MORFEO: Saving Energy in Wireless Access Infrastructures

Karina Gomez[†], Cigdem Sengul[‡], Nico Bayer[§], Roberto Riggio[†], Tinku Rasheed[†], Daniele Miorandi[†]
[†]CREATE-NET, Trento, Italy; *Email: first.secondname@create-net.org*
[‡]Oxford Brookes University, London, UK; *Email: csengul@brookes.ac.uk*
[§]Telekom Innovation Laboratories, Darmstadt, Germany; *Email: nico.bayer@telekom.de*

***Abstract*—Energy efficiency is acknowledged as a pivotal issue for a sustainable development of wireless networking technologies. Traditionally, most works in the area focused on the user equipment, where battery duration represents a key asset. However, as the smartphone and tablet revolution fuels a massive deployment of wireless networks, often in the form of WiFi hotspots, more and more attention is expected to be devoted to the energy-efficient management of wireless access infrastructure. These networks tend to be dense and over–provisioned, which in time leads to significant energy wastage in off–peak conditions. In this paper, we present *MORFEO* a flexible energy–saving decision algorithm to tune the energy consumption of a wireless infrastructure to the actual network conditions in terms of both user density and traffic patterns. Experimental results from a real-life deployment shows that our solution can deliver significant energy savings with minimal degradation in terms of the quality of service provided.**

I. Introduction

Energy efficiency in ICT infrastructures is becoming a top priority for industries, governments and scientific communities alike. A significant portion of the energy consumed in current wireless access infrastructure is not effectively used due to an overprovisioning of resources mandated by the necessity to support peaks in end–users' demand and by the lack of proper network management approaches able to orchestrate the usage of resources in an energy-optimal manner. So far, wireless network architectures and protocols have accounted for the energy efficiency of end–user terminals, which are typically mobile, battery-powered and, thus, have strict constraints in terms of power consumption. On the other hand, less attention has been paid to the energy efficiency of the wireless infrastructure; the rationale being that, as the infrastructure is directly attached to the power distribution grid, energy consumption is an issue of little interest and impact.

However, the recent smartphone and tablet revolution is changing this picture. For instance, the number of WiFi hotspots deployed is increasing exponentially. The trend may be strengthened as WiFi offload technology (which automatically redirects the traffic generated by smartphones and mobile devices from 3G to WiFi networks) gets deployed at scale. Typically, WiFi access points are operated at full power, given that network operators are generally reluctant to shut–down portions of their networks in order to preserve full availability. However, "always available" does not need to mean "always fully powered". Our work starts from the firm belief that there is significant space for improving the energy efficiency of wireless access infrastructures, while at the same time preserving the quality of service delivered to end users.

In our previous works [1], [2], [3], we analysed and experimentally measured the energy consumption of several wireless access devices, including WiFi, WiMAX, and 3G devices. In all cases, it turned out that a significant fraction of the energy consumed by these devices is not traffic dependent. More specifically, we found that injecting traffic in the network until the saturation point results in only a 20% increase in power consumption with respect to the power consumed when there is no traffic in the network besides the regular signalling. In order to achieve significant energy savings, it is imperative to find the means for reducing the power consumed when little or no traffic is present. Similar conclusions are drawn in [4], which studies the energy consumption of 125 WiFi access points (APs) deployed in a typical office environment. APs are densely deployed (one every $5m$) to provide sufficient capacity for four very close users using voice, data and multimedia applications simultaneously. Measurements showed that these APs spend $\approx$ 20 to 80% of their time idling and consume 8.76 MWh of energy per year in total.

In this paper, we propose and experimentally validate a monitoring and control framework that can support different energy-saving strategies in heterogeneous wireless networks. The energy-saving strategies are implemented based on the switching of the operating modes of wireless access devices (i.e., whether they need to remain on, or they can be turned off, or switched to a low-power mode etc.). To make appropriate energy–saving decisions, the framework takes into account the network traffic and usage scenarios, and also by accounting for the cost of switching of the operating modes. The main contributions of this paper are as follows:

- A practical energy–saving decision algorithm, named *MORFEO* which exploits an energy monitoring and control framework. *MORFEO* adapts the energy consumption of a wireless infrastructure to the network conditions in terms of both current user density *and* traffic patterns.

- A measurement–based evaluation of *MORFEO* in a WiFi testbed deployed at Telekom Innovation Laboratories in Berlin, Germany. The implementation and operation of *MORFEO* on a testbed confirms that our framework is flexible enough to support practical use–cases. The experiment results show that *MORFEO* considerably lower the energy footprint of the network with very limited performance penalty for end users.

The paper is structured as follows. Section II discusses the related work. In Section III, we describe the design details of the energy monitoring and control framework and the energy–saving decision algorithm, *MORFEO*. Section IV presents an analysis and measurement of the operation mode switching times, which play an important role in *MORFEO* operation. The experimental settings and measurement results are discussed in Section V. Finally, in Section VI, we draw our conclusions and point out future research directions.

II. Related Work

In this section, we present an overview of the related work on energy efficiency improvements proposed for cellular and enterprise WiFi networks. The majority of the works investigate the most efficient ways to turn off cells/networks or put them in sleep mode, while maintaining good quality of service [5], [6], [7], [8], [9], [10]. In [5], a cooperative approach between two operators, offering service over the same area, is proposed to reduce energy consumption. Under high traffic conditions, both networks are used, while under low traffic conditions, one network is switched off and the second network supports all the traffic in a transparent mode to the users. A similar approach is proposed in [6] using heterogeneous 2G/3G networks of the same operator. The authors introduce algorithms for finding the optimal traffic allocation in cooperative 2G/3G networks in order to power on/off an entire system (2G or 3G) for high or low traffic scenarios, respectively. In [7], [8], [9], [10] several models for the energy–aware management of cellular access networks are introduced. Typically, switching cells off requires shifting the traffic to the cells that remain active in order to guarantee the service availability over the whole area.

