

# Energy savings in opportunistic networks

Sylvia T. Kouyoumdjieva and Gunnar Karlsson  
KTH ACCESS Linnaeus Center and Electrical Engineering  
Stockholm, Sweden  
{stkou, gk}@ee.kth.se

**Abstract**—A major challenge in mobile wireless devices for opportunistic networks is to minimize the energy consumption. The minimization however should not come at a cost of reduced application throughput (i.e. goodput). This work evaluates the potential performance gains for mobile nodes that adopt a duty-cycling scheme in an opportunistic context. We present an analytical framework for evaluating the energy consumption of nodes based on a probabilistic estimation of effective contact durations, and we validate this framework on a mobility scenario. We further perform extensive trace-driven simulations and demonstrate that a duty-cycling scheme considerably improves the performance of opportunistic content distribution systems by decreasing the energy consumption without significantly affecting the goodput.

## I. INTRODUCTION

With more than 11 exabytes of monthly mobile data traffic expected by 2017 [1], wireless mobile networks without an infrastructure (such as ad-hoc networks and opportunistic networks) have been proposed as a viable solution for sharing contents directly among mobile users [2], [3] and offloading network traffic [4]. However, since mobile devices are battery-powered, the development of energy efficient mechanisms and protocols is of great importance. Although there have been efforts towards optimizing the battery lifetime on recent smartphones, the energy consumption of the LAN interface significantly reduces the battery lifetime [5].

We have confirmed this in our implementation of an opportunistic content distribution system for the Android platform, using 802.11 radio interface in ad-hoc mode [6]. Fig. 1 illustrates how the 802.11 radio interface (in ad-hoc mode) affects the battery life on a HTC Hero smartphone. Simply turning on the 802.11 interface (without any data being transmitted or received during the test) reduces the battery life to only 25% of what it is with the interface turned off. This suggests that reducing or eliminating the idle energy cost of the 802.11 interface may be a promising strategy to reduce the overall energy consumption and prolong battery life.

In our previous work [7] we explored a solution based on a dual-radio architecture for minimizing energy consumption. In this paper we look into another strategy for minimizing energy consumption on a single-radio system: *duty-cycling*. Duty-cycling, or sleep scheduling, has been a popular idea for energy savings in sensor networks [8], [9] and delay tolerant networks [10], [11]. Opportunistic networks in urban environments are affected by individual content preferences and by mobility patterns with short contact opportunities which make




Fig. 1. The 802.11 radio interface in ad-hoc mode strongly affects the battery lifetime even when no data is sent or received through it. The rightmost curve illustrates the energy consumption when the 802.11 interface is turned off. The family of curves to the left corresponds to the energy consumption in the presence of 802.11 in ad-hoc mode.

the proposed algorithms inappropriate. The authors of [12] tackle the energy-efficiency problem in opportunistic networks by proposing a pairwise wake-up scheduling algorithm which probabilistically predicts node contacts in the future. However, in urban environments with users on-the-go the assumption that two nodes meet in the future may not hold. Thus, in our evaluation we do not make any assumptions about pairwise interactions of nodes.

This work considers an opportunistic content dissemination application in which content is exchanged when a node (i.e. a pedestrian roaming in a built environment) discovers contents of interest from a peer in its communication range. Thus, the system is characterized by short contacts limited by node movements and not all contacts lead to useful data exchange. The contributions presented here are two-fold. We first develop an analytical model that captures the energy consumption of nodes participating in an opportunistic content distribution system by probabilistic evaluation of effective contact durations. We further support our findings by performing extensive trace-driven simulations incorporating realistic mobility models. Our study provides answers to the following questions:

- Is duty-cycling a feasible mechanism for energy savings in opportunistic content distribution system?
- How is the performance of a duty-cycling scheme in opportunistic networks affected by energy constraints (i.e. limited battery capacity)?
- How does the performance of a duty-cycling scheme in opportunistic networks compare to the performance of previous solutions for minimizing energy consumption in opportunistic networks?

The rest of this paper is structured as follows. Section II briefly reviews existing work and further motivates our study. Section III describes our evaluation scenario and simulation model, and section IV introduces the analytical model. Sections V and VI present the simulation results, and section VII summarizes our conclusions.

