WLAN-Opp: Ad-hoc-less opportunistic networking on smartphones

Sacha Trifunovic [a,*], Maciej Kurant [b], Karin Anna Hummel [a], Franck Legendre [c]

[a] *Computer Engineering and Networks Laboratory, ETH Zurich, Switzerland*
[b] *Google, Zurich, Switzerland*
[c] *Uepaa!, Zurich, Switzerland*

ARTICLE INFO

Article history:
Received 5 March 2014
Received in revised form 21 July 2014
Accepted 27 July 2014
Available online 7 August 2014

Keywords:
Opportunistic networks
Wireless communication
Smartphones

ABSTRACT

Opportunistic networking enables many appealing applications including local social-networking, communication in emergency situations, and circumventing censorship. The increasing penetration of smartphones should, in theory, foster opportunistic networking. In practice, current candidate technologies for opportunistic networking, such as Wi-Fi ad-hoc, Bluetooth, and Wi-Fi Direct, are either not available on current smartphones, or require undesired user interaction to establish connectivity.

To overcome these shortcomings, we propose WLAN-Opp for smartphones. This IEEE 802.11-based technology leverages the tethering mode of smartphones, a feature originally used to share Internet access, which allows smartphones to become WLAN-based access points that provide networks for other smartphones operating as stations. The transitions between WLAN-Opp access point and station mode are randomized as a function of the number of other co-located networks and stations, and depend on duty cycling intervals. We optimize the probabilistic operations in a simulation study and provide a parametrized implementation of WLAN-Opp for out of the box smartphones. By replaying real contact traces in simulation, we find that WLAN-Opp can utilize up to 80% of the contact time while saving up to 90% of the energy Wi-Fi ad-hoc would consume. Finally, we demonstrate in a field trial with 34 users over 5 days that WLAN-Opp can provide a practical solution in a realistic setting.

1. Introduction

Ubiquitous network connectivity is often taken for granted in developed countries. Yet, natural disasters such as earthquakes or floods, which occur more often then commonly believed [5], hamper the performance of communication networks and, in worst case, may destroy network infrastructures. Another example of missing connectivity can be found in developing countries where high-bandwidth connectivity is not provided area-wide. Even if networks are available, authoritarian governments often censor communication by blocking access to information and online social networks (e.g., YouTube, Facebook, Twitter) [9] and can even cause an Internet and mobile phone outage [4].

Opportunistic networks provide an appealing technology to maintain delay-tolerant connectivity under such harsh conditions as well as offload existing infrastructure [8,17], or even support freedom of speech. In opportunistic networks [2,20], mobile wireless devices cooperate to distribute information over spontaneous wireless links whenever mobile devices encounter one another. The increasing penetration of smartphones and tablets favor

* Corresponding author.
E-mail address: trifunovic@tik.ee.ethz.ch (S. Trifunovic).

http://dx.doi.org/10.1016/j.adhoc.2014.07.011
1570-8705/

the feasibility of such networks, however, there is a lack of enabling networking technologies for ad-hoc establishment of wireless links in today's smart mobile devices. Most of all, common smartphones do not support IEEE 802.11 WLAN ad-hoc [12], unless rooted or jail-broken. Communicating via Bluetooth is another option, but this option is limited in terms of communication range and bandwidth as well as human interaction-free discoverability. Although recent developments achieve wider ranges and better transfer rates through hybrid solutions, still WLAN-based networking provides at least 3–10 times wider ranges and 20 times higher rates. Wi-Fi Direct [27], the Wi-Fi Alliance's answer to Bluetooth suffers from cumbersome manual discovery and pairing procedures which in addition are very energy intensive [26].

With WLAN-Opp, we propose an approach to establish connectivity among modern smartphones based on basic WLAN infrastructure functionality, namely access points (APs) and its associated stations (STAs). WLAN-Opp uses the tethering mode of smartphones that allows the device to become an AP. This way, a communication infrastructure is established for co-located devices (devices in contact), which in turn can connect to the provided WLAN network as stations without requiring time-consuming pairing of devices or cumbersome user input on encounters. Further, the functions necessary for WLAN-Opp are provided by the API of common smartphones. These characteristics make WLAN-Opp a very practical opportunistic networking approach.

Conceptually, WLAN-Opp allows to adapt to various topologies of co-located devices by controlling periods of scanning and connectivity, and changing between AP and station mode. The flexibility in changing between AP and station mode allows to share the load of providing AP functionality and, thus to control resource consumption, in particular battery power. Yet, providing the necessary connectivity in a dynamic environment where mobile devices appear and disappear frequently is not straight forward. To solve this problem, WLAN-Opp introduces parametrized randomization of state changes as a core method.

We presented the concept of WLAN-Opp first in [25], where we provided a proof of concept based on simulation. In this paper, we extend our work by a detailed description of a more realistic and flexible WLAN-Opp state machine and a detailed investigation of its major parameters and their impact on *contact utilization*, i.e., the fraction of the time co-located devices can actually communicate via WLAN-Opp. Further we validate the simulated state machine by an implementation of WLAN-Opp on real devices. Overall, we make the following contributions:

- We introduce the details of WLAN-Opp including the state model, randomization, and parameters controlling the behavior of WLAN-Opp (Section 2). We investigate these parameters for varying amounts of co-located devices in a stationary setting using event-driven simulation. As a result of this study, we provide an appropriate configuration of WLAN-Opp for good contact utilization (Section 3).
- By exposing WLAN-Opp to a set of real mobility traces we demonstrate that the occurring contacts are well utilized by WLAN-Opp, i.e., up to 80%, while saving up tp 90% of the energy Wi-Fi ad-hoc would consume (Section 4).
- We implement WLAN-Opp on Android, the currently dominating mobile OS, and make it available as open-source. Using this implementation, we validate our approach by comparing simulation and real measurements and thus show that our simulation model accurately matches the behavior of real smartphones (Section 5).
- Finally, we present a field study lasting over five days with 34 participants and demonstrate the practicality of WLAN-Opp on real devices in an every-day setting. This experiment also shows the usefulness of an additional feature of WLAN-Opp, i.e., leveraging available open access points in the environment to save energy (Section 6).

2. WLAN-Opp

WLAN-Opp uses the WLAN access point feature of mobile phones to enable delay tolerant connectivity. Some devices change into access point mode (AP) and provide the wireless network, while other devices are in idle mode (IDLE) and scan for networks or are connected to a network as a station (STA). In STA mode, devices are still able to scan and discover additional networks. Fig. 1a shows the simplest WLAN-Opp network between one AP and one STA node.

2.1. Challenges

Establishing opportunistic connectivity between mobile devices in proximity is the primary goal of WLAN-Opp. Therefore, being aware of other nodes and coordinated AP mode provisioning are key factors to be considered by the opportunistic networking scheme, which faces the following challenges[1]:

- IDLE nodes can scan for networks to discover APs, but are unaware each other (Fig. 1b). A fraction of nodes may thus be in proximity but not (yet) connected.
- APs are unable to scan for networks and are thus unable to detect each other (Fig. 1c). APs are only aware of their STAs, i.e., nodes that are associated with them. Hence, they might miss opportunities to join other existing networks in proximity leading to partitioned networks (disjoint groups, cf. Fig. 1d).
- Finally, STAs associated with different APs are not aware of each other. This is the second cause for disjoint groups that are unable to communicate with each other (Fig. 1d).