Similar sleeping approaches also exist for WLANs [11], [12]. In [11], an on–demand strategy is proposed to power on or off resources in high–density WLANs. More specifically, the volume and location of the user demand is considered as an input to power on/off WLAN access points (APs) dynamically. A similar approach is also taken in [12], which uses the analyses from traffic in different hours, days and weeks in order to optimize the management of on/off states.

Other approaches to save energy include exploiting the heterogeneity of the access technologies and the interaction with wired networks, and adapting transmission power based on energy and coverage trade-offs [13]. For example, one of the proposals is to decrease the number of active APs by increasing the transmission power and relay messages using ad hoc networking to increase coverage. In [14], the authors describe a set of challenges to minimize power consumption of the whole cellular network architecture and guarantee the QoS at the same time. The future solutions discussed by authors are mainly focusing on energy metrics, energy efficient architectures, multi-hop routing and frequency management. In [15], the authors propose the deployment of small, low-power base stations together with conventional macro sites. The same approach is used in [16] for homogeneous and heterogeneous wireless networks, where the utilisation of low-power micro sites instead of macro sites is considered to enhance throughput, energy efficiency and network coverage.

Finally, in [17], the transmission power of the base stations are optimized to reduce the total energy consumption of the network. The optimization technique takes into account (i) the effect of shadowing, (ii) the presence of thermal noise, and (iii) the impact of the transmission power of base stations on the coverage and the capacity. Similar techniques are also used in [16], [12].

From the literature, it is evident that those schemes that turn off parts of the infrastructure, or put it in sleep mode, have significant potential in terms of reducing energy consumption in wireless networks. In our work, we add to this understanding and show its feasibility in a real network by taking into account hardware and software limitations, the time and energy it takes to switch on and off interfaces and devices, and their impact on network and energy consumption performance.

III. *MORFEO*: Energy–saving Decision Algorithm for Wireless Networks

Wireless networks easily become over–dimensioned during periods of low traffic demands, which directly depend on the deployment area and the network usage guided by human mobility and traffic requirements. For instance, these periods are common at night time and the weekend for enterprise deployments. In order to adapt the network capacity and topology to the actual traffic demands of the users, we propose *MORFEO*, an energy–saving decision algorithm for wireless infrastructure networks. In this section, we first present the energy monitoring and control framework. In our framework, the wireless access devices (WADs) can be in different operation modes, with varying levels of energy consumption. We describe these modes, and then explain how *MORFEO* adapts operation mode per device under different traffic and coverage conditions.

A. Energy Monitoring and Control Framework

In [18], we presented a preliminary framework and hardware–software prototype for enabling the implementation of energy saving approaches. We extended this work to

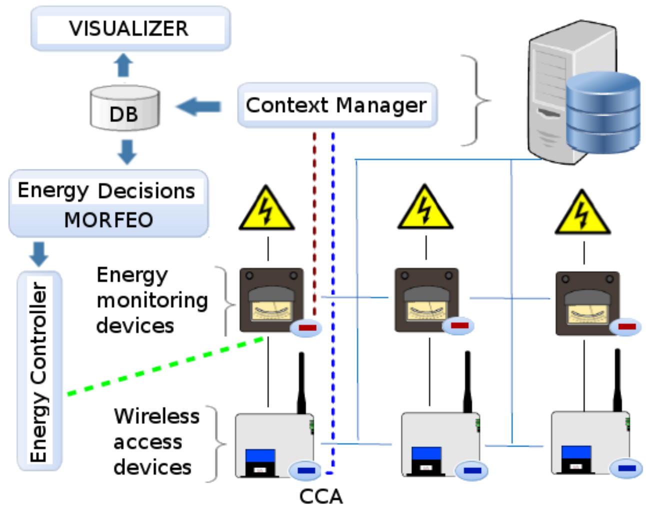

Figure 1: Our energy monitoring and control framework. MORFEO is an example of energy-decision algorithms used to drive energy-saving actions in a wireless infrastructure network.

a more complete and practical framework, depicted in Fig. 1, to support energy–saving algorithms. The main components of our energy monitoring and control framework are:

- **Context Manager (CoMa)** is responsible for gathering relevant information, such as network and power utilization statistics, from the WADs, energy monitoring devices, mobile devices as well as external databases (see [19] for details). The statistics are collected using a Context Collecting Agent (CCA) installed in each WAD and stored in the CoMa database. These statistics mainly include the WAD configurations and locations, the number of connected users, amount and type of traffic for each wireless interface, power consumption of the WAD and other network performance indicators.
- **Energy Decisions (ED)** makes the necessary decisions to switch the operation modes, which will be explained in the next section, and schedules these actions to be executed by the Energy Controller. ED can considered as a repository of energy-efficient algorithms. In this paper, we present *MORFEO* as an example algorithm for energy–saving decisions in a wireless infrastructure network.
- **Energy Controller (EC)** contains the logic for monitoring the energy consumption of the WADs and writing these statistics to CoMa database. EC also triggers the actions scheduled by the ED. For example, the ED can schedule turning off the WADs. The components required for implementing the EC can be either a commercial platform [20] or based on an open–source platform [18]. The actions supported by EC are (i) switching On/Off the device, (ii) switching On/Off the wireless interfaces and (iii) changing the transmission power of each wireless interface.
- **Visualizer** displays a graphical interface of the actions scheduled by the ED and the network statistics about power consumption, users and traffic load collected from the CoMa database (see [21] for details).