## II. RELATED WORK

Energy saving in opportunistic networks is motivated by the limited battery life of mobile devices, as well as the periodic beaconing of nodes attempting to discover peers in their vicinity. Recent studies suggest that energy consumption can be sufficiently decreased by optimizing the beaconing intervals [13] or by utilizing low duty-cycling techniques [8]. The drawback with such approaches as in [13] is that they do not decrease the overall time when the radio listens for a signal from neighboring nodes. Different mechanisms for adaptive scheduling of wake-up intervals have also been proposed in the delay tolerant networks domain [10], [11] however their aim is to reduce the energy consumption while maintaining required connectivity in sparse multihop wireless networks. Instead, opportunistic networks in urban environments are single hop networks characterized by individual content preferences and mobility patterns with short contacts.

Other studies [14], [15], [16] suggest the adoption of low-power radio for performing node discovery. Such dual-radio architectures only wake up the high-power radio when a useful contact has been found. We have also investigated this possibility in a previous work [7], and the results showed that it may prolong the battery lifetime of nodes in opportunistic network. However, since a secondary radio is often not available on personal mobile devices, the applicability is limited.

## III. EVALUATION SCENARIO

We consider energy savings in the context of opportunistic content distribution. Users in our scenario are people equipped with mobile devices walking in an urban area. The opportunistic content distribution system provides a publish/subscribe service, and each user is subscribed to feeds (e.g. topics) of interest. Each feed may hold one or more entries, which are the actual content items, such as media files. When two devices are within communication range, they exchange contents matching their local subscriptions. We envision such system to be highly utilized when users are on the go, therefore our mobility scenarios reflect this.

### A. Mobility scenarios

For capturing human mobility, we use Legion Studio [17], a commercial simulator initially developed for designing and dimensioning large-scale spaces via simulation of pedestrian behaviors. Legion Studio allows the use of open systems, where entities can enter and leave the system according to a predefined pattern. Each simulation run conducted in Legion Studio results in a mobility trace file, containing a snapshot of the positions of all nodes in the system every 0.6 s.




Fig. 2. The simulation scenarios: (a) a part of downtown Stockholm, and (b) a two-level subway station.

Fig. 2 presents the scenarios considered in our evaluation: an outdoor urban scenario, modeling the Östermalm area of central Stockholm, and an indoor two-level subway station.

The urban outdoor scenario consists of a grid of interconnected streets. The active area of the scenario is 5872 m<sup>2</sup>. Throughout their lifetime nodes are constantly moving in the area, therefore it is a high mobility scenario.

The indoor scenario is a subway station with train platforms connected via escalators to the entry-level. Nodes can arrive on foot from any of the five entry points of the subway station, or when a train arrives at the platform. The train arrivals contribute burstiness to the node arrivals and departures. The active area of the scenario is 1921 m<sup>2</sup>.

If not stated otherwise, we have chosen the input parameters of the Östermalm and the Subway scenario such that they result in approximately the same mean node density of 0.09 nodes/m<sup>2</sup> [18].

### B. Application model

The application we evaluate is an opportunistic publish/subscribe content distribution system [6] implemented in the OMNeT++ simulator [19].

Every node periodically broadcasts beacons to inform other nodes within communication range about its presence. When a device encounters another node, it checks in a local cache whether this is the first time it meets the node or, if not, whether the node has obtained new contents since their last contact. If so, the device initiates a request-response communication and at each step it tries to match the remote feeds/entries with its local subscriptions until it finally downloads the actual content items of interest.

### C. Content initialization

In our evaluation scenarios we assume that all nodes carry devices and that there are 100 available content items in the area. Items are organized in 10 feeds, each feed containing 10 items. The popularity distribution of the feeds follows a Zipf distribution with parameter  $\alpha = 0.368$ . Every device is subscribed to a single feed upon entering the observed area, and its cache is initially populated with 5 randomly chosen content items belonging to this feed. Thus, throughout its lifetime in the simulation, each node strives to obtain the rest of the content items that belong to its subscription. Entries have a mean size of 10 kB, and a standard deviation of 2 kB.

#### IV. ANALYTICAL FRAMEWORK: DUTY-CYCLING

In this section we present an analytical framework for evaluation of the mean energy consumption of duty-cycling nodes in an opportunistic network based on their effective link durations. We consider that all nodes in the system adhere to the same duty-cycling interval,  $T_c$ , however the moment when a node initially starts listening is chosen uniformly at random upon the node's creation. Furthermore, each node asynchronously chooses the listening duration in each cycle of its lifetime.

