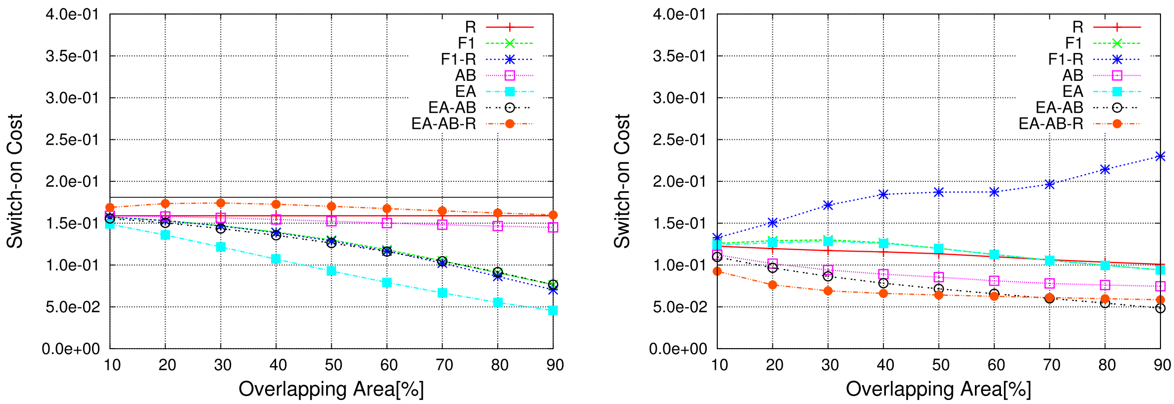

(a) $\rho = 0.4$ (b) $\rho = 0.8$

Fig. 10. switch-on cost (energy consumption due to APs activations) for cost of individual activation $\epsilon = 10$ J.

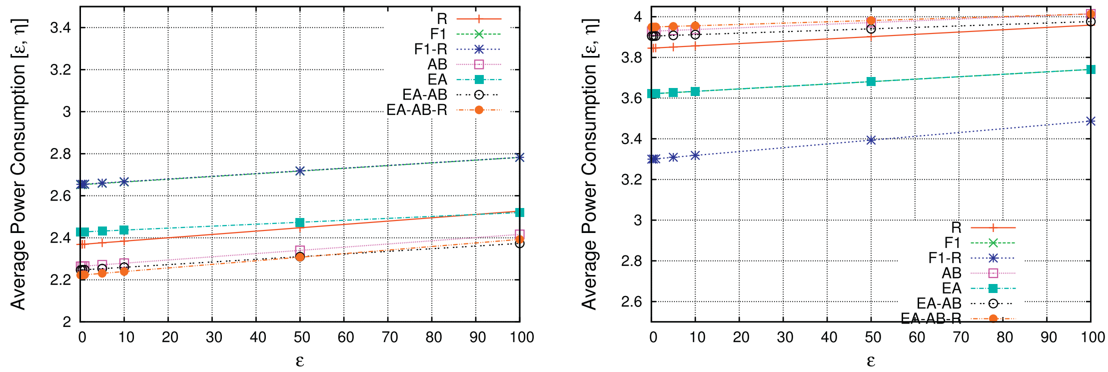

(a) $\rho = 0.4$ (b) $\rho = 0.8$

Fig. 11. Average power consumption for different values of ϵ and 50% of overlapping area.

percentage points of the cost required to power on the two APs that are necessary to provide coverage over the service area (one in each G_1 and G_2).

To explore the influence of different energy costs for the AP switch-on, we show in Fig. 11 the total energy consumption metric (adding the normal operation energy cost to the switch-on cost) for different values of ϵ, with A_0 equal to 50% of the total service area A. We can see that the impact of the switch-on energy cost is quite small at low as well as high load.

For all cases, the largest energy cost is due to the normal AP operation.

3.1.3. Three access points per group

As a variation of the basic scenario, we increase the number of APs in each group to 3, so that the total number of users served simultaneously by each AP group grows to 30. Fig. 12 shows results for $\rho = 0.6$. The algorithms' behavior does not change significantly with respect to the previous case. Further increasing the number of APs becomes critical from the point of view of the state space cardinality, but this situation can be handled with the approximations that we describe later in this paper.