WADs and user terminals also play an important role in this framework. User terminals can be static, nomadic or mobile and they are associated to the WADs via a single wireless hop. They need to send feedback about the network performance to the WADs to be able to make energy efficiency decisions without compromising service quality. A special applications software (APPs) is installed in the user terminal for sending statistics to CoMa.

Based on our measurements, we identify four modes of operation for WADs based on their energy consumption:

- **Full-power Mode (PM):** The WAD and its interfaces are switched on and are operating with highest transmission power level (Tx). This mode provides full coverage and capacity.
- **Active Mode (AM):** The WAD and its interfaces are always on and operating with the default Tx, which is lower than the highest Tx.
- **Partial or Sector Sleep Mode (PSM or SSM):** The WAD is powered on but one or more sectors or interfaces are switched off. Note that similar power-saving modes are already implemented for client-side devices. Here, we introduce this mode for WADs. In this mode, if there are no associated users, sectors or interfaces can be turned off for energy savings. When there are associated users but no traffic, a configurable duty-cycle can be used to periodically turn off the sectors or interfaces. However, the duty-cycle should be automatically and immediately deactivated, when traffic is detected.
- **Off Mode (OM):** The WAD is turned off and only the energy monitoring device is powered. The energy consumed for the energy monitoring device should be always smaller than the power consumed by the WAD in Active Mode. The Off Mode takes advantage of the network over–dimensioning and overlapping coverage provided by different WADs. The goal is to use Off Mode for keeping a minimal set of devices in Active Mode in order to provide full coverage and required capacity with minimal energy consumption.

MORFEO decides which operation modes provide high energy-efficiency and good coverage and capacity based on several context and measurement-based rules. In the next section, we explain the design of *MORFEO* in detail.

B. *MORFEO Design*

MORFEO operates in three parts: *Initialisation*, *Reactive updates*, and *Correction*. In the Initialisation part, *MORFEO* takes decisions about the most appropriate role for each WAD in a *proactive manner* based on the network deployment and context information.

Then, during *Reactive updates*, the statistics about the current network conditions, such as the number of users and the amount of traffic, are collected from CoMa and analysed. Based on this analysis, *MORFEO* decides the most appropriate operation mode for each WAD. To this end, *MORFEO* divides time into decision slots to ensure there is enough time for the CoMa to collect statistics, and also for the EC to switch the WADs to the different operation modes based on the energy-saving decisions by *MORFEO*.

Finally, *MORFEO* switches to the *Correction* part, which takes place after one time-slot duration. Here, if the network performance suffers degradations and the QoS cannot be guaranteed for the current traffic load, the *Correction* algorithm reverts to decisions taken in *Reactive updates*. Also, if there is a possibility to save more energy, *Reactive updates* may be called to update the topology to put more nodes in low-power operation modes. Hence, after its first run, *MORFEO* stays in *Correction* and uses *Reactive updates* as needed and thus, adapts the operation modes, and hence, the energy consumption to user demand.

In the following, we present each of the three parts in more detail.

Initialisation. In this part, *MORFEO* groups nodes into three different groups: *sleep candidates*, *head candidates*, and *special candidates*. More specifically, if the energy consumption can be reduced by switching certain WADs into a low power mode, i.e., Off or, Partial or sector sleep mode, these nodes are marked as *sleep candidates*. Also, the WADs, which can take over the load of other WADs are identified (e.g., due to overlapping coverage) and these nodes are marked as *head candidates*.

Finally, based on the context, under special conditions, a set of nodes may need to be always in Active mode. These conditions mainly appear in scenarios where high amount of traffic is expected but the actual behavior of the expected traffic may not be predicted (e.g., a stadium, where the traffic load is expected only on special events such as concert or soccer matches). In this case, such nodes are marked as *special candidates*.

To make these decisions, *MORFEO* uses different types of information, such as deployment information, which includes the network topology, coverage and network access technologies. Also, the energy cost analysis of the different operating modes of the WADs and the limitations in switching and sojourn times in different operating modes based on the deployed hardware plays an important role in these decisions. In addition, *MORFEO* takes into account the context information, which includes the application scenario, predicted traffic load, and QoS constraints.

Reactive updates. In this part, *MORFEO* evaluates several conditions for each WAD based on network traffic and topology and decides on the operation modes. To make these decisions, *MORFEO* evaluates several conditions for each WAD based on network traffic and topology. We define these conditions as follows:

Algorithm 1 *Reactive updates*

```
1: for all Sleep Candidates ≠ Off Mode do
2:     if WAD in Zero Condition then
3:         WAD → Off Mode
4:     else if WAD in Idle Condition then
5:         WAD → Partial/Sector Sleep Mode
6:     else if WAD in Active Condition then
7:         if ∃ Head Candidate and low traffic then
8:             handover users to Head Candidate
9:             Head Candidate ← Head Condition
10:            WAD → Sleep Candidates
11:        end if
12:    end if
13: end for
14: for all Head Candidates do
15:    if WAD in Head Condition and coverage adaptation needed
       then
16:        WAD → Full-Power Mode
17:        Tilt or Power Adaptation
18:    end if
19: end for
```

- **Active Condition (AC):** The WAD has associated clients sending or receiving traffic.
- **Idle Condition (IC):** The WAD has associated clients but its clients are not sending or receiving traffic.
- **Zero condition (ZC):** The WAD has no associated clients.
- **Head Condition (HC):** The WAD can take over the load of other WADs that are switching to Off mode.

Based on these conditions, *MORFEO* uses the Algorithm 1 to assign operation modes to each WAD in different groups (i.e., sleep candidates and head candidates).