[1] These challenges are mainly given by the smartphone's OS and solvable with access to the Wi-Fi driver. However, this would require rooting the device or installing a custom Android OS.

(a) Normal WLAN-Opp operation with one AP and one STA.

(b) Two IDLE nodes cannot discover each other.

(c) Two APs cannot discover each other.

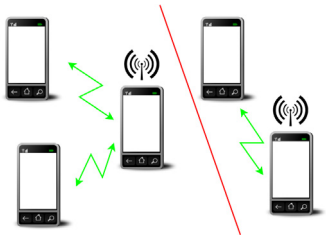

(d) Disjoint WLAN-Opp groups.

Fig. 1. WLAN-Opp operation and challenges.

To prevent the opportunistic network from being caught in these situations, there is need to randomize the time the nodes remain in the respective modes.

2.2. Randomization

To overcome the limits of device discovery and disjoint groups, we randomize the decision of a node to keep or change its mode, and consequently, limit the time spent in each mode. This ensures that two or more devices will most likely not remain in the same mode (all IDLE or all AP) for indefinite time. Hence, mutual discovery is facilitated (solving the problems shown in Fig. 1b and c).

The randomization-based mechanism also introduces *topological dynamics* and reorganizes disjoint groups. In static scenarios, e.g., in an office or at home, the forced topology changes increase the likelihood that co-located devices that are connected to different APs will eventually be able to communicate (solving the problem shown in Fig. 1d). In mobile scenarios, topology changes come for free as devices move in and out of range of each other. Still, randomization has an advantage concerning simplicity and robustness as it provides a simply way to adapt to topology changes without requiring communication. For example, if the AP that currently provides connectivity unexpectedly leaves the group, another random node will take over. For this reason we base the WLAN-Opp state machine on randomized decisions.

2.3. WLAN-Opp state machine

The main part of WLAN-Opp is its state machine with probabilistic transitions, which implements the concept of randomization. The transitions depend on the current situation, i.e., the availability of networks and number of neighbors in the environment, as well as the time having already spent in AP mode. A node's possible states (IDLE, STA, and AP) and the state transitions are depicted in Fig. 2. Table 1 summarizes the variables influencing the state transition probabilities. To ease analysis, implementation, and execution, WLAN-Opp is designed to run in *slotted mode*, with slot time t_{slot}. In every slot, the state machine is executed. This way, we may adjust the behavior of WLAN-Opp to new settings by simply modifying the state transition probabilities. Alternatively, dynamically changing time periods per state may be used, yet, this adds complexity to implementation and analysis. Note that the slots are not synchronized among the nodes. We now detail the operation in the each of the WLAN-Opp states.

IDLE mode: The IDLE state is the initial state of a WLAN-Opp node and the state a node returns to when loosing connection in STA mode or after ceding the AP role. In the IDLE state, the node scans for available networks. This state is left when the node switches to either STA or AP mode.

STA mode: IDLE nodes may eventually find an available network and will connect to it. If multiple networks are discovered simultaneously, one is picked uniformly at random. In STA mode, a node still scans for networks and switches to another available network with switching probability p_{switch}^{STA}, which depends on the current number of neighbors N_c:

$$p_{switch}^{STA} = p_{switch}^{STA}(N_c). \quad (1)$$

In principle, a device aims at staying in a network that provides connectivity to many of its co-located devices. Hence, the probability to switch to another, yet unknown network should decrease with an increasing number of current neighbors N_c.

AP mode: If an IDLE node finds no networks to connect to, it becomes an AP and creates a network by itself with probability p_{on}^{AP}. As nodes in AP mode consume more resources than in the other modes, the probability should depend on the time a node has not been an AP denoted as t_{off}^{AP}. Additionally, as recent neighbors to whom the node has just been connected to (in case the network just disappeared) compete in becoming the next AP, the probability to become an AP should also depend on the number of

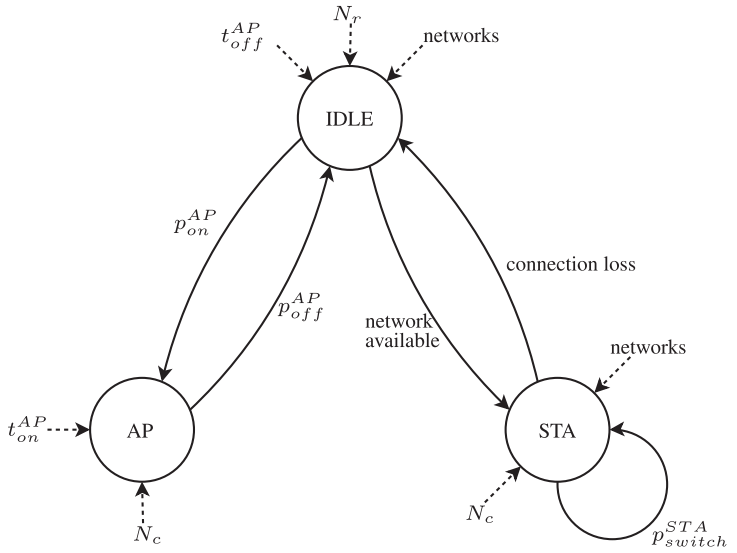

Fig. 2. State transition diagram of WLAN-Opp.

Table 1
Variables that influence the state machine.

Input variables	Description
Networks	Available WLAN networks (# APs in proximity)
N_c	Current neighbors
N_r	Recent neighbors
t_{off}^{AP}	Time since last AP mode
t_{on}^{AP}	Time in AP mode

recent neighbors N_r (intuitively, it should decrease with N_r):

$$p_{on}^{AP} = p_{on}^{AP}\left(t_{off}^{AP}, N_r\right). \quad (2)$$

For the same resource consumption reasons, AP mode should be turned off with probability p_{off}^{AP} that depends on the time already spent in AP mode denoted by t_{on}^{AP}. Another factor should be the number of current neighbors N_c, i.e., the number of stations that are currently served and would loose connectivity:

$$p_{off}^{AP} = p_{off}^{AP}(t_{on}^{AP}, N_c). \quad (3)$$

The probability p_{off}^{AP} should decrease with the number of neighbors N_c and increase with t_{on}^{AP}, the time already spent in AP mode. Note that an actual low battery power level may additionally be considered to limit the time in AP mode. In this work, we only use such limits in the field trial (cf. Section 6).

The choice of the network switching probability p_{switch}^{STA} and the AP turn off probability p_{off}^{AP} highly impact the utilization of a contact, i.e., the fraction of the time co-located nodes may communicate. These parameters are thus studied in detail in Section 3. Differently, the probability to become an AP, p_{on}^{AP}, is defined straightforward to assure that only one co-located station turns into an AP on expectation by design, as introduced in the following.

2.4. Design decisions

The WLAN-Opp state machine has been implemented both in simulation and on real devices to evaluate WLAN-Opp and to demonstrate its practicality. To provide a realistic design, we have to consider a few *real world aspects*. First of all, some operations do not happen instantaneously but take some time on real devices, like scanning for networks and connecting to one of them. Second, devices have a limited energy source, which requires certain design decisions concerning duty cycling and fairness. Last, devices may disappear, either because they crash, run out of battery, or because they move away. This requires a robust design of the AP selection protocol.