##### A. Evaluation of effective link duration

Let us consider a scenario where nodes are in communication range with a specific node, which we call an *observer*. We further carry out the analysis from the point of this observer node. Therefore, without loss of generality we assume that the observer node starts listening (turns on its radio) at time 0 of every period and then chooses a listening period from a distribution  $f_t(t)$ . Other nodes in the system choose their initial listening point uniformly at random in the interval between 0 and  $T_c$ , and choose their listening periods from the same distribution  $f_t(t)$ . For the sake of simplicity we derive results for  $t \sim \text{Uni}(0, T_c)$ , however the analytical model presented below is general enough to include any distribution  $f_t(t)$ . We define the *time until overlap* as the interval of time from the moment the observer turns its radio on until the moment another node starts listening; this definition assumes that the second node starts listening *before* the observer turns its radio off.

Let us denote with  $t$  the duration of a listening period for the observer, and with  $s$  - the duration from the beginning of the observer's listening period until the moment another node starts listening regardless of whether this happens during  $t$  or not. Since the nodes choose uniformly at random their initial listening moment, then  $s \sim \text{Uni}(0, T_c)$ . Then, we can define the *overlap duration*  $u$  as the time from the moment the second node turns its radio on until the moment the observer stops listening:

$$u = \begin{cases} t - s & s < t \\ 0 & \text{otherwise} \end{cases} \quad (1)$$

The distribution of the overlap durations will then be:

$$f_u(u) = \frac{2}{T_c} - 2\frac{u}{T_c^2} \quad \text{for } 0 \leq u \leq T_c. \quad (2)$$

Let us now denote with  $z$  the time till overlap,  $z = t - u$ . (Observe that  $z$  denotes only those nodes that start listening while the observer is on, or  $z = s$  only for  $s < t$ .) Note that since both  $u$  and  $t$  are variables in the interval  $[0, T_c]$ ,  $z \in (-T_c, T_c)$ . By definition,

$$f_z(z) = \int_{-\infty}^{+\infty} f_u(t - z) f_t(t) dt \quad (3)$$

Folding  $f_z(z)$ , for  $0 \leq z \leq T_c$  we obtain

$$f_z(z) = \frac{2}{T_c} - 2\frac{z}{T_c^2}. \quad (4)$$

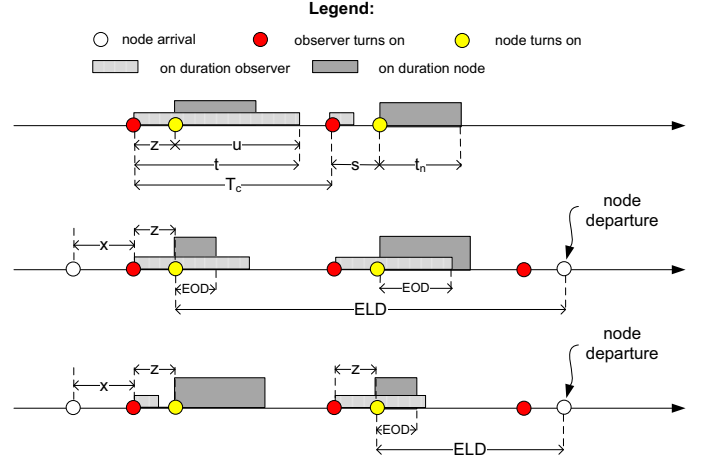


Fig. 3. Top: Visual representation of the terms time until overlap  $z$ , overlap duration  $u$  and cycling interval  $T_c$ . Middle and bottom: Illustration of effective overlap duration (EOD) and effective link duration (ELD) for the OFF-OFF case.

Since there is no synchronization between nodes in the system, the overlap duration  $u$  does not necessarily correspond to the *effective overlap duration*, i.e. the time that two nodes will be simultaneously listening. In order to evaluate the effective overlap distribution, we need to find the distribution of a random variable  $\tau = \min(u, t_n)$  where  $t_n \sim \text{Uni}(0, T_c)$  is the listening duration of the node:

$$F_\tau(\tau) = P(\min(u, t_n) \leq \tau) = F_u(\tau) + F_{t_n}(\tau) - F_{u, t_n}(\tau). \quad (5)$$

Let us now derive the pdf of the *effective link duration* when nodes are mobile. It relies on the statistical information about the physical link durations which has been thoroughly investigated in the literature both analytically and experimentally [20], [21]. When two nodes come in direct communication range, their radios can be in one of the following three states:

- **OFF-OFF**: Both the observer and the node have their radios turned off.
- **ON-OFF**: Either the observer or the node is in listening mode, and the other device has its radio turned off.
- **ON-ON**: Both the observer and the node are listening.