3.2. Unbalanced load

We now consider the case in which the traffic loads on the two areas are different. The total traffic ρ_1 on A_1 is taken to be 20% greater than the traffic on A_2.

Fig. 13 presents the results for the average power consumption metric. Interestingly, since the first group is more loaded, the Fill G_1 first (F1) and Energy-aware policies tend to behave in a similar way for growing load (curves coincide at high load). Since the Random policy assignment does balance the associations when traffic is the same in the two areas, but not in the case of unbalanced traffic, the random policy consumes more than the energy-aware policies. The lower consumption visible in the figure at high load is the side effect of a much higher loss probability (not shown for brevity) exhibited by the Random policy in this scenario.

The total traffic ρ_1 on A_1 is now taken to be 40% greater than the traffic on A_2. Fig. 14 presents the results for the

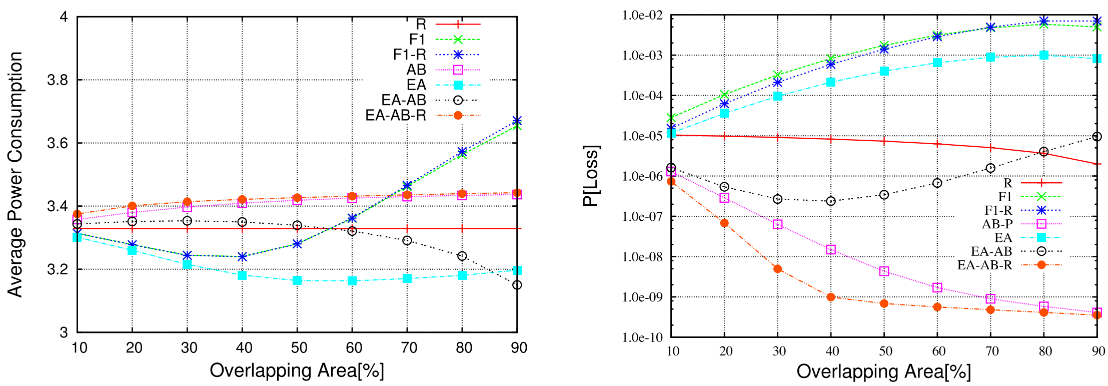

(a) Average Power Consumption (b) Customer Loss Probability

Fig. 12. Average power consumption and customer loss probability for $\rho = 0.6$ with maximum number of users per AP, $K = 10$ and number of AP per group, $n_i = 3$.

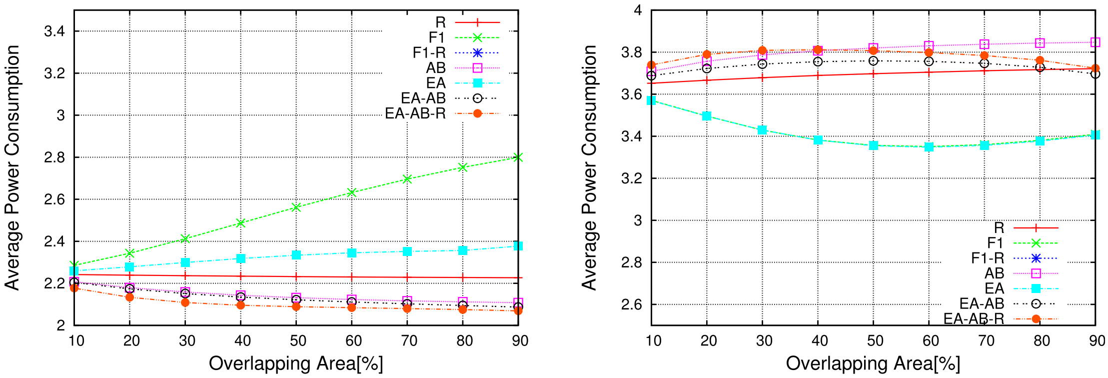

(a) $\rho_1 = 0.4$ (b) $\rho_1 = 0.8$

Fig. 13. Average power consumption for the inhomogeneous traffic to the stations, with traffic on A_1 20% larger than traffic on A_2.

average power consumption metric. For the low load case ($\rho = 0.4$), the policies have the same behavior as the previous case. However, for the high load scenario, the average power consumption values are more spread considering the energy-aware policies. Again, the random policy has a much higher loss probability measure.

