

Efficient Neighbor Discovery in Mobile Opportunistic Networking using Mobility Awareness

Andrea Hess*[‡], Esa Hyytiä*, Jörg Ott*

*Aalto University (Comnet), Email: firstname.lastname@aalto.fi

[‡]University of Vienna, Email: andrea.hess@univie.ac.at

Abstract—To detect peers in mobile opportunistic networks, mobile devices transmit and listen for beacons (“scanning”). If networks are sparse, devices spend quite a bit of energy scanning the vicinity for possible contacts with their radios. Numerous techniques were developed to adapt the scanning intervals as a function of the observed node density. In this paper, we complement such techniques by considering that protocol exchanges between nodes require contacts of a minimal time span and infer scanning opportunities from node mobility. The adaptive beaconing presented in this paper reduces the scanning effort significantly without “losing” many contacts that last long enough to (i) fully establish an ad-hoc connection between two devices and to (ii) transfer a sizeable amount of data. We propose a theoretical model to derive connection probabilities from sojourn times in different mobility settings and evaluate the impact on energy consumption and data forwarding performance using simulations with different mobility models.

I. INTRODUCTION

Mobile opportunistic networks are formed between mobile devices using short-range radio technologies such as IEEE 802.11 WLAN or Bluetooth. To determine if there are peer devices (“contacts”) around to communicate with, a device needs to activate its radio, send beacons, and listen for beacons sent by others. This process is costly in terms of energy consumption, especially if—as generally assumed—the density of devices supporting opportunistic communication is sparse and the number of contacts found per beacon sent is small.¹ This calls for carefully choosing the radio activation periods for reception and the beaconing intervals, e.g., as a function of past observations (see Section II). In this paper, we explore a complementary strategy for adaptive contact probing that uses local knowledge about the velocity at which a node is moving (or if moving at all). We start with two observations:

1) Contact durations between people who are moving will often be short: only people who walk in the same direction at roughly the same speed will remain in contact for a while. For example, even people walking at a slow pace of 0.5 m/s (1.8 km/h) in opposite directions could, assuming 10 m radio range, only be in contact for 20 s. Figure 3(c) provides examples for contact durations for several mobility models, which show that a notable fraction of the contacts will likely be too short for effective information exchange.

2) Experiments have shown that it takes a while to set up a connection between nodes. In addition to the discovery process, two nodes need to run some form of autoconfiguration

(at minimum duplicate address detection), establish a transport connection, and carry out a handshake including an exchange about routing information and buffered messages. Then, they need to choose and transfer a subset of the messages. This implies a lower limit for effectively usable contact durations.

These observations motivate us to limit radio activation and beaconing to those periods, when the mobile nodes are moving very slowly or not at all. When they move faster, even if a node pair established a contact, the residual contact time may not be sufficient for any substantial data exchange. Following this idea enables nodes to focus on those contacts that are likely to allow exchanging larger data volumes (and, at the same time, reduce unsuccessful communication attempts that cost energy and cause unnecessary interference). This will obviously make the nodes miss some communication opportunities, but these false negatives are largely limited to short contacts and thus have only a small impact on the network performance.

In the following, we first review related work in Section II. Section III outlines our adaptive beaconing mechanism, which combines (i) mobility awareness and (ii) expected latencies in connection setup. We give an analytical model and a numerical performance evaluation in Section IV, and a simulation-based evaluation in Section V. Section VI concludes the paper.

II. RELATED WORK

We briefly review related work on adaptive probing for opportunistic, wireless sensor, and mobile ad-hoc networks.

In [1], the authors first derive the optimal probing interval from contact rates analytically. Real-world contact processes are studied in a Bluetooth experiment. The resulting algorithm is compared to two heuristics, (i) an estimation of contact rates based on day time and (ii) an additive-increase-multiplicative-decrease scheme. In [2], an analytical model for message dissemination under different controlled beaconing policies for epidemic forwarding is presented. The energy costs considered are for beaconing and communication (message transmission and reception). The model is validated by first deriving the exponential distributions for inter-contact time in two mobility scenarios, which are then used as input to verify the infection ratio as a function of time against simulation. Based on the Shanghai taxi traces, the optimal beaconing policy is evaluated for vehicular networking. The policy limiting the fraction of nodes (20%) involved in message dissemination is evaluated against a static policy with constant beaconing rate and a benchmark policy with maximal rate. The beaconing protocol in [3] adapts the beaconing interval based on the contact discovery rate in the previous periods. In [4],

¹Note that “scanning” for, e.g., WLAN access points consumes less resources because devices can remain passive while the mains-powered APs transmit beacons.

a regulated beaconing approach for two-hop forwarding is optimized in terms of throughput for an energy-constrained network. Beacon transmission is controlled by introducing three node states: inactive (not participating in communication), active (beaconing), infected (not beaconing anymore). A power saving mechanism using both low- and high-power radio is proposed in [5]. Contacts are discovered by the short-range radio (alternating between sleep and awake modes), while the high-power radio is activated for data transmissions.

Optimal beaconing has been considered also in the context of wireless sensor networks, see, e.g., [6], [7]. When the beaconing intervals are somehow limited, one essentially makes a trade-off between the energy consumption and latency. In [8], a novel scheme to this end is proposed, which has also been implemented for IEEE 802.11 and IEEE 802.15.4 communication in Linux and TinyOS operation systems. Conversely to our work, adaptive versions of MANET protocols with increased message rates at higher node mobility (e.g., HELLO messages in AODV or OLSR) are proposed to keep track of more dynamic network topologies, see, e.g., [9].

III. MOBILITY- AND PROTOCOL-AWARE BEACONING

The basic idea of our adaptive beaconing scheme is simple: Nodes should avoid spending effort and energy on attempting to exploit contacts that will likely be too short for successful message exchanges anyway. As a simple approximation, nodes should only send beacons when they move slowly or are stationary (during sojourn times). We assume that nodes can measure their velocity (e.g., using GPS²) or at least determine when they are moving or stationary (e.g., using accelerometers, observing “visible” cell towers and transitions between them or fingerprints of access points [11]). While we focus node velocity in absolute terms, nodes may also be stationary relative to their environments, e.g., on public transport or an escalator, and sensing algorithms are available to determine a node’s transportation context [12].