Specifically, the current traffic situation of the sleep candidates is analyzed and the WAD is switched to Off Mode if it is under *Zero Condition*. If the WAD is under *Idle Condition*, it is switched to Partial/Sector Sleep Mode. Otherwise, if the WAD is under *Active Condition*, the procedure for balancing the network load is applied. Here, the users are handed over to the head candidates if it is possible. The WAD is added back to the set of sleep candidates to check the possibility of switching to Partial/Sector Sleep or Off Modes (line 10, Algorithm1). Finally, all the WADs in the *Head Condition* are switched to Full-power mode only if coverage adaptations are needed. According to the type of technology, tilting optimization or increasing the transmission power may also be used to extend and improve the coverage area.

Correction. After reactive updates, and one time-slot duration, *MORFEO* transitions to *Correction*, which is detailed in Algorithm 2. Here, *MORFEO* evaluates the network performance and if the QoS constraints are not satisfied, the WADs that are in low power modes are switched to Active Mode. Otherwise, *MORFEO* calls *Reactive updates* in order to discover if more WADs could be put to low power modes.

Algorithm 2 *Correction*

1: **while** MORFEO is active **do**
2: Sleep(Time–slot)
3: **if** Wake up–slot **and** partial–coverage profile **then**
4: **for all** WADs **do**
5: WAD $\rightarrow$ Active Mode
6: **end for**
7: **else**
8: **if** Network performance OK **and** traffic not increasing **then**
9: $\rightarrow$ Reactive updates
10: **else**
11: **if** $\exists$ WAD $\in$ Off Mode **then**
12: WAD $\rightarrow$ Active Mode
13: **else if** $\exists$ WAD $\in$ Partial/Sector Sleep Mode **then**
14: WAD $\rightarrow$ Active Mode
15: **else if** $\exists$ WAD $\in$ Full–Power Mode **then**
16: WAD $\rightarrow$ Active Mode
17: **end if**
18: **end if**
19: **end if**
20: **end while**

Finally, we consider that it might be possible to allow partial coverage rather than full coverage, e.g., in an enterprise deployment, APs might be turned off to provide reduced coverage at night when very low traffic load is expected. In this case, in order to give a chance (i) to the incoming users for connectivity and (ii) to the network to reconfigure the operation modes of the WADs, a wake up slot is introduced. During this time-slot, all the WADs are switched to Active Mode to check whether the operating conditions have changed. The wake-up slot is also executed during *Correction* (line 1-5 in Algorithm 2).

After the first run, *MORFEO* stays in *Correction* and analyses network performance and current conditions based on the statistics collected from CoMa. If the analysis indicates further energy savings, *Reactive updates* evaluate changes in operation modes and goes back to *Correction*, and *MORFEO* continues in this cycle.

IV. Effect of Operation Mode Switching Times

In this section, we present the impact of switching operation modes on the energy efficiency of *MORFEO* decisions. It can be experimentally measured, indeed, that switching between different modes bears a cost. As switching times can be very large (i.e., waking up a WiFi access point from idle mode can easily take 100 s), the cost is non–negligible. This implies that it is not convenient to put a device in low power modes for short time intervals.

Let us first consider the situation in which only two modes (Active and Off Mode) are supported. It is actually possible to define a minimum off–time for which the energy gains (in terms of energy saved with respect to keeping the device in Active Mode) compensates the energy waste (related to the energy spent in switching from Active Mode to Off Mode, from Off Mode back again to Active Mode and Off Mode).

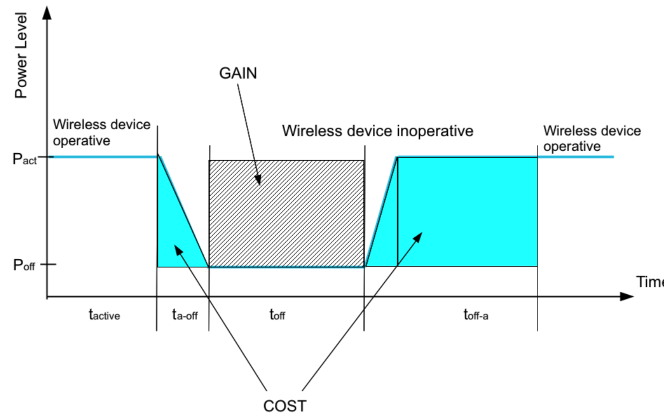

Figure 2: Graphical representation of the energy saved and of that wasted due to switching time. During switching, the device spends additional energy $E_{act-off}+E_{off-act}$, which we refer to as cost.

The situation is graphically depicted in Fig. 2. The *energy gain* equals $(P_{act}-P_{off})\,t_{off}$, where P_{act} and P_{off} represent the power consumption in Active and Off Modes, respectively, and t_{off} is the duration of the time spent in Off Mode. The *switching energy cost* is $E_{act-off}+E_{off-act}$, where $E_{act-off}$ and $E_{off-act}$ represent the *energy* spent in switching from Active to Off, and Off to Active Modes, respectively. Here, compared to the case where the device is always in Active Mode, the *normalised energy cost* is: $E_{act-off}+E_{off-act}-P_{act}\cdot(t_{a-off}+t_{off-a})$, where t_{a-off} and t_{off-a} are the times to switch from Active to Off Mode and from Off to Active Mode, respectively. The minimum viable Off time, t^*_{off}, can be readily computed by equalling the *energy gain* and *normalised energy cost*, leading to:

$$t^*_{off}=\frac{E_{act-off}+E_{off-act}-P_{act}\cdot(t_{a-off}+t_{off-a})}{P_{act}-P_{off}}, \tag{1}$$

In Fig. 3, the measurement results from a sample WAD_1[1] are reported. The WAD_1 takes (i) $\approx$ 6s in order to turn off completely and (ii) $\approx$ 113s in order to be on and completely operational again. Since for this WAD_1, the Active Mode and switching mode powers are almost the same, based on the Equation 1, $t^*_{off}\approx 0$. Once the WAD is put in Off Mode, *MORFEO* does not evaluate the conditions that can lead to its wake-up before $t^*_{off}+t_{a-off}+t_{off-a}$, which is approximately 120 seconds, for WAD_1.