Below we examine these three cases and provide analytical expressions for the effective link duration for each of them. We assume that every contact opportunity leads to content exchange, so in all cases we consider the effective link duration to be the time from the moment an overlap occurs until the end of the physical contact.

1) *The OFF-OFF case*: Fig. 3 illustrates the case when both the observer and the node have their radios turned off upon establishing a physical contact. Let's denote by  $x$  the random variable of time from the moment a node enters the observer's communication range until the next moment when the observer starts listening. We can represent  $x$  as a stationary process, and according to Palms theorem we know that whenever the process is uniformly distributed, the time  $x$  will also be uniformly distributed in  $[0, T_c]$ . If  $D$  is the physical

contact duration and it has distribution  $f_D(D)$ , then we can obtain the distribution  $f_L(L) = f_D(D - x)$  where  $L = D - x$  is the duration from the moment the observer starts listening until the end of the physical contact. From the moment the observer starts listening we already know the distribution of time until overlap, which is given by  $f_z(z)$  in Eq. (3). We should note that it is possible that  $L > T_c$ , thus we can divide  $L$  in  $j$  intervals of length  $T_c$  with respect to the cycling of the observer node,  $j = \lfloor \frac{L}{T_c} \rfloor$ . If we denote with  $\tau_e$  the effective link duration, we can then express  $z = L - jT_c - \tau_e$ . Then,

$$f_{\tau_e|L}(\tau_e|L) = \sum_{j=0}^{\lfloor \frac{L}{T_c} \rfloor} q_j^j p_j^{j+1} f_z(L - jT_c - \tau_e), \quad (6)$$

where  $q_j = F_t(L - jT_c - \tau_e)$  is the probability that there was no overlapping between the observer and the node in the  $j$ th cycle, and  $p_j = F_t(T_c - (L - jT_c - \tau_e))$  is the probability that there was no overlapping between the node and the next occurrence of the beginning of the listening period of the observer, and  $F_t(\cdot)$  denotes the distribution of listening durations.

Integrating over all possible  $L$ , gives us:

$$f_{\tau_e}(\tau_e) = \int_0^\infty f_{\tau_e|L}(\tau_e|L) f_L(L) dL \quad (7)$$

2) *The ON-OFF case:* The ON-OFF case refers to a situation where a node that enters the observer's communication range already has its radio turned on, and the observer turns on its radio before the node goes to sleep.

$$f_{\tau_e|D}(\tau_e|D) = \int_0^\infty \int_{D-L}^{T_c} f_z(z) dz f_L(L) dL \quad (8)$$

Then,

$$f_{\tau_e}(\tau_e) = \int_0^\infty f_{\tau_e|D}(\tau_e|D) f_D(D) dD \quad (9)$$

3) *The ON-ON case:* The ON-ON case refers to a situation where both the observer and the node entering the observer's communication range already have their radios turned on upon physical contact establishment. Thus, the nodes can utilize the whole physical contact for content exchange. The probability that nodes can utilize the whole physical contact,  $f_x(x)$ , can be expressed as:

$$f_x(x) = \int_0^{T_c-x} f_z(z) (1 - F_\tau(T_c - z - x)) dz \quad (10)$$

where  $x$  is the time interval from the beginning of the physical contact until the next listening period of the observer, and  $F_\tau(T_c - z - x)$  is the distribution of the effective overlap duration in (5).

Thus, the distribution of effective link durations that span over the whole physical contact duration is:

$$f_{\tau_e}(\tau_e|x) = \int_0^\infty f_x(x) f_D(D) dD. \quad (11)$$

The distribution of effective link durations is then given by Eq. (7), (9) and (11).