A crucial factor in successful message exchange is the time necessary to establish an ad-hoc connection between two nodes. Experiments showed that we expect a delay of 10–30s can be from scanning the vicinity and discovering a neighboring node until the connection is established and ready for data exchange. For WLAN ad-hoc networks, Roy measured a duration of some 15s from scanning start to IP address verification, shown in Table I [13]. For Bluetooth, Pietiläinen et al. [14] measured a device discovery delay of some 10s plus additional several seconds per discovered device for connection setup and service discovery. They also report low connection success rates (some 15%) for the Bluetooth experiments, leading to additional variable delays caused by each failing connection attempt (14s on average with similar high standard deviation).

In DTN, beacons are usually sent in the order of a few seconds. While the length of a beaconing interval is not specified in the DTN IPND protocol description, the RFC for the MANET NHDP protocol [15] proposes an interval

²One approach to avoid high energy consumptions is to sample GPS coordinates only while the person is moving, indicated through a low-energy sensor (e.g., the phone’s accelerometer), and only in low rates (e.g., 20 seconds) as proposed in [10].

TABLE I. DURATION OF AD HOC CONNECTION ESTABLISHMENT

WLAN		Bluetooth	
Scanning phase	< 1 s	Discovery phase	10 s
Joining network	4 s	Connection setup	3 s p. device
Time gap	10 s	Service discovery	1 s p. device
Selecting/verifying IP address	1 s		
	15 s		10 + 4 × x s

of 2s. This beaconing interval is in our approach adapted to different *mobility-awareness configurations*. We will study configurations where nodes broadcast beacon messages if they move with a velocity v that is less than or equal to a velocity threshold θ_v . In the case of $\theta_v = 0$ m/s, nodes scan their neighborhood only while they are sojourning at a location, $v = 0$. Additionally, we explore thresholds for slow movement. Note that θ_v is chosen configured for pedestrian velocity, but could be adapted to other mobility scenarios (and device radio capabilities). This threshold is also applied in the *radio (de)activation* scheme, where a node’s radio interface is turned off as soon as θ_v is reached. This deactivation setting is evaluated against an *always-on* setting comprising mobility-aware beaconing and, while $v > \theta_v$ holds, passive responsiveness to beaconing.³

IV. ANALYTICAL OBSERVATIONS

In this section, we consider the proposed scheme analytically. For simplicity, we assume $\theta_v = 0$, i.e., that nodes communicate only when they are stationary. We derive expressions for the probability of a successful contact when a node stops, and for the mean rate of such contacts.

A. Pause Times and Random Observer

Suppose that the velocity threshold for communication is set to zero, $\theta_v = 0$. Let us further assume that the nodes move independently and the pause time durations are also independent and identically distributed $\tau_i \sim \tau$. That is, the length of the pause does not depend on the location, time, or node id. Then $f(t)$ denotes the PDF of the pause time distribution, so that, e.g., the mean is $E[\tau] = \int_0^\infty t f(t) dt$.

Suppose that a node A has just stopped and it sends a beacon signal in hopes of finding other nodes currently stationary in the vicinity. Let random variable N denote the number of nodes currently within the transmission range. The pause time of node A obeys the original pause time distribution, $\tau_A \sim \tau$. However, the key observation is that the remaining pause time of a node B does not. This is the well-known *hitchhiker’s paradox* due to the fact that it is more likely to hit a long interval than a short. It follows that the pause time distribution of node B, τ_B , has the PDF of

$$g(t) = \frac{t f(t)}{E[\tau]}.$$

The “random observer”, i.e., the node A, has arrived at a random moment, that is the remaining pause time of node B, denoted by R_B , is uniformly distributed on $(0, \tau_B)$. Letting Δ denote the time required for a successful beaconing, *the probability of a successful beaconing* is,

$$\beta(\Delta) \triangleq P\{\tau_A > \Delta\} \cdot P\{R_B > \Delta\},$$

³A practical implementation would choose some hysteresis function to prevent oscillation for velocities around θ_v .

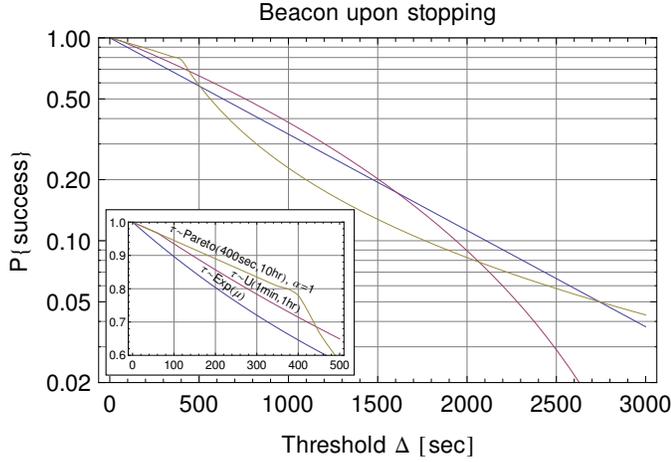


Fig. 1. The probability of a successful beacon upon stopping (per node in the vicinity), p , for three pause time distributions: uniform(1min,1hr) and truncated Pareto in interval (400,36000) with $\alpha = 1$. Note that each have approximately the same mean pause time of about half an hour.

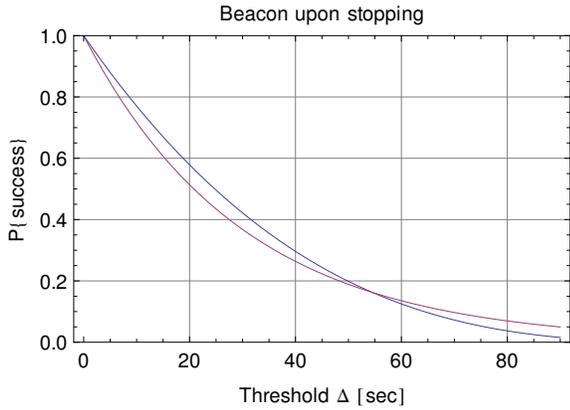


Fig. 2. Here $E[\tau] = 1$ min and distributions are uniform and exponential.