The Fig. 4 reports an additional example from WAD_2[2]. The WAD_2 takes (i) $\approx$ 5s in order to turn off completely and (ii) $\approx$ 60s in order to be on and completely operational again. Since for this indoor WAD_2, the Active Mode and switching

[1]The WAD used is a commercially available $Saxnet\ Meshnodes\ III$, which are specialized for outdoor usage under extreme environmental conditions. It is part of a testbed [21] deployed at Deutsche Telekom Laboratories in Berlin, Germany, see additional details in Section V.

[2]The WAD used is a commercially available $Saxnet$ for indoor usage

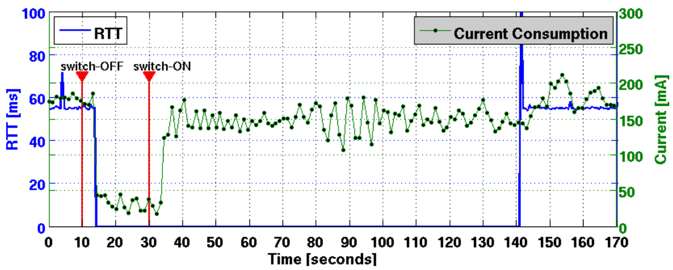

Figure 3: Switching time: measurements for a sample WAD.

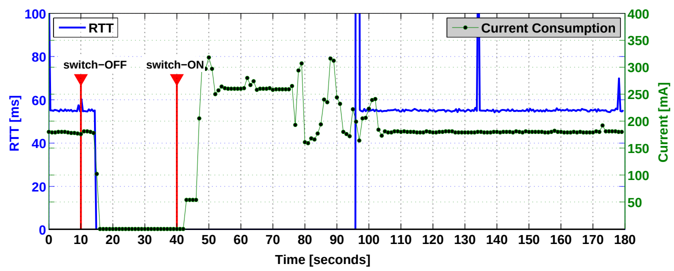

Figure 4: Switching time: measurements for a sample WAD_2.

mode powers are different, based on the Equation 1, t^*_{off} is $\approx$ 15 seconds. Therefore, once the WAD_2 is put in Off Mode, *MORFEO* does not evaluate the conditions that can lead to its wake-up before approximately $\approx$ 80 seconds, for this indoor WAD_2. As a conclusion, we can note that the energy cost to put WAD_2 on is higher than the energy cost of the device in Active mode.

A similar reasoning applies to the case in which also a Partial Sleep Mode is supported. Also in this case it is possible to define a minimum time the device has to spend in Partial Sleep Mode to amortise the cost of switching. We experiment with different configurations for Partial Sleep Mode with a WAD with two wireless interfaces. First, we configure the WAD in full Partial Sleep Mode by switching off all the wireless interfaces and measure the power consumed as $\approx$ 10.38 W. Then, we configure the WAD such that one wireless interface is up and the other is down, and measure the power consumed as $\approx$ 11 W . Finally, we configure the WAD in Active Mode by switching on all the wireless interfaces of the WAD. The power consumed for Active Mode is 11.63 W. Therefore, the energy cost of each wireless interface is constant and approximately equal to 0.62 W. We repeated the experiment in the reverse sequence (i.e., from two interfaces up to only one interface up and both interfaces down) and observed similar results. The switching times were less than 10 ms[3]. Based on these experiments, we concluded that the power cost for switching between Active and Partial Sleep Modes were negligible. This trivially requires that in our system, in order to have any energy gain, the energy spent in Active Mode should be greater than the energy spent in Partial Sleep Mode. As this always holds true, the minimum time required for saving energy to switch a WAD to Partial Sleep Mode is $t^*_{psm} = 0$, i.e., there is no condition on the minimum time to spend in PSM.

[3] 10 ms is the time granularity of our measurement framework.

The minimal amounts of time to be spent in Off and Partial Sleep Modes defines a limit on the dynamics of *MORFEO*, i.e., on its ability to track variation in context and traffic. Yet, the values we experimentally measured for WiFi-based WADs are much shorter than what can be found in cellular network technology where, e.g. switching Partial Sector Mode–Active Mode or Off Mode–Active Mode may be necessary every tens of minutes or even hours.

V. Performance Evaluation

We implemented *MORFEO* in python and installed the energy monitoring and control framework in a 8-node testbed deployed at Deutsche Telekom Laboratories in Berlin, Germany [21] . In this section, we present the experimental evaluation of *MORFEO* under several practical scenarios. To this end, we first describe our experiment methodology, and conclude with energy measurement results and network traffic performance. The goal of our evaluation is to show that *MORFEO* can save energy under different conditions without degrading service quality.

A. Experiment Methodology

Network setup. The testbed is composed of 8 custom IEEE 802.11a/b/g Access Point (AP) and 13 clients. Fig. 6 shows the network topology. The testbed covers an area of approx. $9600m^2$, which is divided into four separate courtyards. The APs are commercially available Saxnet meshnode III equipped with multiple wireless interfaces. They are also connected via Ethernet to a wired backbone network. The WiFi antennas are sector type and the operating frequency was set to 5.18GHz for wireless backbone links and 2.24 GHz for wireless access links. The rate adaptation algorithm has been set to the default auto, and the transmission power has also been left as the default value of 17 dBm (50.12 mW) for all experiments.