Temporal operation granularity: On real mobile devices, operations take a certain amount of time. For example, Android is set to scan for networks every five seconds and it takes around 4.5 s to turn-on AP mode on a Google Nexus One device. To make sure all operations can complete during one slot, the slot time should be at least five to ten seconds. On the other hand, the slot time should not be too long, otherwise changes in the network environment are missed. For example, the device may not be aware of the loss of neighbors or the availability of new networks, and actions are delayed. Additionally, the slot time should not be equal on all nodes to avoid synchronized state changes that prevent nodes from reaching a stable setting. Also, the time slots do not need to be of equal length in all states. Generally, to achieve a good trade-off, we select the slot duration uniformly at random in the range from 10 to 15 s for the real implementation and for the simulation: $t_{slot} = \mathcal{U}[10, 15]$. However, in the IDLE state, we achieve a slightly better performance if we scan for new networks and initiate association attempts to available networks twice in every time slot, i.e., every 5–7.5 s.

Fairness and duty cycling: The AP mode is the most energy consuming mode in WLAN-Opp since a node has

to send SSID beacons about every 100 ms and to relay traffic. For the sake of fairness we limit the time a node is allowed to stay in AP mode (t_{on}^{AP}) to $t_{on,max}^{AP}$ and we duty cycle the AP operation by introducing a minimum time $t_{off,min}^{AP}$ a node has to wait before becoming an AP again; we denote the time a node is off-duty by t_{off}^{AP}. To achieve higher network stability, a longer $t_{on,max}^{AP}$ is preferred, but it comes at the cost of high energy consumption of one single device. However, nodes with an external power supply may choose to provide AP functionality indefinitely, e.g., to aid neighbor connectivity in popular hotspots. In case battery power is used, a smaller $t_{on,max}^{AP}$ will be chosen, which improves energy consumption fairness among a set of devices. While the WLAN-Opp platform allows for such a per device configuration of the maximum AP time, the performance analysis of the WLAN-Opp state machine was performed with a constant value $t_{on,max}^{AP} = 10$ min. We thoroughly analyze how to adaptively choose the maximum AP time to optimize fairness and efficiency depending on the anticipated remaining contact time among devices in [26].

The duty cycle mainly impacts neighbor discovery. While a larger $t_{off,min}^{AP}$ saves power, we miss more connection opportunities. To improve this trade-off, duty cycling should depend on the context. Hence, if neighbors are present, we set the cycles to a short duration, $t_{off,min}^{AP} = 10$ s, that will exponentially increase each time a devices becomes an AP, but no STA associations happen; $t_{off,min}^{AP}$ increases up to a configurable maximum time denoted by $\max\left(t_{off,min}^{AP}\right)$. This way, WLAN-Opp can efficiently adapt to changes in node densities such as given by day-night cycles.

Randomized AP selection: Communication always requires an AP, but multiple co-located APs result in disjoint groups – a situation we aim to avoid by the following approach: If sufficient time has elapsed since the last time the node was in AP mode, i.e., $t_{off}^{AP} > t_{off,min}^{AP}$, a node may consider becoming an AP depending on the estimated availability of other nodes with potential to change to AP mode, which we call *candidates*. In case the node was just connected to another AP, which might have moved away or turned itself off, the estimated number of candidates C is the number of recent neighbors N_r. If no recent neighbors are known, i.e., $N_r = 0$, we set $C = 2$ as this is the minimal number of candidates to justify setting up a network and becoming an AP. More formally, we have

$$C(N_r) = \begin{cases} N_r & \text{if } N_r > 0 \\ 2 & \text{otherwise.} \end{cases} \tag{4}$$

The probability p_{on}^{AP} is set to be inversely proportional to the estimated candidates $C(N_r)$, i.e.,

$$p_{on}^{AP}\left(t_{off}^{AP}, N_r\right) = \begin{cases} \frac{1}{C(N_r)} & \text{if } t_{off}^{AP} > t_{off,min}^{AP} \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

With this definition, we will have exactly one device becoming an AP on expectation. For a randomized protocol, this is the ideal tradeoff between finding the next AP as fast as possible while minimizing the probability of multiple APs and disjoint groups. Further, it is ensured that the AP of the recent network will not immediately become an AP again. Whereas such a simple random AP handover is inherently less efficient than a deterministic solution, this random scheme has two advantages: (i) it does not require any additional communication among the devices and (ii) it is very robust if the current AP or the soon to be AP suddenly becomes unavailable due to mobility. For these reasons, WLAN-Opp implements this simple randomized AP selection scheme, while more sophisticated schemes are studied in future work.

In the following, we will detail and study the probabilities to switch off AP mode and to switch between networks in station mode, with respect to their impact on the utilization of contacts.

3. Parameter study

WLAN-Opp is controlled by its state transition probabilities, which depend on external factors such as available networks or neighbors (cf. Table 1). To study how these factors should impact the probabilities and thus the performance of WLAN-Opp, we model WLAN-Opp in an event-driven simulator. Similar to real devices, each node in the simulation implements its own clock and runs the state machine (cf. Fig. 2) in every time slot. The simulation period consists of 100 h, resulting in 25,000–36,000 events per node. For this parameter study we assume a *static scenario*, where a given number of devices are constantly in range. They are, however, not aware of their total number.

First, we define the performance metrics used for the study. Then, we use the simulator to find the parameter settings for the transition probabilities p_{switch}^{STA} and p_{off}^{AP} that perform well for the given metrics.

3.1. Performance metrics

The ultimate aim of WLAN-Opp is to provide a connection whenever there is a contact. To evaluate how well this goal is met, we introduce the metric *utilization ratio*, which is the fraction of the total pairwise contact time $t_{contact}$ a pair is actually able to communicate:

$$r = \frac{t_{com}}{t_{contact}}, \tag{6}$$

where t_{com} is the connection time of the node pair. The average utilization ratio among all node pairs, is denoted as $\bar{r}$.

In the case of a static scenario, this metric indicates how well WLAN-Opp can break up disjoint small groups to combine them to larger groups (largest possible groups). At the same time, it reflects the disconnection times due to restructuring. In dynamic scenarios involving mobility, the utilization ratio is mostly impacted by duty cycling of the discovery process, i.e., how often nodes become an AP to be discoverable.

3.2. Switching between networks

The station switching probability p_{switch}^{STA} determines the flexibility of a node in STA mode to associate with another

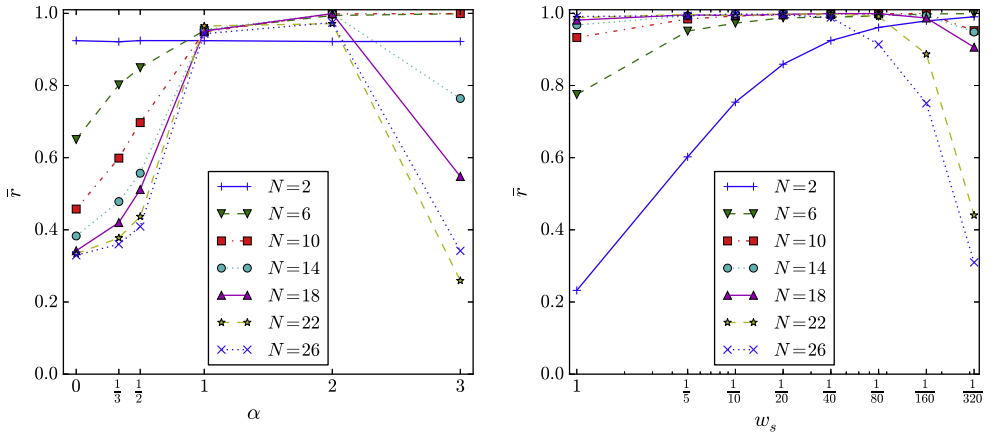

Fig. 3. Station switching probability (p^{STA}_{switch}): impact of parameters α and w_s on average utilization ratio $\bar{r}$ for different numbers of co-located nodes N, number of networks is 3 (standard deviation is negligible and thus omitted): (a) $\bar{r}$ as a function of α (weight $w_s = \frac{1}{40}$) and (b) $\bar{r}$ as a function of w_s (exponent scaling factor $\alpha = 2$).

network, if available. During our study of p^{STA}_{switch}, we assume that multiple networks are available which can be used by all nodes. Every node under investigation always stays in STA mode and must not become an AP. This way, we can analyze p^{STA}_{switch} in isolation.