4. Approximations

As we can see from (1), the cardinality of the state space S of the queuing system grows roughly with the fourth power of the maximum number of users that can be served by a group of APs (assuming that the two groups can serve about the same number of users). If the number of AP groups is m, and each group can serve at most N users, the state space cardinality is of the order of N^{2m}. The Markovian model that underlies the queuing system suffers from a combinatorial explosion of the number of states for both growing number of users and growing number of AP groups.

Dealing with very large state spaces implies both technical problems in the implementation of the model solution, and numerical errors in the computation of the performance measures of interest. In this section, we present approximate models that provide accurate estimates of the performance metrics, while generating state spaces of much smaller size than the models discussed so far, and even allowing a closed form expression for the limiting probabilities.

In the following, the description and the evaluation of the approximate models focus only on the Random CMN algorithm for two reasons. First, the approximations are simpler to describe in this case, and, second, the R policy seems quite a reasonable compromise between performance and simplicity of implementation. The extension to the other CMN algorithms discussed in the previous section is tedious, but straightforward.

4.1. Single queue approximation

As a first step, we develop an approximation based on considering just one of the two queues in our system, say queue 1. The idea behind this approach is to approximate the customer arrival rate γ_1 at queue 1 from the area

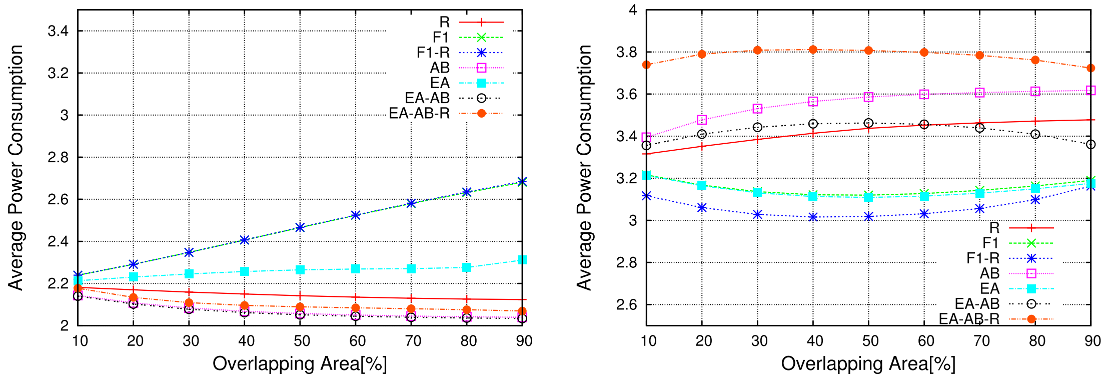

(a) $\rho_1 = 0.4$ (b) $\rho_1 = 0.8$

Fig. 14. Average power consumption for the inhomogeneous traffic to the stations, with traffic on A_1 40% larger than traffic on A_2.

$A_0 = A_1 \cap A_2$, independently on the state of the other queue (queue 2). The performance metrics of interest are then computed for queue 1, independently from the state of queue 2.

The customer arrival rate in A_0 is equal to $\alpha_0\lambda$. According to the Random CMN algorithm, this rate is equally split between the two queues, so that queue 1 receives from this area a customer flow at rate $\alpha_0\lambda/2$. However, queue 1 also receives all customers that arrive in area A_0 when queue 2 is full, so that we must also account for an additional customer arrival rate equal to $P[full_2]\alpha_0\lambda/2$, where $P[full_2]$ is the probability that queue 2 is full. So, we can write:

$$\gamma_1 = (1 + P[full_2])\frac{\alpha_0\lambda}{2}. \tag{10}$$

The total arrival rate at queue 1 is then:

$$\lambda_1 = \alpha_1\lambda + (1 + P[full_2])\frac{\alpha_0\lambda}{2}. \tag{11}$$

Assuming that the two queues are equal (same number of APs, same size of the served area, same maximum number of users per AP), or very similar, we can approximate the value of $P[full_2]$ with the value $P[full_1]$. If the parameters of the two queues are indeed identical, we only introduce an error deriving from the fact that we are assuming independence in the behavior of the two queues. If the parameters of the two queues differ, an additional error is introduced by assuming $P[full_2] = P[full_1]$.