as the two nodes behave independently. The first factor is given, and for the second we obtain

$$\begin{aligned} P\{R_B > \Delta\} &= \int_{\Delta}^{\infty} g(t) \cdot P\{R_B > \Delta \mid \tau_B = t\} dt, \\ &= \frac{1}{E[\tau]} \int_{\Delta}^{\infty} f(t) \cdot (t - \Delta) dt. \end{aligned}$$

Thus, the probability of a successful beaconing is

$$\beta(\Delta) = \frac{1}{E[\tau]} \cdot \int_{\Delta}^{\infty} f(t) dt \cdot \int_{\Delta}^{\infty} f(t) \cdot (t - \Delta) dt. \quad (1)$$

Note that if the pause times obey exponential distribution, $\tau \sim \text{Exp}(\mu)$, then, due to the lack of memory property of the exponential distribution, one trivially obtains $\beta(\Delta) = e^{-2\mu\Delta}$. Similarly, if $\tau \sim U(0, \tau_{\max})$, then one obtains

$$\beta(\Delta) = \begin{cases} (1 - \Delta/\tau_{\max})^3, & \Delta < \tau_{\max}, \\ 0, & \Delta \geq \tau_{\max}. \end{cases}$$

Fig. 1 illustrates how the probability p behaves with three sample distributions as a function of the threshold Δ . In Fig. 2,

the mean pause time is considerably shorter and the probability of successful beaconing upon stopping decreases sharply.

B. Mean contact rate

Then we assume that the locations where nodes stop are uniformly distributed with density n_p on some region with area A , i.e., there is an average $A \cdot n_p$ stationary nodes. As the nodes move independently, the mean number of nodes within the transmission range of a stopping node is then (approximately due to neglecting the boundary),

$$N_p \approx \pi d^2 \cdot n_p.$$

Then we note that under these ideal assumptions, new contacts can emerge only when a node stops and it discovers its new neighborhood. Letting λ_p denote the rate at which nodes stop moving, the rate of contacts longer than Δ is given by

$$\lambda_C \triangleq \lambda_p \cdot N_p \cdot \beta(\Delta). \quad (2)$$

Given the pause time distribution, mean number of neighbors and the frequency of pauses, (2) and (1) give us the rate at which successful contacts between two nodes are established in the given area.

C. Mobility models and beaconing

We choose a number of mobility models from literature to evaluate the impact of our adaptive beaconing mechanism. We include a broad spectrum ranging from the old but well-known Random Waypoint model to map-based models as well as local movement models:

Map-based model. Here, the nodes walk between Points of Interest (POIs) on the shortest path on the Helsinki city center map (4500x3400 m). The pause times are uniformly distributed between 1 min and 1 h. Figure 3(a) and 3(b) depict the Complementary Cumulative Distribution Functions (CCDFs) for pause time and velocity prevailing in each model.

Working Day Movement (WDM) model [16]. In the WDM model, the nodes move between home, office, and evening activity sub-models as pedestrians who might also use a public means of transport (in case the trip duration can be reduced by using buses arriving at a nearby station). The Helsinki city map is here used as well and locations are assigned to each activity. The pause time depends on the activity, e.g., during working time nodes stay continuously at the office for up to 4 h, during evening activities up to 2 h at a POI.

SLAW model [17]. The Self-Similar Least-Action Walk model is based on properties of human movement found in GPS traces (heavy-tail flights and pause times, heterogeneously bounded mobility areas, etc.). Pedestrians move between POIs placed in clusters on the area. Each node is assigned a set of POIs which are visited in a total trip length minimizing manner. For simulating the SLAW model we generated traces using BonnMotion [18] for the same rectangular 4500x3400 m area. The pause time obeys a truncated Pareto distribution with minimum 30 s and maximum 700 min.

Movement Activity (MA) model [19]. The MA model mimics the movement behavior of activities, such as way to work, shopping, tourists, and evening. The pause time follows

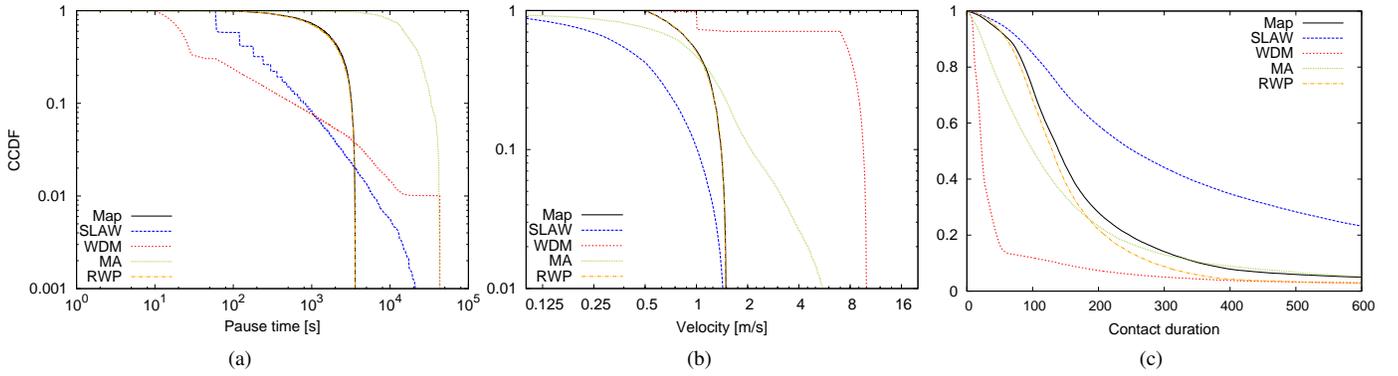


Fig. 3. CCDF for (a) pause time, (b) velocity, and (c) contact duration observed in each mobility model.

a Pareto distribution, its parameters depend on the movement activity of a node ($41 \leq s \leq 123$; $1.3 \leq \alpha \leq 4.3$).

Random WayPoint (RWP) model. This model generates paths between randomly selected waypoints on the simulation area. The velocity is chosen from a uniform distribution. The optional pause times at the waypoints are as well uniformly distributed. For comparison purposes, we use the same minimum and maximum values as for the Map-based model.

The contact duration distributions emerging from these models are summarized in Figure 3(c). While the longest contacts are established in the SLAW setting (the long pause times allow contact durations of up to 7 h), the (faster) public transport mobility in WDM leads to a high fraction of short contacts (69% below 60 s). Due to the joint pause time and velocity setting, the RWP curve closely follows the one for Map, but diverges at some point (upper 20% are above 246 s in Map, above 162 s in RWP). In addition, the difference in absolute contact numbers has to be born in mind, which is about 50 000 for Map and 20% less for RWP.