As introduced in Eq. (1), p^{STA}_{switch} is a function of the number of current neighbors N_c. In case $N_c = 0$, the node should switch networks with probability $p^{STA}_{switch} = 1$ as there is currently no node available to communicate with. Otherwise, the probability should decrease with the number of current neighbors, scaled by an exponential factor α and weighted by w_s:

$$p^{STA}_{switch}(N_c) = \begin{cases} w_s \cdot N_c^{-\alpha} & \text{if } N_c > 0 \\ 1 & \text{otherwise.} \end{cases} \quad (7)$$

To find an appropriate setting of α and w_s, we perform a simulation with varying parameters for different amounts of co-located users N. We explore the range of $N \in \{2, 6, 10, 14, 18, 22, 26\}$, which is feasible for opportunistic networks that may appear in practice. Another input variable is the number of available networks. However, the number of networks does not influence the utilization ratio, given there is more than one network.[2] In the simulation, three available networks are used. We find that the utilization ratio $\bar{r}$ heavily depends on the number of co-located users N. The parameter space exploration yields a good utilization ratio $\bar{r}$ for $\alpha = 2$ and $w_s = \frac{1}{40}$ for all N.

The detailed impact of α is shown in Fig. 3a. For $N = 2$ nodes, the maximum number of current neighbors is 1 and thus not affected by α. For $\alpha < 1$ the utilization declines rapidly with increasing number of nodes. The intuition behind this observation is that a node's desire to stay in a group of neighbors scales slower than the group size. For this reason, α should be at least 1. Using values greater than 1 makes bigger groups more attractive and, at the same time, makes the network more stable by automatically clustering at bigger group sizes. In particular for higher N and larger values of α such as $\alpha = 3$, the flexibility of the network is hampered as small stable groups are formed, which leads to network partitioning and thus reduces the utilization ratio $\bar{r}$.

The influence of the weight w_s is shown in Fig. 3b. Here again, the parameter settings for a good utilization ratio $\bar{r}$ depend strongly on N. Smaller groups prefer more stability, denoted by a small w_s but larger groups need to remain flexible. All in all, for $w_s = \frac{1}{40}$ the simulation yields a generally good result for all group sizes.

3.3. Switching AP functionality

The AP turn off probability p^{AP}_{off} controls the time a node stays in AP mode. On the one hand, this probability determines the stability of the network while on the other hand it ensures fairness among a group of co-located devices in terms of energy consumption. By forcing a change in AP provisioning, WLAN-Opp aims not to drain a single device's battery.

To study the impact of p^{AP}_{off}, we focus on the AP mode operation, i.e., switching on and off AP mode. Whenever nodes are in STA mode, they are configured with the parameters found above (Section 3.2). As introduced in Eq. (3), the AP turn off probability p^{AP}_{off} is a function of t^{AP}_{on}, the time operating in AP mode, and the number of current neighbors N_c. If t^{AP}_{on} exceeds the time allowed to stay in AP mode $\left(t^{AP}_{on,max}\right)$, the node changes its state to IDLE. Otherwise, the probability to change to IDLE depends on the current number of neighbors N_c. Without having neighbors, $p^{AP}_{off} = 1$ in order to turn off AP mode. Thus, the probability to switch off AP mode is given as follows determined by the exponent scaling factor β and weight w_a:

$$p^{AP}_{off}\left(t^{AP}_{on}, N_c\right) = \begin{cases} w_a \cdot N_c^{-\beta} & \text{if } t^{AP}_{on} < t^{AP}_{on,max} \text{ and } N_c > 0 \\ 1 & \text{otherwise.} \end{cases} \quad (8)$$

[2] We verified that with various amounts of networks. The intuition behind this is simple: nodes tend to cluster in very few networks anyway.

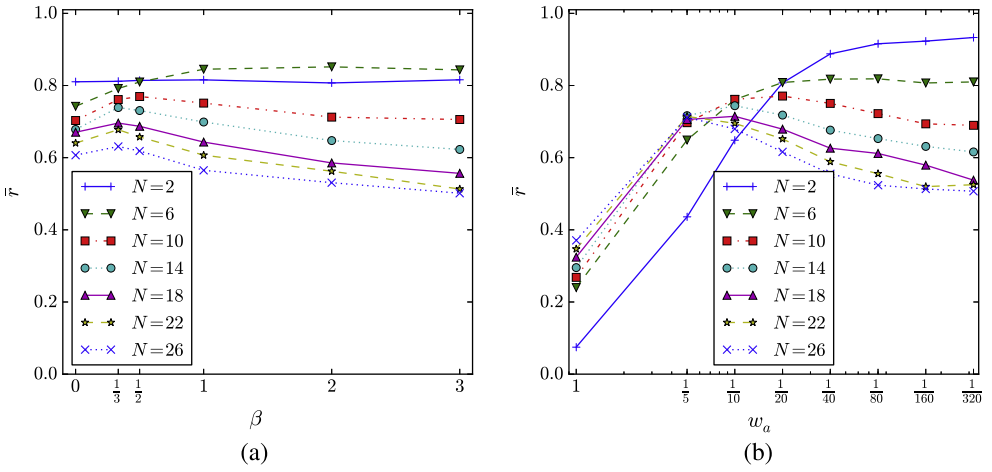

Fig. 4. AP turn off probability $\left(p_{off}^{AP}\right)$: impact of parameters β and w_a on average utilization ratio $\bar{r}$ for different numbers of co-located nodes N (standard deviation is negligible and thus omitted): (a) $\bar{r}$ as a function of β (weight $w_a = \frac{1}{20}$) and (b) $\bar{r}$ as a function of w_a (exponent scaling factor $\beta = 0.5$).

To find a good setting of β and w_a, we simulate the whole parameter space for different amounts of co-located users N. As before, we explore the range of $N \in \{2, 6, 10, 14, 18, 22, 26\}$. Nodes operating in STA mode use the parameter values $\alpha = 2$ and $w_s = \frac{1}{40}$ (cf. Section 3.2). For reasons of fairness, we set the maximal AP on time to $t_{on,max}^{AP} = 600$ s (see Section 2.4). For duty cycling we use a minimum AP off time of 10 s that exponentially increases up to $\max(t_{off,min}^{AP}) = 600$ s. The parameter space exploration resulted that a good utilization ratio $\bar{r}$ is achieved for $\beta = \frac{1}{2}$ and $w_a = \frac{1}{20}$ for all N.[3]

The detailed influence of the exponent scaling factor β is shown in Fig. 4a. Its impact on the values of $\bar{r}$ is actually nearly negligible. The utilization however varies depending on the number of nodes N due to increasing partitions of the network: the number of APs and, thus, emerging networks (network partitions) is higher when more devices are present. All in all, taking into account the whole range of N, a good setting for β is $\frac{1}{2}$.