### B. Evaluation of energy consumption

When turned on the radio interface is typically in one of the following states: *idle* (listening for an incoming signal), *transmit* and *receive*. We here present a simplified model which takes into account the physical link durations and the effective overlap durations of an observer and estimates the mean energy consumption for it. Fig. 4 illustrates the general idea. The time  $t_{busy}$  corresponds to the effective overlap duration, which in turn translates into the energy consumed by a node when it either transmits or receives data. We choose to use the effective overlap duration instead of the effective link duration, since the former represents the minimum guaranteed time when two nodes can exchange data.  $t_{idle}$  is the time when the radio is in idle state, and the node is scanning the area for possible peers to associate with. In the case of duty-cycling,  $t_{idle}$  is the time during which a node has its interface turned on, but there are no nodes of interest in its range. In the case of a single radio which is always turned on,  $t_{idle}$  corresponds to the time during which a node is not communicating with a peer. Thus, we can represent the energy consumed by a node in the system as:

$$E = e_1 \sum t_{busy} + e_2 \sum t_{idle} \quad (12)$$

where  $e_1$  denotes the energy spent for sending/receiving information and  $e_2$  is the energy consumed by listening.

The distribution of the *total effective overlap duration*  $x = \sum t_{busy}$  can be expressed as a sum over all possible combinations of effective overlap durations  $\tau_k$  that an observer can experience throughout its lifetime:

$$f_x(x) = \sum_{k=1}^N \sum_{\forall (\tau_1 \dots \tau_k)} \prod_{i=1}^k f_\tau(\tau_i), \quad (13)$$

$\sum \tau_k = x$

where  $N$  is the number of nodes the observer meets throughout its lifetime.

Since the intervals of time when a node is listening follow a uniform distribution, the distribution of the *total listening duration*  $y = \sum_{k=0}^n t_{idle}$  where  $n = \lfloor \frac{L}{T_c} \rfloor$  for the duty-cycling case can be obtained by the Irwin-Hall's uniform sum distribution [22].

$$f_y(y, n) = \frac{1}{2(n-1)!} \sum_{k=0}^n (-1)^k \binom{n}{k} \left( \frac{y}{T_c} - k \right)^{n-1} \text{sgn} \left( \frac{y}{T_c} - k \right) \quad (14)$$

Thus, we can present the mean energy consumption of a duty-cycling node as:

$$\bar{E}_{dc} = e_1 E[f_x(x)] + e_2 E[f_y(y, n)] \quad (15)$$

and the mean consumption of a node without duty-cycling as:

$$\bar{E}_{on} = e_1 E[f_x(x)] + e_2 (T - E[f_x(x)]) \quad (16)$$

where  $T$  is the lifetime of the node in the system.

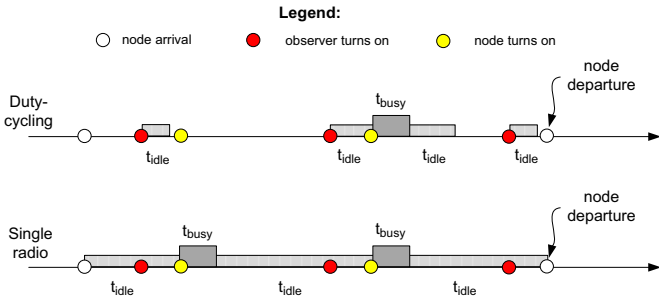


Fig. 4. Mapping time to energy consumption for an off-the-shelf single radio system, and a single radio with duty-cycling.




Fig. 5. Street scenario: (a) cdf of effective link durations and (b) cdf of effective overlap duration.

## V. SIMULATION RESULTS: STREET SCENARIO

In this section we present simulation results to validate the proposed analytical model. We introduce a simple mobility scenario which we further refer to as ‘street scenario’. As the name suggests, it represents a street segment. Nodes enter into the street from both ends, and traverse the distance at a constant speed  $v = 0.9$  m/s. Each node acts as an observer while traversing the street. We here present results for the analytic model and a simulation of a street with length  $L = 100$  m, communication range  $r = 10$  m and symmetric Poisson arrival rates  $\lambda = 0.01$  nodes/s at both ends of the street. We note that we only consider effective link and overlap durations for nodes moving in opposite directions since these contacts bring fresh contents in the system and assist the content dissemination process. At a granularity of 0.6 s (to correspond to the granularity provided by Legion Studio) each node that has its radio turned on records information about the overlap link durations with other listening peers in range and the potential effective link durations with the same peers, as long as any are available. Fig. 5 illustrates the results for the effective link durations and the effective overlap durations. The results in Fig. 5(a) are compared to the overall effective link durations distribution in Eq. (7), (9), and (11), while the results in Fig. 5(b) are compared to the effective overlap durations distribution in Eq. (5).