For the model solution, in this case we must use an iterative algorithm, because of the circular dependence of the loss probability on the arrival rate, and of the arrival rate on the loss probability. At each step of the iteration, the probability distribution vector $\vec{\pi}$ is computed, as well as $P[full_1]$, for a given value of the customer arrival rate λ_1. The new value of λ_1 is then computed from $P[full_1]$, and a new iteration is run. The algorithm stops when either a maximum number of iterations is reached, or the relative error between the values of $P[full_1]$ at two consecutive steps of the iteration is smaller than a given tolerance. The performance measures of interest are computed from the final probability distribution vector.

Note that the queue we are using in this approximation is an $M/M/N_1/0$, with arrival rate equal to λ_1, so that:

$$\pi_{i,1} = \frac{\left(\frac{\lambda_1}{\mu}\right)^i \frac{1}{i!}}{\sum_{k=0}^{N_1}\left(\frac{\lambda_1}{\mu}\right)^k \frac{1}{k!}} \tag{12}$$

for $i = 0, \ldots, N_1$, and

$$P[full_1] = \pi_{N_1,1} \tag{13}$$

Hence, the iterative algorithm is extremely fast.

This approach can be easily extended to cases comprising more than two queues. Numerical results for this approximation are presented in Section 4.3.1, and validated against the exact model and simulation. As expected, this single queue approximation yields very accurate results, especially when the queues are equal.

4.2. Multi-queue approximations

The approximation procedure described above can be easily improved, by considering all the queues in the system, but still assuming independence in their behaviors. We can thus model each queue separately, computing the customer arrival rate λ_i at each queue. The computation of λ_1 in the case of 2 queues is thus as in (11), with $P[full_2]$ computed in this case from the analysis of queue 2 in isolation. Also in this case, for the computation of the performance metrics, an iterative algorithm must be used. At each step of the iteration, the probability distribution vectors for the two queues are computed from their analytical expression, as well as $P[full_1]$ and $P[full_2]$, using for the customer arrival rates the values of λ_1 and λ_2. The new values of λ_1 and λ_2 are then computed from $P[full_1]$ and $P[full_2]$, and a new iteration is run. As before, the algorithm stops when either a maximum number of iterations is reached, or the relative error between the full queue probability values at two consecutive steps of the iteration are smaller than a given tolerance.

Also in this case, the approximation can be easily extended to cases comprising more than two queues. Numerical results for this approximation are presented in Section 4.3.2, and validated against simulation. As expected, this approximation yields better results than the single-queue approximation, especially when queues are different.

4.3. Validation of approximations

4.3.1. Single-queue systems

We start by validating the proposed single-queue approximation, considering queuing systems where all queues are subject to the same load, and the Random CMN algorithm is used. We start with systems that can be solved with the detailed model, so that the approximate results can be compared against the results of the original model; we then let the number of users and the number of areas grow.

Let us start with two overlapping areas, two APs per group and 10 MTs on each AP at most. To assess the accuracy of the approximate solution, we evaluate the relative error, δ, between the approximate and exact computation of the average power consumption. Table 3 reports, for $\rho = \{0.4, 0.6, 0.8\}$, the minimum and maximum values of δ observed when varying the fraction of overlapping area between the service areas A_1 and A_2. Relative errors are very small.

For the cases with a larger number of users or larger number of areas, since the computation of the exact solution is not feasible, the approximation is compared against results obtained by simulation. Letting the maximum

Table 3
Relative error δ (2 APs per group, $K = 10$ MTs, 2 areas).