The results with varying weights w_a are shown in Fig. 4b. We basically observe again the absence of a clear optimum and a strong dependence on the group size due to the resulting network partitions. Taking into account the full range of N, a good setting for w_a is $\frac{1}{20}$.

3.4. Influence of amount of co-located devices

As shown before, the number of co-located users N greatly impacts the utilization ratio $\bar{r}$. We thus observe different outcomes of WLAN-Opp for different N, depending on the parameter settings as visualized in Fig. 5. For very small amounts of co-located nodes, such as $N = 2$, a static configuration is beneficial as reconfiguration attempts are in vain. For very large groups, e.g., $N = 26$, the most flexible configuration performs best, as reconfigurations to join so far disjoint groups are desired. Finally, it can be seen that for our selected parameters resulting from the previous parameter study, WLAN-Opp performs constantly well over a wide range of group sizes. However, there is clearly room for improvement. The most straight forward candidates are deterministic AP handover schemes requiring minimal communication among the nodes.

[3] Varying the maximal AP time has actually a slight impact on the optimal parameters. When significantly increasing $\max\left(t_{off,min}^{AP}\right)$, the optimum slightly moves to more stable parameters (bigger β and smaller w_a). In practice, maximum AP times should be chosen based on battery life and fairness considerations and not their minimal impact on the utilization ratio.

4. Evaluation of WLAN-Opp under mobility

Whenever mobility comes into play, a contact of two or more devices becomes a precious good in opportunistic networks. It is thus crucial for a neighbor discovery and communication establishment protocol to exploit the time in proximity as much as possible. In this section, we evaluate how well WLAN-Opp utilizes the contacts given in real world mobility traces. The fundamental mechanism that influences the contact utilization ratio is the duty cycling of operations for connectivity provisioning such as becoming an AP. These operations consume a considerably large amount of energy. In the following, we evaluate the impact of duty cycling on the energy consumption and utilization ratio of different contact traces.

4.1. Contact traces

We evaluate the WLAN-Opp state machine under mobility by replaying four real world contact traces presented in the following. These traces differ in terms of size, duration, and contact frequency. Their main characteristics are summarized in Table 2.

H06: During Infocom 2006, contact traces of 78 conference attendees wearing Bluetooth devices were collected for the Haggle project [3]. The granularity of the direct contacts collected on four continuous days is given by the two minute scanning interval of the devices. In total there are 128,979 contacts.

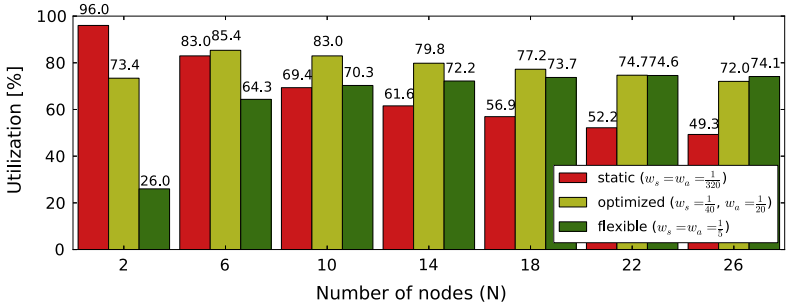

Fig. 5. Utilization ratio depending on number of co-located users and 'flexibility' due to the chosen parameter setting: red bars ("static", left) show a static configuration, green bars ("dynamic", right) show a very flexible configuration, and yellow bars ("optimum", middle) show the parameter setting of Sections 3.2 and 3.3.

Table 2
Properties of contact traces.

	H06	MIT	ETH	SF
# Nodes	78	96	20	536
Time period	93 h	14.9 weeks	104 h	24 days
Type	Bluetooth	Bluetooth	Wi-Fi ad-hoc	GPS
Scanning interval	2 min	5 min	2 s	30 s
# Contacts total	128,979	75,425	18,293	1,332,261
# Contacts/node	1654	786	915	2486

MIT: The MIT Reality Mining project [7] collected Bluetooth contacts with a five minute scanning interval of 96 students and staff members on the MIT campus during several month. We extract the 15 weeks of the trace with the highest contact density. We looked at a total of 75,425 contacts.

ETH: The ETH trace was collected on the ETH campus in 2005, from 20 researchers carrying a PocketPC [16]. Scanning was done via beacons over Wi-Fi ad-hoc with a two second interval. There are a total of 18,293 recorded contacts.

SF: The San Francisco taxicab trace contains GPS position recordings for over 500 taxicabs over a period of a month [21]. The contacts in this trace are inferred based on a conservative transmission range of 30 m. This results in a total of 1,332,261 inferred contacts.

4.2. Duty cycling: contact utilization vs. energy consumption

WLAN-Opp is duty cycled by design, as being discoverable and performing a discovery is not possible at the same time, thus, these phases have to alternate. We can however tune the length of the cycles by adjusting the minimal AP off time $t^{AP}_{off,min}$, i.e., the limit of its exponential back-off max $\left(t^{AP}_{off,min}\right)$. By increasing this time, a device is less frequently in the energy hungry AP state but on the downside it is less discoverable leading to missed contact opportunities. In order to analyze the effect of duty cycling on the contact utilization and energy consumption we apply the simulation model to real world contact traces resulting in the duration each node spends in one of the three states IDLE, STA, or AP. Yet, we also need to know the energy consumption of these states. Table 3 summarizes battery lifetime measurements performed on a Nexus One for all the required states. While the exact amount of energy might vary among phones and batteries, the ratios of the different states provide a more stable and significant measure. The battery consumption in STA mode is 5.97 times higher than in IDLE mode, the one of AP mode is 4.35 times higher than in STA mode. Similar factors were obtained from energy measurements on other phones [26].

Using these energy values, Fig. 6 shows the energy consumption of a node under the mobility defined by the four traces depending on the duty cycle. The energy is expressed as a fraction of the energy the traditional Wi-Fi

Table 3
Energy consumption of different operations expressed as the percentage of a Nexus One battery they consume in one hour.

Operation	Battery consumption
802.11 scanning (IDLE)	0.2%/h
UDP beaconing (STA)	1.19%/h
802.11 SSID beaconing (AP/Ad-Hoc)	5.19%/h

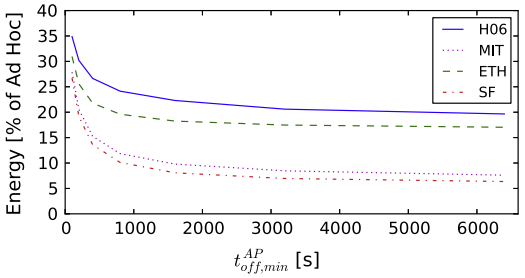

Fig. 6. The effect of duty cycling on the energy consumption for the different traces.

ad-hoc protocol would require. We assume that a node in ad-hoc mode consumes comparable amount of energy to the AP mode as it also advertises a WLAN network by sending out SSID beacons. We can clearly see that for a $\max\left(t^{AP}_{off,min}\right) > 20$ min (1200 s), the energy consumption is almost constant. While the actual consumption depends on how many contacts are present in the trace, the overall consumption is significantly lower than the traditional IEEE 802.11 ad-hoc protocol. While this is not that surprising, given that a node in Wi-Fi ad-hoc mode would need to constantly beacon, we will now investigate the effect of duty cycling on contact utilization.