Fig. 6 presents the distributions of the total listening durations and the total effective overlap durations. The total listening durations distribution, in Fig. 6(a), is compared to the CDF of the Irwin-Hall distribution in Eq. (14), while the total effective link overlap durations distribution, in Fig. 6(b),

corresponds to Eq. (13).

In order to evaluate the energy consumption for the street scenario, we construct the following simple content exchange scenario. We assume that nodes entering at the left-hand side of the street are content seekers, while nodes entering at the right-hand side are content providers. There is only one feed to which all nodes are subscribed, and this feed holds infinitely many unique items. Furthermore, content seekers request items as long as a content provider is in range and is listening. In this way we assure that every node uses its full effective overlap duration for content exchange: content seekers constantly request contents, and content providers constantly send out data. Fig. 7 presents the results. The current drawn in each radio state is taken from the simulation parameters and plugged into Eq. (15) and (16) from the analytical framework. We see that the analytical results slightly underestimate the actual performance of the system. The reason is that the energy framework [23] of the OMNeT++ simulator implements a detailed profile of the energy consumption, and captures even the energy spent for decoding incoming frames. Such level of detail is omitted in the analytical framework, thus the framework presents a lower bound of the mean energy consumption.

## VI. SIMULATION RESULTS: URBAN SCENARIOS

We here investigate the performance of more complex mobility scenarios: the Östermalm grid of streets and the Subway station. We explore the behavior of a duty-cycling scheme with a cycling interval of  $T_c = 10$  s when nodes have infinite and finite battery capacity, and we contrast this behavior to other proposed solutions for decreasing energy consumption in opportunistic networks.

### A. Performance metrics & configurations

We focus on two performance metrics: *energy consumption* and *goodput* (i.e. application throughput) from a system perspective. Since we study an open system, it is important that the metrics are normalized with respect to the nodes’ sojourn time in the simulation. The system goodput is simply the sum of the number of bytes downloaded  $B_i$  by each node divided by the sum of the lifetimes of nodes in the simulation  $t_i$ , or  $G = \sum B_i / \sum t_i$ . We only count bytes of fully downloaded content items, so the goodput is a measure of the system usefulness for the users, i.e. how much content it can provide. The energy consumption of the system is defined as the sum of the energy consumed by the radio interface(s) normalized with respect to the sum of nodes’ sojourn times. Thus, the energy consumption can be expressed as  $E = \sum E_i / \sum t_i$ .

We compare the performance of the following configurations with respect to the above performance metrics:

- 1) **Single radio (SR-ON)** — an off-the-shelf solution as it is currently implemented in today’s smartphones; a single radio interface is used for both service and content discovery as well as for actual content exchange.
- 2) **Single radio with duty-cycling (SR-DC)** — a single-radio system as proposed herein.






Fig. 6. Street scenario: (a) CDF of total listening durations and (b) CDF of total effective overlap duration.

- 3) **Dual radio (DR)** — a dual-radio system where a primary high-power, high-bitrate radio (IEEE 802.11) is used for data transfer and a secondary low-power, low-bitrate radio (IEEE 802.15.4) is used as a control channel for performing neighbor and service discovery. The high-power radio is completely suspended when not in use; it is associated with a modest (2 s) delay for turning on, but does not consume energy when suspended. A detailed description of the dual-radio can be found in [7].
- 4) **Benchmark (BM)** — an idealized system in which global knowledge is assumed for the location and the subscriptions of each node in the system; such a system concentrates on evaluation of the energy consumed for the actual data exchange and abstracts the discovery phase, thus it provides a lower bound for the energy consumption and an upper bound for the goodput.

### B. Infinite battery

Let us now examine the performance of the system as a whole. We first consider a case where the battery of each of the nodes in the area is enough to support exchange of data between nodes throughout their lifetime in the system. Fig. 8 and 9 present results for the Östermalm and the Subway scenario when the communication range is  $r = 10$  m and  $r = 50$  m, respectively. Although the single-radio system with duty-cycling does not decrease the energy consumption as much as the dual-radio system, at 50 m communication range (Fig. 9) it outperforms the dual-radio system in terms of goodput for both scenarios. This is more visible in the Subway scenario: although the energy consumed by the dual-radio architecture is only 20% of that of the single-radio SR-ON, the goodput has dropped with some 40%. The data in Table I reveals that the drop in goodput is actually due to the fact that many of the nodes in the Subway scenario do not obtain any content throughout their lifetime when operating with a dual-radio. This is due to the characteristics of the low-power 802.15.4 radio which is unable to perform the node and service discovery in areas with high density. (For more explanation see [7].) However, in Fig. 9(b) we see that with the single radio system with duty-cycling we achieve the same goodput as the single-radio system, while we only consume 40% of the energy of the single radio. Thus, in the single-radio system we have at least 60% of energy that comes to waste.