Traffic Generation and Power Consumption Monitoring. Traffic is generated using the Iperf traffic generator and is injected into the network from APs to clients. The power consumption is measured using the EPC [20] power meter. EPC is a commercial solution for real–time monitoring of power consumption. The power consumption statistics are logged by EPC with a granularity of 0.1 W and a sampling period of 1 second.

Testing Methodology. To test the network under different traffic conditions, we generate synthetic traffic in the form of single UDP or TCP flows with a duration of 300 seconds. During this time, we collect power consumption measurements and network performance statistics from each AP. We considered the following network profiles:

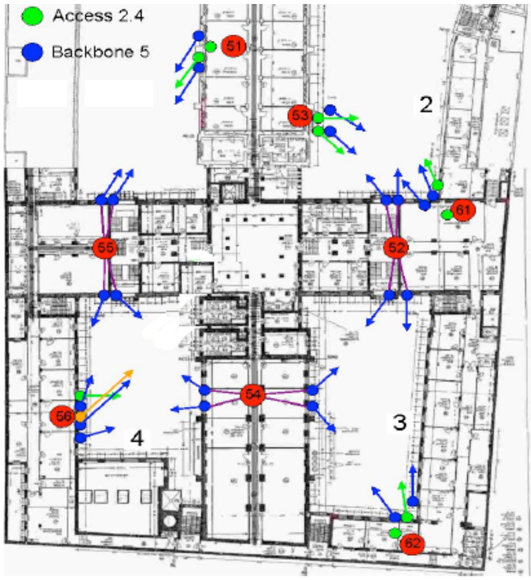

Figure 5: Network topology for testing scenarios. The red dots represent the WiFi APs used in our experiments. The green and blue dots and arrows show the antenna directions for access and backbone interfaces, respectively.

i) Normal Profile: The network is working in its default setting with all nodes in Active mode, and *MORFEO* is not activated.
ii) Isolated Profile: Partial coverage is allowed and the APs work in an isolated manner without any overlapping coverage with neighboring APs.
iii) Partial–coverage Profile: Partial coverage is allowed and the network coverage is achieved with 2 head candidates.
iv) Full–coverage Profile: Full coverage is required and the network coverage is again achieved with 2 head candidates. It is important to note that in the Full–coverage profiles, the head candidates are able to entirely cover the coverage area of the sleep candidates and hence, operate with higher transmission powers.

In our experiments, we activated *MORFEO* during 15 slots. We calculated the minimum time–slot, based on how long CoMa takes collecting measurements and how fast the operation modes can be switched by EC. We measured this time $\approx$ 40 seconds. We also set the slot 4 as the wake–up slot for the partial–coverage scenarios.

For the *Initialisation* part, for the Partial–coverage and Full–coverage profiles, we set 3 APs as a sleep candidate (52,61,62), 2 APs as head candidates (54,55), 2 APs with a special candidates (51,56). . For the Isolated profiles, 6 APs were set as sleep candidates and 2 APs as special candidates (51,56).The situation of node 53 is special, as it is isolated from the other APs, and hence, even if it is a sleep candidate, it needs to stay in Active Mode if a user associates as it is not able to handover its users to a head node. Each experiment duration is set for 700 seconds and we run each experiment 3 times. The following experiments were performed for each profile:

i) Scenario 1 - No users: The network is working without users.
ii) Scenario 2 - No traffic: There are users associated to the network, but they do not generate traffic.
iii) Scenario 3 - UDP traffic: There are users in the network generating UDP traffic. Each user generates a flow of 2Mbps.
iv) Scenario 4 - TCP traffic: There are users in the network generating TCP traffic.

B. Results

Based on the measurement results, we first look at the average power consumption of each AP under the four different scenarios and profiles. The Fig. 6 show that the average power consumed in the normal profile is in most cases higher than the power consumed by the APs with *MORFEO* in the other profiles. The results show that while sleep candidates 52,61 and 62 can significantly save energy, the head candidates 54 and 55 are able to save energy in Scenario 1, with no users. Table I shows the number of interfaces and the user distribution per AP (51 and 56 are omitted as they do not have any users). The results show that, if partial coverage is allowed, the sleep candidate nodes are able to save significant energy.

In Fig. 7, normalized values of the average network power consumption for the different testing scenarios are presented. We calculated the normalized network power consumption by considering the power consumption in the normal profile as the reference index for the scenarios. Both Fig. 6 and Fig. 7 show that the energy savings with *MORFEO* vary from 3-45% of the total of power consumed in the normal profile, even if the only a number of nodes could be sleep candidates . As expected, the best opportunity for energy savings occurs when the network has no associated users. In this case, *MORFEO* can save between 15-30%. When there are users associated to the network, the Isolated profile experiences the worst case in terms of energy saving due to the fact that the network topology does not allow the presence of head candidates. Hence, *MORFEO* can only opt for Partial or Sector Sleep Mode for the APs. Instead, in the Partial–coverage and Full–coverage profiles, the presence of the 25% of head candidates allows the network to save energy while maintaining the coverage and capacity.