D. Numerical example: RWP model

The main strength of RWP is the simplicity which facilitates analytical results, and consequently, also rational values for the mobility parameters. Let ℓ denote the mean leg length. The mean travel time between two waypoints is

$$E[T_\ell] = \ell \cdot E[1/v],$$

where v denotes the leg specific velocity [20]. Letting K denote the number of nodes, we have

$$\lambda_p = \frac{K}{E[\tau] + \ell \cdot E[1/v]}.$$

Consequently, the mean contact rate is

$$\lambda_C = \lambda_p \cdot (K - 1) \cdot q \cdot \beta(\Delta),$$

where q denotes the probability that another node B is stationary and within the transmission range,

$$q \approx \frac{\pi d^2}{A} \cdot \frac{E[\tau]}{E[\tau] + \ell \cdot E[1/v]}.$$

Assuming $4500 \text{ m} \times 3400 \text{ m}$ area, velocity $v \sim U(0.5, 1.5)$ m/s and $\tau \sim U(1 \text{ min}, 1 \text{ h})$ pauses (cf. Section IV-C), we have $\ell = 2070 \text{ m}$, $E[1/v] \approx 1.1 \text{ m/s}$, and $P\{\text{stationary}\} \approx 0.45$, i.e.,

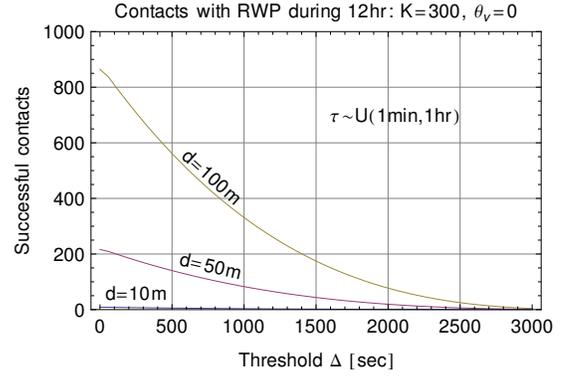


Fig. 4. The mean number of contacts longer than Δ with RWP for 300 nodes in $4500 \text{ m} \times 3400 \text{ m}$ area, $v \sim U(0.5 \text{ m/s}, 1.5 \text{ m/s})$ and $\tau \sim U(60 \text{ s}, 3600 \text{ s})$.

about 45% of the nodes are stationary at a random point of time. Moreover, $q \approx 9.16 \times 10^{-8} d^2$. For $\beta(\Delta)$, we have

$$\beta(\Delta) = \begin{cases} 1 - \Delta/1830, & \Delta \leq 60, \\ \frac{(\Delta - 3600)^3}{45865656000}, & 60 < \Delta \leq 3600, \\ 0, & \text{otherwise.} \end{cases}$$

For example, assuming $\Delta \leq 60$, we have explicitly

$$\lambda_C \approx 2.23 \times 10^{-11} K(K - 1)(1 - \Delta/1830).$$

Fig. 4 depicts the mean number of successful contacts during a 12 hours for the transmission radii of $d = \{10, 50, 100\} \text{ m}$. We can observe that only with $d = 100 \text{ m}$ there seems to be a reasonable number of sufficiently long contacts. The same observation will be made also in the next section with simulation experiments.

V. SIMULATION-BASED EVALUATION

We now present the results of an extensive simulation study done with the ONE Simulator [21] – Table II summarizes the parameters. We choose the following metrics: 1) *Total usable contact time*: the overall duration of established contacts relative to the overall duration of established contacts in the setting without any mobility awareness (w/o). 2) *Contact time per beacon ratio*: the relative improvement for contact duration per beacon message ratio γ for introducing mobility awareness (ma): $\gamma_{ma}/\gamma_{w/o}$. 3) *Energy consumption for beaconing and active radio interface*: the mean power consumption \bar{C}_{ma} for

the *ma*-setting compared to the power consumed without *ma* $\bar{C}_{w/o}$. 4) *Message delivery ratio*: the delivery probability in two message forwarding scenarios, Spray-and-Wait [22] and Epidemic forwarding [23].

TABLE II. GENERAL PARAMETER SETTING.

General	
Simulation time	12 h
Node number	300
Area size	4500 m x 3400 m
Transmission range	50/100 m
Connection establishment	
Beaconing interval t_b	2 s/6 s
Connection setup delay t_{setup}	0/10/20/30 s
Successful connection threshold Δ	$t_{setup} + t_{data}$

A. Total usable contact time

Fig. 5 and 6 show the fraction of *usable connection time* for mobility-aware (*ma*) probing relative to the results without mobility awareness for $\Delta = 0..300$ s and two transmission ranges. This metric gives us an idea of the overall network capacity as, given a constant transmission rate, connection time can be translated to pairwise transmission capacity. The plots on the left-hand side depict the usable connection time fraction for mobility awareness with radio deactivation for a velocity threshold θ_v of 1.0 m/s (Fig. 5(a)) and θ_v of 0.0 m/s (Fig. 5(c)), while Fig. 5(b) and 5(d) show the numbers without radio deactivation for these θ_v values.

The main visible trend suggests that the fraction of usable contact time increases with connection threshold Δ , as the number of (too) short contacts should be reduced by our beaconing approach. The contact time for the least restrictive setting (radio always-on and velocity threshold 1.0 m/s) shown in Fig. 5(b) deviates within only 5% from the total possible contact time. In the most restrictive setting (Fig. 5(c)) with deactivation and θ_v of 0.0 m/s, the curves are highly variable and only 20% of the possible contact time would be available in some mobility models (Map, RWP) if there is no connection success threshold. With increasing Δ the available share rises to 60% and more, but stays low for SLAW mobility. When setting the transmission range to 50 m (Fig. 6), a very similar picture of usable contact time fractions emerges. However, it has to be noted that the absolute total contact time for the 50 m range is about one-fifth of the 100 m configuration.