Fig. 7. Street scenario: mean energy consumption for a duty-cycling radio (SR-DC) and a single radio which is always listening (SR-ON).

When the communication range is reduced to only  $r = 10$  m, we still see the same trend in gains in goodput for the Subway scenario, Fig. 8(b). However, we also see that in the Östermalm scenario the single-radio system with duty-cycling performs worse than both the dual-radio and the single-radio configurations when it comes to goodput. This is due to the differences in the structure of the Östermalm and Subway scenarios [18], thus when nodes duty-cycle they sometimes miss useful contacts. We should also note that the number of nodes that do not obtain contents in this scenario does not differ much among the configurations (see Table I).

Scenario	BM	SR-ON	SR-DC	DR
Östermalm, $r = 10$ m	2.3%	3.6%	4.7%	3.8%
Östermalm, $r = 50$ m	0.3%	0.6%	0.5%	1.2%
Subway, $r = 10$ m	0.5%	0.6%	0.6%	3.4%
Subway, $r = 50$ m	0.0%	0.2%	0.1%	22.7%

TABLE I  
PERCENTAGE OF NODES THAT DO NOT OBTAIN A SINGLE CONTENT ITEM DURING THEIR LIFETIME.

### C. Finite battery

In this section we investigate the effects of duty-cycling on nodes with an energy budget. We consider the Östermalm scenario, and we vary the arrival rate at each of the entry points from  $\lambda = 0.01..0.9$  nodes/s. Upon arrival in the system each node picks an energy budget that it can spend throughout its lifetime in the observed area. Since the lifetime of nodes in the area is rather short (i.e. on average 300 s), we scale the battery capacity so that it corresponds to the lifetime. Currently batteries of mobile devices have an average capacity of  $C = 2000$  mAh. If a battery should last some  $T = 16$  hours between charges, then we are left with  $C_h = \frac{C}{T} = 125$  mAh. Since our nodes live on average 5 min, then a capacity of approximately 10mAh would be sufficient for most nodes to support their radios constantly turned on while in the system. However we would like to challenge the system, and allow some nodes to drain their batteries before leaving the observed area. Therefore, uniformly at random each node is allowed to choose an energy budget in the range of (1, 5) mAh.

The first question we aim to answer is how duty-cycling affects the system performance upon scaling the system? Fig. 10 presents the results for the Östermalm scenario with  $r = 10$  m communication range across different arrival rates.




Fig. 8. Normalized energy consumption and normalized goodput for the Östermalm scenario (a) and the Subway scenario (b) with a 10 m communication range.




Fig. 10. Effects of energy budgeting: (a) energy consumption and (b) goodput for the Östermalm scenario with  $r = 10$  m and different arrival rates.

Each point on the graphs corresponds to the performance metrics collected from a single trace. Across all densities (i.e. different arrival rates) the energy consumption of the duty-cycling scheme is lower than the energy consumption of the SR-ON scheme due to the decreased amounts of time when the radio interface is simply listening. Moreover, the duty-cycling scheme virtually extends the lifetime of nodes with energy budgeting thus allowing for more contact opportunities and yielding as much as twice the goodput of a SR-ON scheme even in sparsely populated areas. As the arrival rates increase, nodes get more chances of meeting peers that could provide them with contents of interest. Thus, the amounts of data obtained by both schemes become comparable, however the duty-cycling scheme still provides better performance at a lower energy cost.

The second question we address is how does duty-cycling on nodes with an energy budget perform in comparison to our benchmark with infinite battery and global knowledge for all useful content exchanges? We choose to investigate this for three different values of the arrival rate, i.e.  $\lambda = 0.01, 0.15, 0.9$ , to examine what happens when we have very sparse to very dense configurations. One could possibly relate these to an early morning, a lunch-time and a rush-hour configuration, respectively. The results are presented in Fig. 12 and Table II.