The results also show that the Partial–coverage profile saves the highest energy since this profile allows some coverage holes in the network for short periods. This is because, in the Full–coverage profile, the head candidates are sharing completely the coverage area of the sleep candidates, and hence the wake-up slot is not activated. The APs are put in Off Mode only if they are sharing the coverage area with head candidates. Clearly, this condition limits the

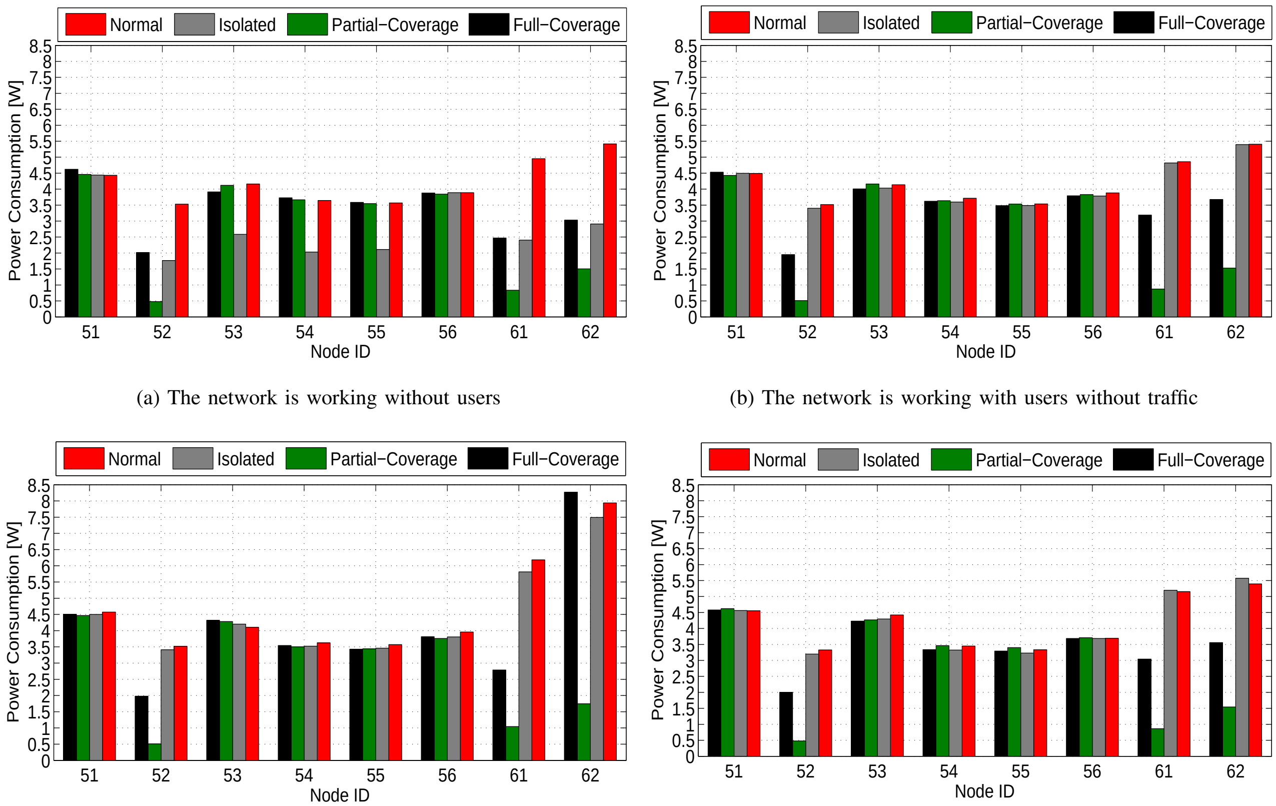

(a) The network is working without users

(b) The network is working with users without traffic

(c) The network is working with users and UDP traffic

(d) The network is working with users and TCP traffic

Figure 6: Average power consumption of each AP in the network working under different scenarios.

amount of energy that can be saved compared to Partial–coverage, where the APs are put in Off Mode even if they are not sharing the same coverage area with the head candidates. This limitation holds true for all the experiments even when the network is operating without users (see Fig. 7). However, note that, in the Partial–coverage profile, the wake–up slot is activated every 4^{th} slot in order for the incoming users to associate to the network. Therefore, in this profile, there is a cycle between periods of full and partial–coverage. Obviously, lower number of wake-up slots enables more energy savings, but with the consequence of fewer full–coverage periods. Therefore, wake–up slot frequency must be carefully chosen and should represent the network dynamics in terms of both users and traffic. Finally, we observe that the most effective way to save energy is by putting the AP in Off Mode as much as possible. Therefore, an accurate network coverage map will allow more effective planning in terms of head and sleep candidates, which will eventually lead to higher energy savings.

Finally, the Table I reports the iperf output for each AP under different profiles. In the case of UDP traffic, the packet loss was less than 1% in all the experiments. In the case of TCP traffic, we observed the same performance for both scenarios. Therefore, the network performance does not suffer any degradation when *MORFEO* is operating. It is important to note the UDP and TCP results are presented in a single column since the obtained results are the same across all the profiles.

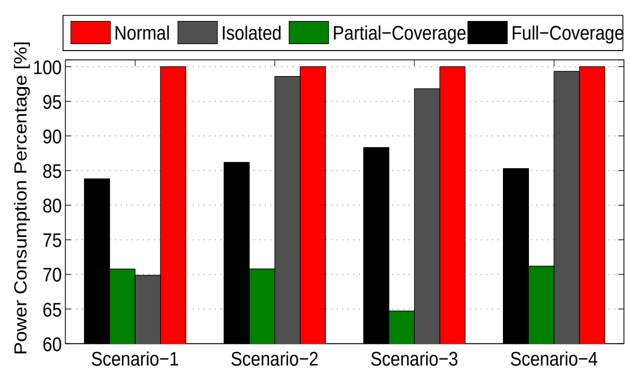

Figure 7: Normalized Network Power Consumption for the different testing scenarios.