ρ	Minimum δ	Maximum δ
0.4	8.445375e−6	1.689047e−5
0.6	3.014154e−5	2.529892e−3
0.8	1.990460e−3	7.565440e−3

number of users that can be served by each queue grow, considering 40, 60, 80, 100 users on each area, $K = 10$ and $\rho = \{0.4, 0.8\}$, Table 4 shows the mean energy consumption of group 1, computed by the approximate solution and through simulation; for simulation results we report the confidence interval at 95% confidence level obtained by repeating 10 independent simulation runs. Values in boldface refer to the cases in which the approximate solution falls out of the confidence interval obtained by simulation. When the load is low the approximate solution is very accurate and always falls in the confidence interval; when the load is high, the approximate solution that assumes independence among queues slightly overestimates the energy consumption. Still, estimations are pretty accurate in most of the cases.

In what follows, we consider the case of a queuing system comprising a larger number of queues, i.e., several AP clusters which serve multiple overlapping areas. For the numerical results, consider the case of four groups of APs (four queues in the model), that form one overlapping area that is common to all the four areas. Each group includes two APs, and each AP can serve up to 10 MTs. Table 5 shows the relative errors between the single-queue approximate model and simulation. Also in this case, relative error values remain very small, indicating that the proposed approximation is accurate even in more complex cases, which would be problematic to handle with an extension of the exact model that we presented in the first part of this paper.

Let us now turn our attention to the case of a WLAN with eight groups of APs, and eight overlapping service areas. Each group includes two APs and each AP can serve up to 10 MTs. Table 6 shows the relative errors between the single-queue approximate model and simulation. Again the relative error values remain small.

4.3.2. Multi-queue systems

Let us now consider the case of unbalanced load, due to different arrival rates on each area. We focus on the basic scenario, with two groups of two APs each, and each AP serving at most 10 users; the probability α_1 that an arriving customer is constrained to proceed to station 1 is 25% greater than the probability α_2.

Although the single-queue approximation can also be applied in this scenario, the multi-queue approximation is expected to provide better accuracy. Indeed, while the single-queue approximation, for $\rho = 0.4$, increases both the maximum and minimum errors to 5.491816e−3 and 4.923707e−2, respectively, the 2-queue approximation leads to errors which are two orders of magnitude smaller, as we can see in Table 7. Note that this improved accuracy is paid with a negligible increase of the solution complexity.

Table 8 shows the relative errors between the single-queue approximate model and simulation for the case of four areas. The relative error remains small, indicating that the proposed approximation performs well.

4.3.3. Non-exponential association times

Finally, in this section we consider the case of systems with non-exponential association times. The distribution

Table 5
Relative error δ (20 MTs, 4 areas).

ρ	Minimum δ	Maximum δ
0.2	4.966846e−6	1.689021e−4
0.4	1.224378e−4	1.724102e−3
0.6	1.005664e−3	1.675102e−2
0.8	7.299118e−3	1.319239e−2

Table 6
Relative error δ (20 MTs, 8 areas).

ρ	Minimum δ	Maximum δ
0.2	6.954102e−5	2.943501e−2
0.4	1.493859e−3	4.938786e−2
0.6	2.987414e−2	9.235088e−2
0.8	3.345820e−2	9.756801e−2

Table 4
Approximate solution and the simulation confidence intervals, case of two groups.

APs	A_0 (%)	$\rho = 0.4$		$\rho = 0.8$	
		Approx.	Conf. inter.	Approx.	Conf. inter.
4	20	2.055	[2.053, 2.058]	3.558	[3.549, 3.560]
	40	2.055	[2.051, 2.056]	**3.571**	[3.564, 3.570]
	60	2.055	[2.052, 2.058]	**3.586**	[3.576, 3.584]
	80	2.055	[2.051, 2.056]	**3.602**	[3.581, 3.592]
6	20	2.853	[2.849, 2.872]	5.187	[5.172, 5.203]
	40	2.854	[2.838, 2.859]	5.198	[5.186, 5.212]
	60	2.854	[2.848, 2.862]	5.210	[5.195, 5.227]
	80	2.853	[2.838, 2.858]	**5.223**	[5.191, 5.213]
8	20	3.650	[3.642, 3.684]	6.807	[6.797, 6.828]
	40	**3.651**	[3.614, 3.641]	6.813	[6.774, 6.818]
	60	3.651	[3.630, 3.672]	6.822	[6.807, 6.852]
	80	3.650	[3.625, 3.651]	6.834	[6.824, 6.865]
10	20	4.451	[4.444, 4.470]	8.420	[8.385, 8.430]
	40	**4.452**	[4.425, 4.450]	**8.423**	[8.381, 8.417]
	60	4.452	[4.428, 4.463]	8.429	[8.385, 8.439]
	80	4.451	[4.432, 4.471]	8.439	[8.431, 8.471]

Table 7
Relative error δ – multi-queue systems (20 MTs, 2 areas).