The contact utilization ratio (Eq. (6)) is measured as the percentage of the aggregated pairwise contact times during which WLAN-Opp manages to connect the pair. Fig. 7 shows that the contact utilization ratio only slightly decreases with a longer duty cycle. If we assume a cycle of 20 min as proposed above, the utilization is between 50% and 80%, depending on the trace. In contrast to Fig. 5 where the utilization ratio was determined by the speed of the AP handover, here the utilization loss is dominated by the effect of duty cycling the discovery process which leads to undetected short contacts. The reason why the SF trace performs so poorly is thus its massive amount of short contacts. Over 20% of the total contact time are short contacts that last less than one minute, which is below the granularity of the MIT and H06 traces. Such contacts can hardly be utilized by WLAN-Opp at all. Generally, traces with fewer short contacts result in a better utilization. In any case, the tremendously reduced energy consumption, resulting in about 10–25% of the amount Wi-Fi ad-hoc networking consumes, justifies the small decrease in contact utilization.

5. Model validation: simulation vs. real world

WLAN-Opp has not only been implemented in simulation, but also on Android-based smartphones. While WLAN-Opp on Android was implemented as a framework that enables the deployment of opportunistic applications, we use it here to validate the simulated state machine of WLAN-Opp. In order to compare the behavior of the state machine in the simulation and the real world, we compare the pairwise *connection* and *inter-connection time* distributions, as it allows best to capture the dynamics of the system. We conduct the experiment on up to ten Galaxy Nexus devices running Android 4.1.2.

We show results based on the parameters chosen in Section 3 (i.e., $\alpha = 2,\ \beta = 0.5,\ w_s = \frac{1}{40},\ w_a = \frac{1}{20}$). Additional experiments with other parameter settings were performed, which further confirm the match of the simulation and the real world implementation. We set the maximal time to stay in AP mode to $t^{AP}_{on,max} = 600$ s in order to get a rather dynamic scenario also for larger numbers of co-located nodes N. The pairwise connection times are thus always below 600 s.[4]

The simulation-based and real world implementations of WLAN-Opp both implement the same state machine. However, the simulation model does not capture all real world phenomena, such as interference. The distributions thus show a high similarity, but are not an exact match of one another. By introducing an additional interference factor to simulation, we show that it is further possible to adjust the simulation to the real world. Since modeling interference is difficult we only simulate the most prominent effect. The interference we implement corresponds to the association process, which commonly takes longer to succeed if more devices try to associate at the same time. Figs. 8 and 9 show the distributions of the *real world* implementation and the simulation-based implementation (*simulation*, cf. Sections 3 and 4), and the simulation with interference introduced above (*simulation (interference)*).

How the connection times are distributed among those 600 s depends on the number of co-located nodes N as shown in Fig. 8. While for all N the simulation closely follows the real world, the impact of real world imperfections are more visible the more co-located nodes are present ($N = 10$). The same is true for the inter-connection time distribution visualized in Fig. 9. Here we can observe a clear improvement in similarity to the real world when applying simulated interference. While generally the differences increase with an increasing N, the simulation manages to closely match the reality.

6. Field trial

To show the feasibility of WLAN-Opp in a real setting, we ran a field trial for five subsequent working days (from Monday morning until Friday evening) in the urban area of Zurich, Switzerland. A total of 34 users equipped with smartphones participated consisting mainly of employees of two neighboring research labs, four employees of a small startup with ties to one of the research labs, and two external persons. Overall, 19 distinct device models from five different vendors were used. A total amount of 94,215 connections were recorded, which is an average of 2771 connections per node. As expected, the number of neighbors of devices varies heavily during one day. Fig. 10 shows the daily patterns for the accumulated amount of connected neighbors of all devices. Office hours of each day are clearly visible. During peak time, the maximum

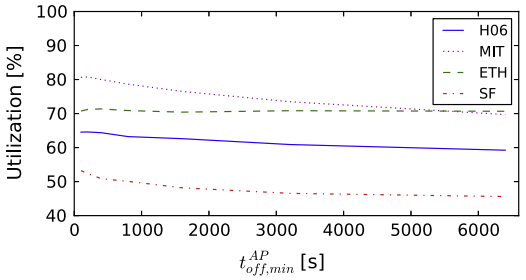

Fig. 7. Contact utilization depending on the duty cycle for the different traces.

[4] Actually, connection times up to 630 s are possible as up to two slot times may not be added to t^{AP}_{on} (one at the beginning and one at the end).

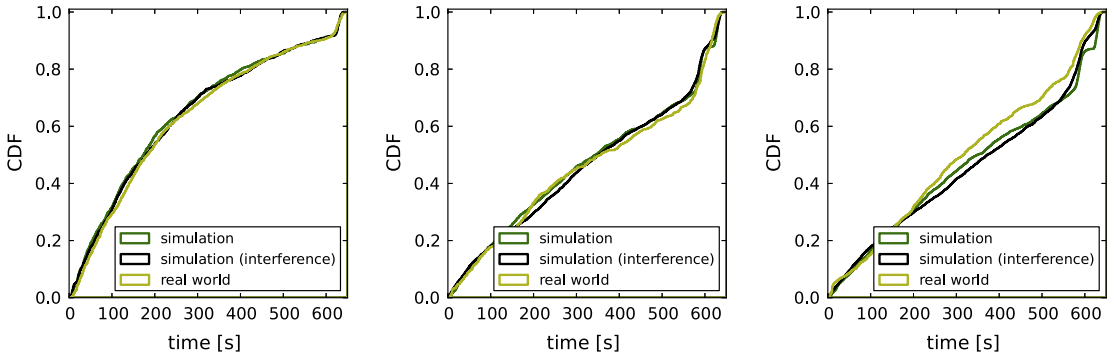

(a) $N = 2$. (b) $N = 6$. (c) $N = 10$.

Fig. 8. Comparison of the pairwise connection time distribution.

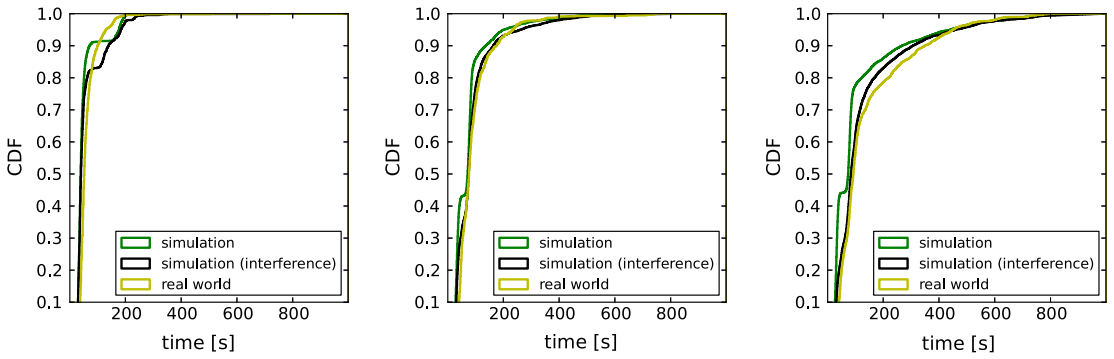

(a) $N = 2$. (b) $N = 6$. (c) $N = 10$.