B. Adjusting the beaconing interval t_b

Fig. 7 gives the hourly number of beacons sent with its standard deviation. These numbers basically correspond to the fraction of time nodes are stationary with each mobility model, or moving with less than 1 m/s respectively. Since our approach aims to be efficient by detecting long contacts with less beaconing effort, it is reasonable to adapt also the beacon interval t_b to observable sojourn times. We therefore define a rule of thumb based on two values: (i) the length of pauses of interest (we use here the upper 95% fraction of pause times) and (ii) the connection setup time:

$$t_b \leq \frac{\max\{t^*, t_{setup}\}}{10}, \quad \text{and} \quad P(\text{pause} < t^*) = 5\%.$$

We additionally evaluate the case of $t_b=6$ s, since 5% of the pause times are smaller than 60 s.

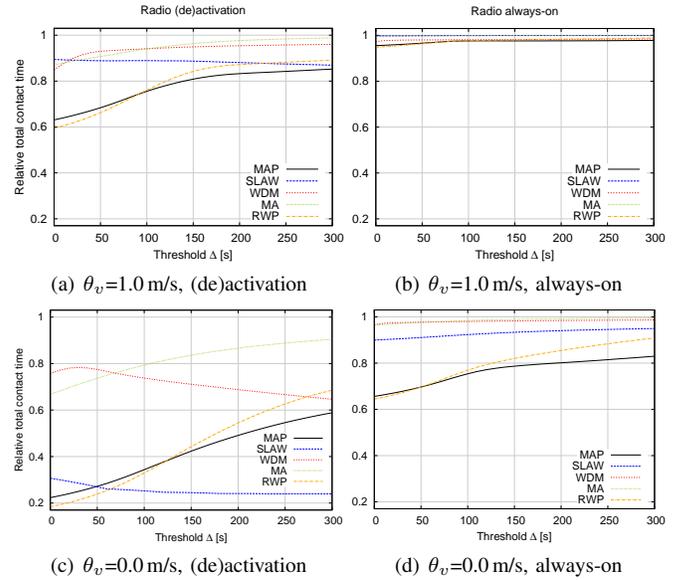


Fig. 5. 100 m range: share of total duration of established contacts longer than threshold Δ .

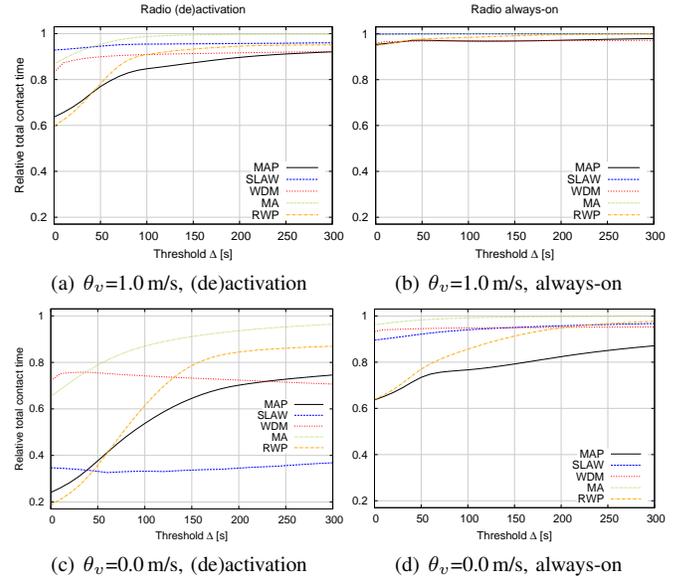


Fig. 6. 50 m range: share of total duration of established contacts longer than threshold Δ .

The resulting fraction of usable contact time, compared to the *w/o ma*-setting with beaconing interval t_b of 2 s exhibit a similar curve picture as in Fig. 5 suggesting that only few contacts are shortened because detected later. For comparing the two t_b configurations, we take a look at the average *transmission capacity loss* per contact given a bitrate of 2 Mbps for two *ma*-settings, namely $\theta_v=1.0$ m/s deactivation and $\theta_v=0.0$ m/s always-on. Fig. 8 points out that increasing the beaconing interval has only a small impact in order of kB. In our configuration, up to 13 MB of transmission capacity might drop away while the numbers reduced by an increased connection threshold. The capacity loss shows high variations, when calculating the average over all node pairs. For example, Map at threshold $\Delta=0$ exhibits in the first setting (Fig. 8(a)) a capacity loss of 12.38 MB with a standard deviation of 11.85.

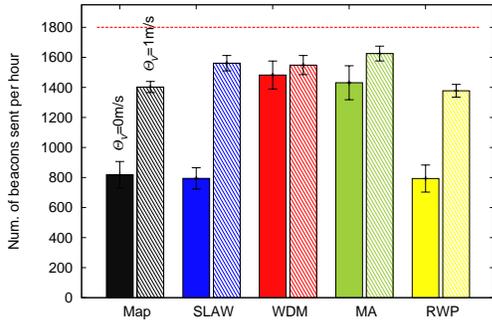


Fig. 7. Average number of beacons sent by each node within one hour. The red line indicates the number without mobility awareness (i.e., $1h/t_b$).

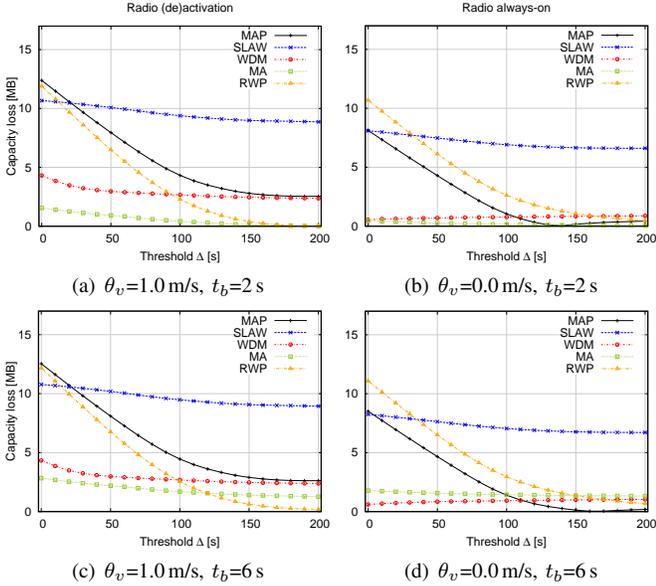


Fig. 8. Average transmission capacity loss per node pair at increasing threshold Δ with radio (a,c) deactivation and (b,d) always-on.