VI. Conclusions

In this paper, we presented a practical energy-saving decision algorithm for improving the energy efficiency of

Table I: Network performance for each AP working under $Scenario_3$ and $Scenario_4$

AP ID	Int.	Users	UDP [Mb/s]	Packet Loss [%]	TCP [MBytes]
52	3	2	4	0.01	246
53	1	1	2	0.02	322
54	4	4	8	0.00	1324
55	3	3	6	0.00	756
61	2	1	2	0.00	322
62	2	2	4	0.01	419

wireless access networks. The *MORFEO* algorithm, together with the energy monitoring and control framework, allows appropriate energy saving decisions to adapt the energy consumption of a wireless infrastructure to the actual network conditions in terms of both user distribution and traffic patterns. The experimental evaluation of *MORFEO* in a real network deployment confirmed that the proposed solution can provide significant energy savings between with minimal degradation of network performance. We have designed *MORFEO* to take advantage of context and measurement-based network information, and adapt network operation in any wireless network. Therefore, as future work, we plan to extend our evaluation to other access network infrastructures, and in particular heterogeneous network deployments. We believe *MORFEO* can enable even higher energy savings taking into account the energy-performance-coverage trade-offs of different technologies.

REFERENCES

[1] K. Gomez, D. Boru, R. Riggio, T. Rasheed, D. Miorandi, and F. Granelli, "Measurement–based modelling of power consumption at wireless access network gateways," *Computer Networks*, vol. 56, pp. 2506–2521, Jul. 2012.

[2] K. Gomez, R. Riggio, T. Rasheed, F. Granelli, M. D., and C. I., "Analysing the energy consumption behaviour of WiFi networks," in *Proc. of IEEE GreenCom*, Sep. 2011.

[3] R. Riggio and D. J. Leith, "A measurement-based model of energy consumption in femtocells," in *IEEE/IFIP Wireless Days*, Dublin, Ireland, 2012.

[4] A. Jardosh et al., "Green WLANs: On-demand WLAN infrastructures," *Mobile Networks and Applications*, vol. 14, no. 6, pp. 798–814, Dec. 2009.

[5] M. A. Marsan and M. Meo, "Energy efficient management of two cellular access networks," *SIGMETRICS Perform. Eval. Rev.*, vol. 37, pp. 69–73, Mar. 2010.

[6] L. Saker, S. Elayoubi, and H. Scheck, "System selection and sleep mode for energy saving in cooperative 2G/3G networks," *Vehicular Technology Conference Fall,*, pp. 1–5, Sep. 2009.

[7] M. Marsan, L. Chiaraviglio, D. Ciullo, and M. Meo, "Optimal energy savings in cellular access networks," in *Proc. of IEEE ICC*, Dresden, Germany, 2009.

[8] G. Micallef, P. Mogensen, and H. O. Scheck, "Cell size breathing and possibilities to introduce cell sleep mode," *European Wireless Conference*, pp. 111–115, April 2010.

[9] S.-E. Elayoubi, L. Saker, and T. Chahed, "Optimal control for base station sleep mode in energy efficient radio access networks," in *INFOCOM, 2011 Proceedings IEEE*, April 2011, pp. 106–110.

[10] K. Dufkova and, M. Bjelica, B. Moon, L. Kencl, and J.-Y. Le Boudec, "Energy savings for cellular network with evaluation of impact on data traffic performance," in *Proc. of IEEE EW*, Lucca, Italy, 2010.

[11] A. P. Jardosh, K. Papagiannaki, E. M. Belding, K. C. Almeroth, G. Iannaccone, and B. Vinnakota, "Green WLANs: On-demand WLAN infrastructures," *Mob. Netw. Appl.*, vol. 14, pp. 798–814, Dec. 2009.

[12] J. Lorincz, A. Capone, and M. Bogarelli, "Energy savings in wireless access networks through optimized network management," in *Proc. of the 5th IEEE international conference on Wireless pervasive computing*. Mondena, Italy: IEEE Press, 2010, pp. 449–454.

[13] Y. Al-Hazmi, H. De Meer, K. A. Hummel, H. Meyer, M. Meo, and D. Remondo, "Energy-efficient wireless mesh infrastructures," *IEEE Network Magazine*, vol. 25, pp. 32–36, Mar.-April 2011.

[14] S. Armour, T. OFarrell, S. Fletcher, A. Jeffries, D. Lister, S. McLaughlin, J. Thompson, and P. Grant, "Green radio: Sustainable wireless networks," *Mobile VCE Conference*, 2009.

[15] A. J. Fehske, F. Richter, and G. P. Fettweis, "Energy efficiency improvements through micro sites in cellular mobile radio networks," in *Proc. of the IEEE Global Communications Conference*, 2009, pp. 1–5.

[16] F. Richter, A. J. Fehske, P. Marsch, and G. P. Fettweis, "Traffic demand and energy efficiency in heterogeneous cellular mobile radio networks," *Vehicular Technology Conference*, pp. 1–6, May 2010.

[17] K. J.-M., C. M., and F. Marache, "Limiting power transmission of green cellular networks: Impact on coverage and capacity," *Communications (ICC) IEEE International Conference*, pp. 1–6, May 2010.

[18] R. Riggio, K. M. Gomez, T. Rasheed, and C. Sengul, "Energino: Energy saving tips for your wireless network," *ACM Sigcomm 2012*, August 2012.

[19] S. Göndör, A. Uzun, and A. Küpper, "Towards a dynamic adaption of capacity in mobile telephony networks using context information," in *Proceedings of the 11th International Conference on ITS Telecommunications*. St. Petersburg, Russia: IEEE, Aug 2011, pp. 606–612.

[20] "Power Control," http://www.gude.info/index.php.

[21] A. Uzun, S. Göndör, N. Bayer, L. Kollecker, and A. Küpper, "Visualizing the effects of power management algorithms for mobile networks under realistic conditions," in *Proceedings of the Conference on Electronics Goes Green 2012*. Berlin, Germany: IEEE, Sep. 2012.