ρ	Minimum δ	Maximum δ
0.2	4.981047e−6	2.492212e−5
0.4	4.906312e−5	4.524259e−4
0.6	6.060947e−4	4.891496e−3
0.8	2.269933e−3	1.149481e−2

Table 8
Relative error δ – multi-queue systems (20 MTs, 4 areas).

ρ	Minimum δ	Maximum δ
0.2	4.730487e−4	2.981811e−3
0.4	5.875441e−3	5.439721e−2
0.6	1.028257e−2	7.844002e−2
0.8	4.151188e−3	1.149481 − 2

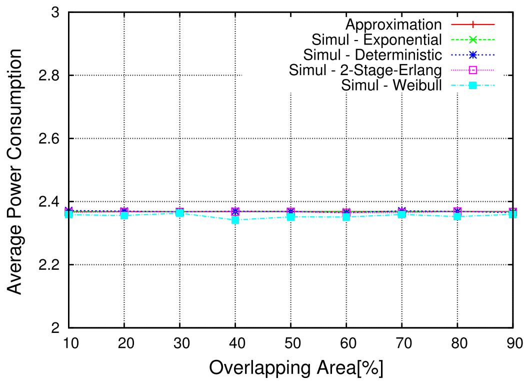

Fig. 15. Average power consumption for different distributions of the association times.

of association times depends on many factors, such as the scenario, the user behavior, the environment and it is, in general, difficult to derive. The objective of this section is to evaluate the accuracy of the predictions we obtain with our approximate model when various distributions are considered. Good approximations are indications of a high degree of representativeness of our approximate results in different possible scenarios. In particular, we compare the results obtained from the proposed model against simulation results obtained by considering systems with different distributions and the same mean value. We consider again the case of two areas, $K = 10$, and system load $\rho = 0.4$. In Fig. 15 the results are plotted for the cases of association times distributed according to exponential, deterministic, Erlang (2 stages) and Weibull distributions. The results are clearly almost insensitive to the distribution of the association time. Thus, we can consider our model and our results representative of more general cases.

5. Conclusions

In dense WLANs, where a large number of APs is deployed to provide large capacity, the use of sleep modes for the APs can make the network energy efficient. The idea is to use sleep modes for a portion of the APs when the number of mobile terminals associated with the network is small and not all the capacity is needed.

Based on this idea, in this paper, we presented a set of algorithms to make the capacity adaptive to the number of associated terminals and, in particular, for the association of active users with APs, with the final objective of reducing the amount of energy consumed in dense, centrally managed, WLANs. We investigated the energy-performance trade-off with a model based on two coupled queues. We also introduced computationally efficient and accurate approximations for the analysis of large systems.

Two main lessons can be learn from the results. First, the amount of energy that can be saved is quite considerable, especially at low loads; this suggests that approaches like those presented in this paper are interesting for large enterprise networks that include large numbers of APs. Second, results indicate that differences in energy efficiency among CMN algorithms remain limited, so that the use of simple approaches seems reasonable and, thus, even a limited state information is sufficient to make the algorithm effective.

Acknowledgment

The research leading to these results has received funding from the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement no. 257740 (Network of Excellence "TREND").