C. Contact time per beacon ratio

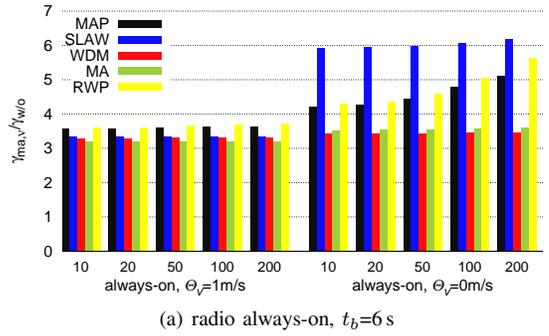
Fig. 9 depicts the improvement of the contact time per beacon ratio for five exemplary connection success thresholds. While the ratio improvement is only slightly influenced by the connection threshold in the $\theta_v=1.0$ m/s settings, a growth is observable if $\theta_v=0.0$ m/s for Map, SLAW, and RWP, particularly in the always-on setting (Fig. 9(a) on the right).

D. Energy consumption

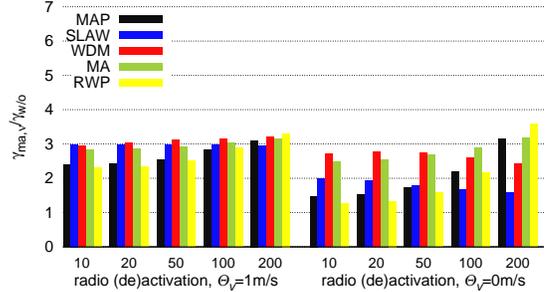
We define a consumption model to calculate the energy needed for sending beaconing messages and for an active radio interface. Three different states are considered: *radio off* (*state 0*), *radio on* (*state 1*), and *transmitting beaconing messages* (*state 2*). The mean power consumption \bar{C} is then

$$\bar{C} = \sum e_i \times p_i.$$

where e_i denotes the fraction of time a node is in state i and p_i is the power consumed in state i . The conditions for switching to a particular state applied in our model are summarized in Table III. For example, without any mobility awareness the mobile device is permanently in *state 2*, and with mobility awareness the state depends on the current velocity v .



(a) radio always-on, $t_b=6$ s



(b) radio (de)activation, $t_b=6$ s

Fig. 9. Relative improvement for contact time per beacon ratio for threshold Δ of 10 s, 20 s, 50 s, 100 s, and 200 s.

TABLE III. ENERGY CONSUMPTION MODEL PARAMETERIZATION

	<i>state 0</i>	<i>state 1</i>	<i>state 2</i>
w/o ma	–	–	$\forall v$
Radio always-on	–	$v > \theta_v$	$v \leq \theta_v$
Radio (de)activation	$v > \theta_v$	–	$v \leq \theta_v$
<i>Active radio</i>			
WLAN ad hoc	–	48.7 mWs	48.7 mWs
Bluetooth	–	89.5 mWs	89.5 mWs
<i>Beaconing</i>			
WLAN ad hoc (UDP)	–	–	1554 mWs / t_b
Bluetooth (RFCOMM)	–	–	634 mWs / t_b

The computation of the energy consumed by the radio interface during beaconing is based on power consumption values measured in [24]. In the study, measurements with six mobile devices (four Android-based, two Windows Mobile-based phones) of different vendors were done. The power consumption for transferring data packets was not measured on the Android-based phones, since WLAN ad hoc mode is not supported there. The power consumption values averaged (i) over six phones for active radio and (ii) over two phones for message transmission are also given in Table III. Without mobility awareness the power consumption \bar{C} for 12 hours activity like in the simulation setting is 9.91 Wh (WLAN ad hoc) and 4.88 Wh (Bluetooth) per device.

Fig. 10 shows the relative power consumption values arising from the four example *ma*-settings. In general, it can be said that deactivating the radio brings little gain, but degrades the contact time per beacon ratio heavily. Comparing top to bottom diagram shows that radio deactivation has a smaller effect on the improvement (\approx few percentages) since beaconing requires more energy than active radio. With a larger beaconing interval $t_b=6$ s, this impact is less significant but still exists. In the always-on case (Fig. 10(a)) the improvement for WLAN is more significant, since Bluetooth consumes more power for listening than WLAN. The reduction of energy

TABLE IV. PARAMETERS FOR PERFORMANCE EVALUATION

General		Spray & Wait	
Message TTL	5 hours	Mode	Binary
Buffer size	5 MB	Message copies	20% of node num.

consumption for WLAN and Bluetooth is almost equal if the radio is deactivated (Fig. 10(b)). Inherently, mobility patterns with high sojourn fraction exhibit only small consumption decreases (Fig. 7). E.g., with WDM, a device would constantly scan the neighborhood while spending the day in the office.

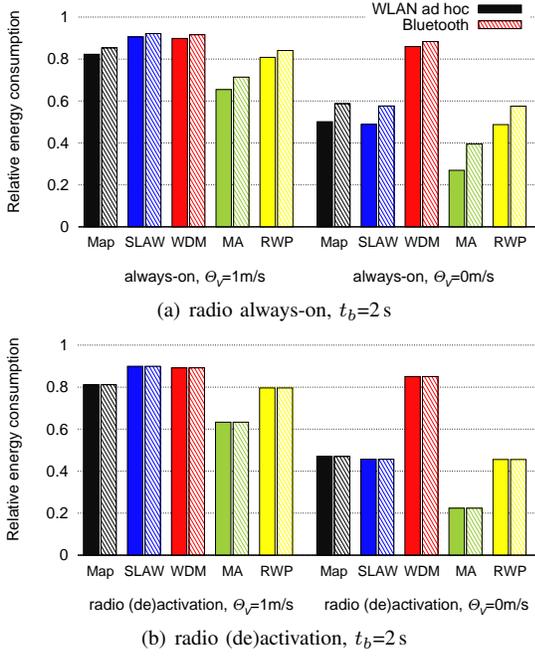


Fig. 10. Relative energy consumed per node for WLAN and Bluetooth beacons and active radio by using mobility awareness.