Fig. 9. Comparison of the pairwise inter-connection time distribution.

number of simultaneously discovered neighbors by a particular device is 15.

In urban settings, many Wi-Fi networks are actually available and some of them do not require credentials for a connection. Such open access points seamlessly integrate with WLAN-Opp as the WLAN-Opp state machine does not differentiate between networks generated by a device running WLAN-Opp or other available and open networks (with the exception that WLAN-Opp knows for sure it will find a neighbor on a WLAN-Opp network). Many of such open access points allow for two associated devices to communicate even if Internet access is restricted by a captive portal. By using open access points, connectivity can be provided in an energy-efficient way as no smartphone has to take the AP role.

A study by Kärkkäinen et al. [13] shows that out of 50 randomly selected open access points in the city of Helsinki, 60% can be used for opportunistic communications.[5] We verified this study on a larger scale in Zurich downtown. Out of 2103 analyzed access points in the city, we were able to communicate through half of them when testing with both, multicast and unicast packet transmissions. Fig. 11 shows all usable networks in Zurich downtown. Given this high coverage of hotspots in Zurich, in our study, WLAN-Opp smartphones operated a considerable time in station mode utilizing infrastructure WLANs to communicate with other WLAN-Opp smartphones.

The main observations of the study are: (i) WLAN-Opp was able to provide opportunistic communication, (ii) it is possible to operate WLAN-Opp without disturbing the user as its only active in the background when the screen is off, (iii) the slightly faster battery depletion was acceptable but should be improved by smarter duty cycling, (iv) in a campus setting or in urban areas, available open access points allow devices to operate most of the time in the less energy intensive station mode, and (v) during AP mode, disabling 3G is required to avoid unwanted tethering costs.

7. Discussion and related work

In the previous sections, we demonstrated that WLAN-Opp is a mature enabling technology for opportunistic networks that can overcome limitations of current technologies. Here, we discuss related opportunistic networking applications, security implications, and energy efficiency issues.

Opportunistic networking: Opportunistic networks have been researched on for more than ten years and

[5] If all possible forms of communications are considered, such as unicast IPv4 and multicast IPv6.

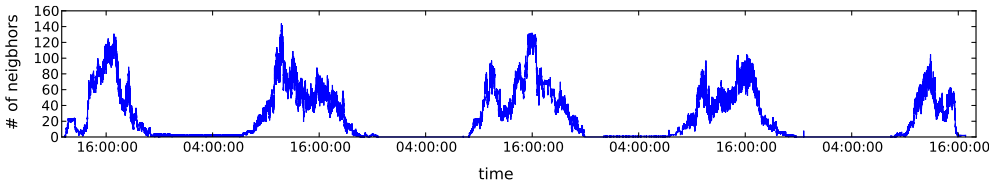

Fig. 10. Daily pattern: Total neighbors encountered by all devices during the field trial.

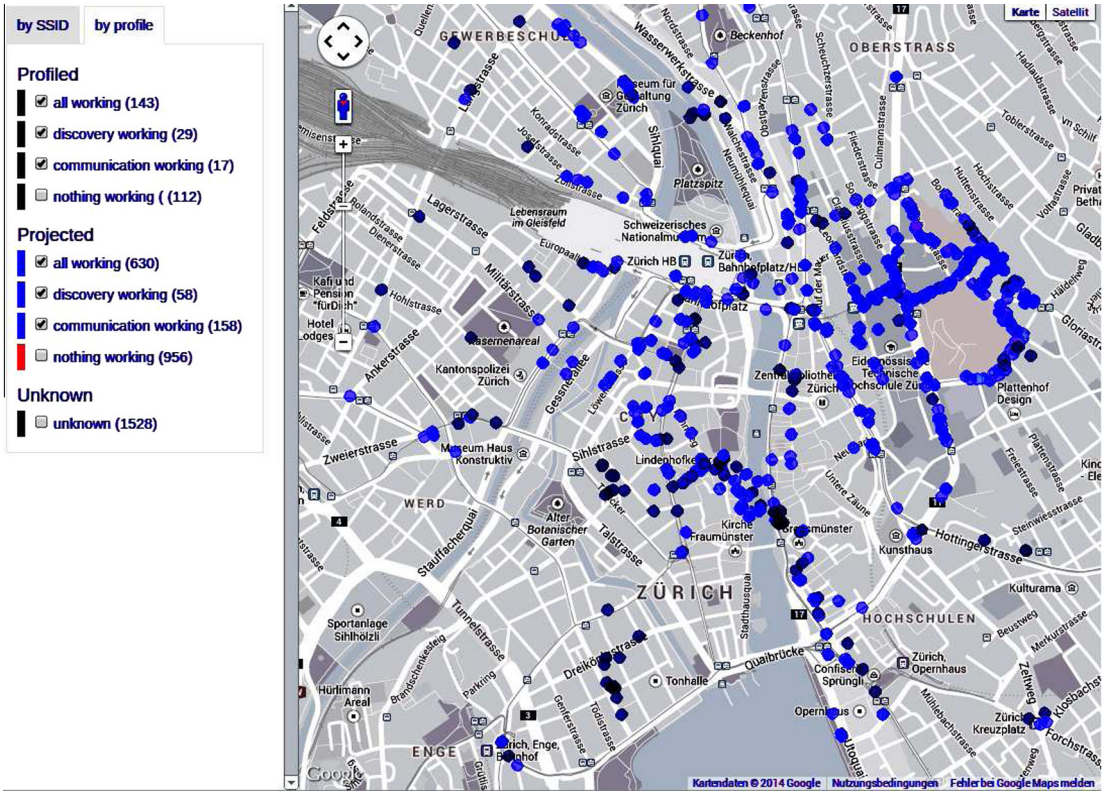

Fig. 11. Open access points in Zurich downtown, which allow local communications.

important theoretic and simulation-based results have been achieved [10,11,22,24]. From the perspective of real applications, no opportunistic application has reached public usage. There are, however, a few notable research prototypes. The three projects 7DS [19], Haggle [23], and PodNet [15], built different types of opportunistic media sharing applications based on Wi-Fi ad-hoc and Bluetooth. However, as these communication technologies are either unavailable on off-the-shelf smartphones or have unrealistic setup requirements involving human interaction, it was difficult for those projects to reach the public. Here, we see the need for a technology such as WLAN-Opp to bridge this gap.

Security: Some ad-hoc protocols such as Bluetooth and Wi-Fi Direct have a strong emphasis on security by enforcing secure pairing. While secure pairing in general provides some degree of security, it hinders the fast and transparent establishment of opportunistic networks. Further, it was shown that Bluetooth does not provide sufficient security despite secure pairing [18]. Thus, with WLAN-Opp we push security from the networking to the application layer as done in the Internet and rely on applications to use end to end security measures, such as TLS [6].

Energy efficiency: WLAN-Opp already provides means to conserve battery consumption by limiting the time a node remains in the energy intensive AP mode. In addition, if no current application requires communication with a neighbor, WLAN-Opp can easily duty cycle neighbor discovery and save battery power. Smartphone measurements have shown that transmitting and receiving data causes nearly three times the energy consumption of the idle state [1]. Further, context awareness could help to reduce the energy required for neighbor discovery [14]. For example, low power sensors, such as the accelerometer could be used to detect whether a node is moving and the discovery intervals may be adapted accordingly.