References

[1] http://www.metageek.net/products/inssider.
[2] J. Lorincz, A. Capone, D. Begusic, Optimized network management for energy savings of wireless access networks, Computer Networks Journal 55 (3) (2011).
[3] A.P. Jardosh, K. Papagiannaki, E.M. Belding, K.C. Almeroth, G. Iannaccone, B. Vinnakota, Green WLANs: on-demand WLAN infrastructures, Mobile Networks and Applications 14 (2009) 798–814.
[4] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, A simple analytical model for the energy-efficient activation of access points in dense WLANs, in: e-Energy 2010, Passau, Germany, April 2010.
[5] A. Wierman, L.L.H. Andrew, A. Tang, power-aware speed scaling in processor sharing systems, in: Infocom 2009, Rio de Janeiro, Brazil, April 2009.
[6] A. Gandhi, V. Gupta, M. Harchol-Balter, M.A. Kozuch, Optimality analysis of energy-performance trade-off for server farm management, in: Performance 2010, Namur, Belgium, November 2010.
[7] A. Gandhi, M. Harchol-Balter, I. Adan, Server farms with setup costs, in: Performance 2010, Namur, Belgium, November 2010.
[8] M. Ajmone Marsan, L. Chiaraviglio, D. Ciullo, M. Meo, Optimal energy savings in cellular access networks, in: ICC'09 – Green Communications Workshop, Dresden, Germany, August 2009.
[9] F. Richter, A.J. Fehske, G.P. Fettweis, Energy efficiency aspects of base station deployment strategies for cellular networks, in: Proceedings VTC'09, September 2009.
[10] M. Ajmone Marsan, M. Meo, Energy efficient wireless internet access with cooperative cellular networks, Elsevier Journal Computer Networks 55 (2) (2011) 386–398.
[11] K. Son, B. Krishnamachari, SpeedBalance: speed-scaling-aware optimal load balancing for green cellular networks, in: Proc. IEEE Infocom 2012, Orlando, Florida, USA.
[12] M. Herlich, H. Karl, Reducing power consumption of mobile access networks with cooperation, in: e-Energy 2011, Columbia University, New York, 2011.

Ana Paula Couto da Silva received the B.Sc degree in Computer Science from Federal University of Juiz de Fora, Brazil, in 1999 and the M.Sc and D.Sc degrees in Computer and System Engineering both from Federal University of Rio de Janeiro, Brazil, in 2001 and 2006 respectively. She was with IRISA-Rennes, France, in 2005 and 2007 and with Politecnico di Torino, Italy, in 2008. Since February 2009, she is with Computer Science Department at Federal University of Juiz de Fora, Brazil, as Assistant Professor. Her areas of interest are in the field of: modeling and analysis of computer systems, reliability analysis and P2P application analysis.

Michela Meo graduated from Politecnico di Torino with a summa cum laude degree in Electronic Engineering in July 1993. In July 1997 she obtained the Ph.D. degree at Politecnico di Torino in Telecommunications Engineering. In November 1999, she became an Assistant Professor at Politecnico di Torino; and since November 2006, she is Associate Professor. She coauthored more than 150 papers, in international journals or conferences. She edited five special issues of international journals, including ACM Monet and Performance Evaluation Journal. She was general chair and program chair of ACM International Workshop on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM); she was also program chair of the IEEE International Workshop on QoS in Multiservice IP Networks (QoS-IP), the 1st IEEE International Workshop on Mobile Vehicular Networks (MoVeNet) and of the 14th IEEE Symposium on Computers and Communications (ISCC); and she was in the program committee of about 80 international conferences, including Sigmetrics, Infocom, ICC and Globecom. She was TPC co-chair of a workshop titled "Communications and Control for Sustainable Energy Systems: Green Networking and Smart GridsG held in conjunction with IEEE Infocom 2012.

Marco Ajmone Marsan is a Full Professor at the Electronics and Telecommunications Department of Politecnico di Torino, in Italy, and a part-time Chief Researcher at IMDEA Networks in Madrid, Spain. He holds degrees in Electronic Engineering from the Politecnico di Torino and the University of California, Los Angeles. He was at Politecnico di Torino Electronics Department from November 1975 to October 1987 - first as a researcher and then as an Associate Professor. He was a Full Professor at the University of Milan Computer Science Department from November 1987 to October 1990. From September 2002 to March 2009 he was the Director of the Institute for Electronics, Information and Telecommunications Engineering of the National Research Council. During the summers of 1980 and 1981, he was with the Research in Distributed Processing Group, Computer Science Department, UCLA. During the summer of 1998 he was an Erskine Fellow at the Computer Science Department of the University of Canterbury in New Zealand.