E. Effects on network performance

Finally, we evaluate how the radio activation and velocity threshold settings affect the performance of two forwarding mechanisms: Spray-and-Wait [22] and Epidemic forwarding [23] – see Table IV for the protocol parameters. Fig. 11 shows the numerical results.

Spray and Wait (S&W). A positive result is that thresholds of up to 100 s cause only slight losses for Map and SLAW. The results for WDM suggest that a TTL of 5 h might be too short for scenarios where people might spend the whole day at office hindering message relaying. The figures also confirm that the setting with radio deactivation at 0.0 m/s is too restrictive and achieves only for SLAW an acceptable ratio.

Epidemic forwarding. Here, each node attempts to transfer all messages currently stored in its buffer during established connections. This leads to long lasting transmissions that are likely aborted by connection teardowns. Further, nodes carrying a fully loaded buffer cannot accept messages of new encounters. Hence, the message delivery is less successful than with S&W. Fig. 11 also shows that the delivery ratio may improve with increased connection setup delay (e.g., for SLAW between $\Delta=0$ and $\Delta=20$), since buffers are less loaded then.

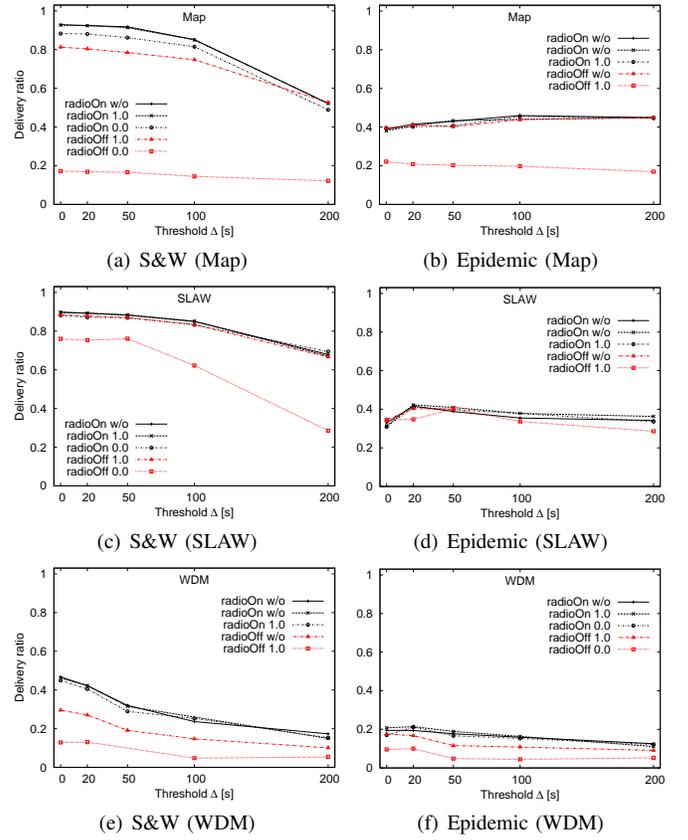


Fig. 11. Delivery ratio in Spray-and-Wait (left) and Epidemic (right) forwarding scenario for three mobility models.

F. Special mobility situations and future work

There are special situations in which people may move a bit, but we still would like our algorithm to work. We envision two different settings: (i) people being stationary relative to their environment (like a train or a bus), and (ii) train or subway stations where a lot of people meet and permanently move although they might be stationary relatively on escalators. For the first setting, context detection works (see, e.g., [12]) are available, where transportation modes are deduced from acceleration readings of mobile devices. In the station setting, the times people are on an escalator would need to be recognized. For a preliminary evaluation of this scenario we rely on a subway station mobility model:

Subway station model [25]. The Station mobility model pictures the pedestrian dynamics within a subway station. Nodes may either arrive through the station entrance and board a train or may exit a train stopping on the track and leave the station through the entrance door. The area of activity has the size of 1921 m². We extracted the movements of 500 nodes from a set of station traces for our experiments, which results in a total simulation time of 1200 seconds.

We modify our algorithm in such a way that, if a node is on an escalator although possibly leading to velocity above the threshold, beaconing is not stopped nor is the radio deactivated. In this mobility scenario, the nodes move faster than the pedestrian velocity assumed in the previous sections (56% of velocity values are above 1 m/s, 20% above 1.5 m/s). Fig. 12(a)

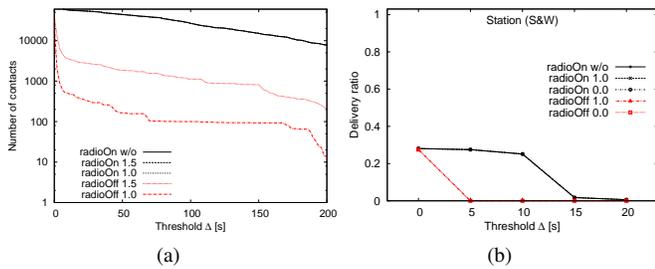


Fig. 12. (a) Number of contact pairs and (b) S&W delivery ratio for Station mobility traces.

reveals that the velocity threshold θ_v has almost no effect in the always-on setting and most faster moving nodes hear beacons of slower moving nodes in the dense station area. Radio (de)activation drops the number of contact pairs with growing connection success threshold drastically, e.g., when $\theta_v=1.5$ from 27 164 ($\Delta=0$ s) to 496 ($\Delta=10$ s).

Fig. 12(b) plots the delivery ratio for Spray and Wait for delays of up to 20 s. In addition to short contact durations, the short activeness time of nodes is impeding data dissemination in this mobility scenario as nodes are for very short time (on average only for 200 s) present in the station area. At most a delivery ratio of 37% is achieved, but the ratio converges to 0 as soon as delays are greater than 10 s. With radio deactivation, an acceptable ratio can only be achieved without any delay. This experiment essentially demonstrates that there are environments where data dissemination does not work well, or even not at all if establishing a link takes more than 10 s.

VI. CONCLUSION

We have proposed a simple adaptive beaconing scheme to improve energy efficiency for peer discovery in mobile opportunistic networks. The scheme makes beaconing a function of node movement, limiting node activity to low (relative) velocities and thereby avoids investing resources to establish contacts that would likely be too short for (substantial) data exchange. While this obviously yields a lower number of contacts in total, we find that longer (and thus more useful) contacts are less affected. We are able to reduce beaconing and the associated energy consumption notably across all mobility models we studied, with limited impacting on the communication performance only. The algorithm complements related work and can be implemented in practice in an energy-efficient way using sensors readily available in mobile devices.