8. Conclusion and future work

We have proposed WLAN-Opp, a viable and energy-efficient WLAN-based protocol that enables opportunistic communications by exploiting the mobile access point mode of today's smartphones. WLAN-Opp may flexibly switch between access point and station mode depending on other available networks and number of neighbors, constrained by time thresholds that assure that a device is not exploited by the resource consuming access point mode. Randomization of state transitions between the modes of WLAN-Opp prevents from missing contact opportunities and network partitioning. This way, WLAN-Opp can utilize contacts originating from sampled real mobility up to 80% while saving up to 90% of the energy Wi-Fi ad-hoc would consume.

We have implemented WLAN-Opp on Android, currently the most used smartphone platform. In a field trial with 34 participants observed over five days, WLAN-Opp demonstrated its practicality in real scenarios. While we finally have a platform that can clearly promote opportunistic networks, there is still room for future work on optimizing WLAN-Opp. We already tackled aspects of energy fairness and efficiency in WLAN-Opp [26]. A next step would be to analyze the impact WLAN-Opp has on delay tolerant routing and on traffic forwarding protocols exposed to real traffic demands.

Acknowledgements

This work is funded by ETH Zurich under Research Grant ETH-20 09-1 and by the EU Commission under the SCAMPI (FP7-258414) FIRE Project and the FP7 Marie Curie IEF program (PIEF-GA-2010-276336MOVE-R).

References

[1] A. Carroll, G. Heiser, An analysis of power consumption in a smartphone, in: USENIX Annual Technical Conference, 2010.
[2] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott, Pocket Switched Networks: Real-world Mobility and Its Consequences for Opportunistic Forwarding. Technical Report 617, University of Cambridge, 2005.
[3] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, J. Scott, Impact of Human Mobility on Opportunistic Forwarding Algorithms, vol. 6, IEEE TMC, 2007.
[4] T.M. Chen, Governments and the executive 'internet kill switch', IEEE Netw. 25 (2) (2011) 2–3.
[5] M. Dekker, C. Karsberg, Annual Incident Reports 2011, Technical Report October, ENISA, 2012.
[6] T. Dierks, E. Rescorla, The Transport Layer Security (TLS) Protocol, RFC 5246, 2008.
[7] N. Eagle, A.S. Pentland, Reality mining: sensing complex social systems, Person. Ubiquitous Comput. 10 (4) (2006) 255–268.
[8] B. Han, P. Hui, M.V. Marathe, G. Pei, A. Srinivasan, V. Anil, Cellular traffic offloading through opportunistic communications: a case study, in: Chants, 2010, pp. 31–38.
[9] M. Helft, D. Barboza, Google Shuts China Site in Dispute over Censorship, The New York Times, 2010. March(22)..
[10] O.R. Helgason, F. Legendre, V. Lenders, M. May, G. Karlsson, Performance of opportunistic content distribution under different levels of cooperation, in: European Wireless, 2010.
[11] T. Hossmann, T. Spyropoulos, F. Legendre, Know thy neighbor: towards optimal mapping of contacts to social graphs for DTN routing, in: INFOCOM, 2010.
[12] IEEE-SA, IEEE 802.11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications, 2007.
[13] T. Kärkkäinen, M. Pitkänen, J. Ott, Enabling ad-hoc-style communication in public WLAN hot-spots, in: Chants, 2012.
[14] K.-H. Kim, A.W. Min, D. Gupta, P. Mohapatra, J. Pal Singh, Improving energy efficiency of Wi-Fi sensing on smartphones, in: INFOCOM, April 2011.
[15] V. Lenders, G. Karlsson, M. May, Wireless ad hoc podcasting, SIGMOBILE Mob. Comput. Commun. Rev. 12 (1) (2008) 65–67.
[16] V. Lenders, J. Wagner, M. May, Measurements from an 802.11b mobile ad hoc network, in: IEEE EXPONWIRELESS, 2006.
[17] F. Malandrino, M. Kurant, A. Markopoulou, C. Westphal, U.C. Kozat, Proactive seeding for information cascades in cellular networks, in: INFOCOM, March 2012, pp. 1719–1727.
[18] T.C. Niem, Bluetooth and its inherent security issues, SANS InfoSec Read. Room (2003).
[19] M. Papadopouli, H. Schulzrinne, Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices, in: MobiHoc, 2001.
[20] L. Pelusi, A. Passarella, M. Conti, Opportunistic networking: data forwarding in disconnected mobile ad hoc networks, Commun. Mag. 44 (11) (2006) 134–141.
[21] M. Piorkowski, N. Sarafijanovic, M. Grossglauser, A parsimonious model of mobile partitioned networks with clustering, in: IEEE COMSNETS, 2009.
[22] T. Spyropoulos, K. Psounis, C.S. Raghavendra, Efficient routing in intermittently connected mobile networks: the multiple-copy case, IEEE/ACM ToN 16 (1) (2008) 77–90.
[23] J. Su, J. Scott, P. Hui, J. Crowcroft, Haggle: seamless networking for mobile applications, in: UbiComp, 2007.
[24] E. Talipov, Y. Chon, H. Cha, Content sharing over smartphone-based delay-tolerant networks, IEEE TMC (2011).
[25] S. Trifunovic, B. Distl, D. Schatzmann, F. Legendre, WiFi-Opp: ad-hoc-less opportunistic networking, in: Chants, 2011, pp. 37–42.
[26] S. Trifunovic, A. Picu, T. Hossmann, K.A. Hummel, Adaptive role switching for fair and efficient battery usage in device-to-device communication, ACM SIGMOBILE Mobile Comput. Commun. Rev. 18 (1) (2014) 25–36.
[27] Wi-Fi-Alliance, Wi-Fi Peer-to-Peer (P2P) Technical Specification, 2011. <http://www.wi-fi.org/discover-wi-fi/wi-fi-direct>.

Sacha Trifunovic received his MS in Electrical Engineering at ETH Zurich in 2010. Since then he is a PhD candidate at the Communication Systems Group at ETH Zurich. His research interest lie in opportunistic networking.

Maciej Kurant received his PhD degree at EPFL, Lausanne, Switzerland, in 2009. Next, he spent 2 years at University of California, Irvine. In 2012 he was a Senior Researcher and Lecturer in the Communication Systems Group (CSG) of ETH Zurich. Currently he is a Software Engineer at Google, Zurich.

Karin Anna Hummel is a senior researcher and lecturer at ETH Zurich, Communication Systems Group. She received her PhD in Computer Science from the Vienna University of Technology in 2005. Her main research interests include ad-hoc and opportunistic networking, energy-efficient wireless networking, aerial communications, and mobility characterization. She is author of more than 60 peer-reviewed works on mobility-aware computing, mobility modeling, and wireless networking.

Franck Legendre received a degree in Telecommunication engineering from the Institut National des Télécommunications (INT), Evry and a MSc in Computer Networks from the Université Pierre et Marie Curie Paris Universitas, in 2002. He received his PhD from Université Pierre et Marie Curie Paris Universitas in 2006. PhD works were conducted in the Networks and Performance Analysis team under supervision of Professor Serge Fdida and CNRS research scientist Marcelo Dias de Amorim. From 2007 to 2013 he was a senior researcher and lecturer in the Communication Systems Group at ETH Zurich lead by Prof. Bernhard Plattner. He now is CTO at Uepaa!, Zurich.