While velocity dependence appears basically workable, further study is needed in a number of directions, including: actually combining our and other algorithms to evaluate the combined effect (and identify potential feature interactions); assessing scenarios in which people move together on public means of transport or escalators; and dealing with errors in pause and movement detection. Finally, we seek doing an implementation for mobile devices to carry out practical experiments to validate our results and calibrate future simulations.

ACKNOWLEDGMENT

This work was supported by the Academy of Finland in the PDP project (grant no. 260014), and by the European Commission within EIT ICT Labs and the FP7 project SCAMPI.

REFERENCES

- [1] W. Wang, V. Srinivasan, and M. Motani, "Adaptive contact probing mechanisms for delay tolerant applications," in *ACM MobiCom*, 2007.
- [2] Y. Li, Z. Wang, D. Jin, L. Su, L. Zeng, and S. Chen, "Optimal Beaconing Control for Epidemic Routing in Delay-Tolerant Networks," *IEEE Trans. on Vehicular Technology*, vol. 61, no. 1, pp. 311–320, 2012.
- [3] B. J. Choi and X. Shen, "Adaptive exponential beacon period protocol for power saving in delay tolerant networks," in *Proc. of IEEE ICC*, 2009.
- [4] E. Altman, A. P. Azad, T. Başar, and F. De Pellegrini, "Optimal activation and transmission control in delay tolerant networks," in *IEEE INFOCOM*, 2010.
- [5] H. Jun, M. H. Ammar, M. D. Corner, and E. W. Zegura, "Hierarchical power management in disruption tolerant networks with traffic-aware optimization," in *Proc. of ACM CHANTS workshop*, 2006.
- [6] R. Musaloiu-E, C.-J. Liang, and A. Terzis, "Koala: Ultra-Low Power Data Retrieval in Wireless Sensor Networks," in *Int. Conf. on Information Processing in Sensor Networks (IPSN)*, 2008, pp. 421–432.
- [7] V. Dyo, S. A. Ellwood, D. W. Macdonald, A. Markham, C. Mascolo, B. Pásztor, S. Scellato, N. Trigoni, R. Wohlers, and K. Yousef, "Evolution and sustainability of a wildlife monitoring sensor network," in *ACM Conf. on Embedded Networked Sensor Systems (SenSys)*, 2010.
- [8] J. A. B. Link, C. Wollgarten, S. Schupp, and K. Wehrle, "Perfect difference sets for neighbor discovery: energy efficient and fair," in *Extreme Conference on Communication (ExtremeCom)*, 2011.
- [9] Y. Huang, S. N. Bhatti, and S.-A. Sørensen, "Adaptive MANET Routing for Low Overhead," in *WOWMOM*, 2007.
- [10] Y. Wang, J. Lin, M. Annavaram, Q. A. Jacobson, J. Hong, B. Krishnamachari, and N. Sadeh, "A framework of energy efficient mobile sensing for automatic user state recognition," in *Proc. of ACM MobiSys*, 2009.
- [11] S. Nawaz, C. Efstratiou, and C. Mascolo, "ParkSense: A Smartphone Based Sensing System For On-Street Parking," in *Proc. ACM MobiCom*, September 2013.
- [12] S. Hemminki, P. Nurmi, and S. Tarkoma, "Accelerometer-Based Transportation Mode Detection on Smartphones," in *ACM Conf. on Embedded Networked Sensor Systems (SenSys)*, 2013.
- [13] S. K. Roy, "Characterizing Opportunistic Contacts," Special Assignment, Aalto University, Tech. Rep., May 2013.
- [14] A.-K. Pietiläinen and C. Diot, "Experimenting with Opportunistic Networking," in *Proc. of ACM MobiArch Workshop*, 2009.
- [15] T. Clausen, C. Dearlove, and J. Dean, "RFC6130: Mobile Ad Hoc Network (MANET) Neighborhood Discovery Protocol (NHDP)," 2011.
- [16] F. Ekman, A. Kernen, J. Karvo, and J. Ott, "Working Day Movement Model," in *ACM SIGMOBILE Workshop on Mobility Models*, 2008.
- [17] K. Lee, S. Hong, S. J. Kim, I. Rhee, and S. Chong, "SLAW: self-similar least-action human walk," *IEEE/ACM Trans. Netw.*, vol. 20, no. 2, 2012.
- [18] N. Aschenbruck, R. Ernst, E. Gerhards-Padilla, and M. Schwamborn, "BonnMotion: a mobility scenario generation and analysis tool," in *Proc. ICST SIMUTools*, 2010.
- [19] K. A. Hummel and A. Hess, "Movement activity estimation and forwarding effects for opportunistic networking based on urban mobility traces," *Wireless Communications and Mobile Computing*, vol. 13, no. 3, pp. 343–360, 2013.
- [20] E. Hyytia, P. Lassila, and J. Virtamo, "Spatial node distribution of the random waypoint mobility model with applications," *IEEE Transactions on Mobile Computing*, vol. 5, no. 6, pp. 680–694, 2006.
- [21] A. Keränen, J. Ott, and T. Kärkkäinen, "The ONE Simulator for DTN Protocol Evaluation," in *Proc. ICST SIMUTools*, 2009.
- [22] T. Spyropoulos, T. Turetli, and K. Obraczka, "Routing in Delay-Tolerant Networks Comprising Heterogeneous Node Populations," *IEEE Transactions on Mobile Computing*, vol. 8, pp. 1132–1147, 2009.
- [23] A. Vahdat and D. Becker, "Epidemic Routing for Partially-Connected Ad Hoc Networks," Duke University, Tech. Rep. CS-200006, 2000.
- [24] R. Friedman, A. Kogan, and Y. Krivolapov, "On Power and Throughput Tradeoffs of WiFi and Bluetooth in Smartphones," *IEEE Transactions on Mobile Computing*, vol. 12, no. 7, pp. 1363–1376, 2013.
- [25] O. Helgason, S. Kouyoumdjieva, and G. Karlsson, "Does mobility matter?" in *Proc. of IEEE WONS*, 2010, pp. 9–16.