In sparsely populated environments, in Fig. 12(a), the restrictions on the battery capacity significantly reduce the goodput of the system, with the reduction being even more notable when the nodes have their radios constantly turned on. We note that the duty-cycling scheme does improve both the energy consumption and the goodput of the system, however




Fig. 9. Normalized energy consumption and normalized goodput for the Östermalm scenario (a) and the Subway scenario (b) with a 50 m communication range.




Fig. 11. Effects of energy budgeting: normalized energy consumption and normalized goodput for the Subway scenario with  $r = 10$  m.

we also note that the performance is still far below the performance of the benchmark scenario. Interestingly, even when there are no restrictions on the battery capacity (SR-ON, Infinite Battery), the goodput still cannot match the goodput provided by the benchmark configuration. The reason for this poor performance is the low density in the observed area. Table II shows that when  $\lambda = 0.01$  nodes/s, even the benchmark cannot satisfy the needs of more than 50% of the nodes in the system.

As the density of nodes increases, the performance of the duty-cycling scheme in terms of goodput approaches that of the single radio scheme with infinite battery capacity, while it still keeps the energy consumption approximately 60%-70% lower than that of the single radio scheme. These results demonstrate that duty-cycling not only virtually extends the battery lifetime of nodes, but it also increases the opportunities of nodes to obtain more content items of interest throughout their stay in the system.

Arrival rate	SR-ON (IB)	SR-ON (FB)	SR-DC (FB)	BM (IB)
$\lambda = 0.01$ n/s	57.9%	92.9%	84.0%	50.5%
$\lambda = 0.15$ n/s	3.6%	41.4%	15.9%	2.3%
$\lambda = 0.90$ n/s	1.4%	5.3%	2.2%	0.3%

TABLE II  
PERCENTAGE OF NODES THAT DO NOT OBTAIN CONTENT IN THE ÖSTERMALM SCENARIO WITH  $r = 10$  m AND DIFFERENT ARRIVAL RATES.

We also examine the effects of energy budgeting for the Subway scenario with  $r = 10$  m. The results are presented in Fig. 11. When an energy budget is introduced for the nodes, the duty-cycling scheme consumes around 60% of the energy consumed by a single radio, without compromising

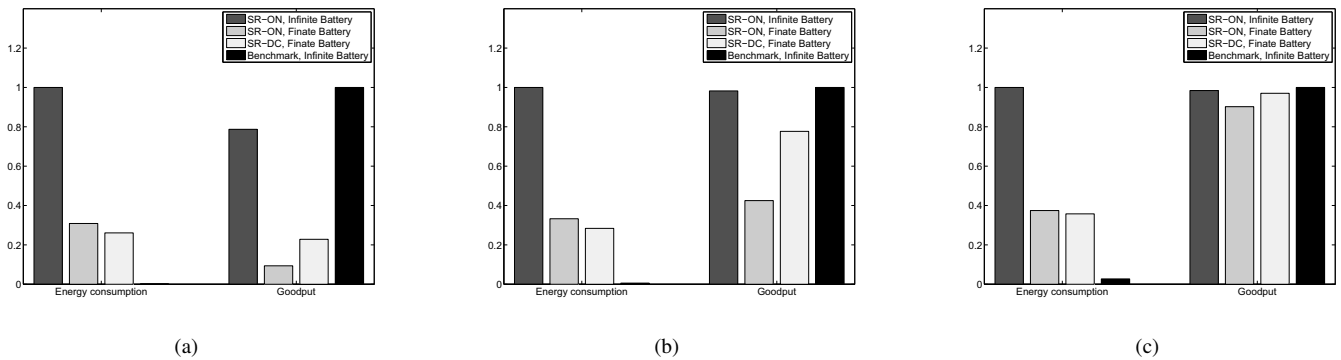


Fig. 12. Effects of energy budgeting: normalized energy consumption and normalized goodput for the Östermalm scenario with  $r = 10$  m and different arrival rates: (a)  $\lambda = 0.01$  nodes/s, (b)  $\lambda = 0.15$  nodes/s, and (c)  $\lambda = 0.90$  nodes/s.

the goodput performance. We note that although the Subway scenario has the same average node density as the Östermalm scenario with  $\lambda = 0.15$  nodes/s, here the performance of the duty-cycling scheme in terms of goodput is better than the performance of the same scheme in the Östermalm scenario, Fig. 12(b). Due to the differences in the structure of the two scenarios, nodes in the Subway scenario have more opportunities of meeting peers that can provide them with contents of interest.

## VII. CONCLUSION

In this paper we investigated the energy savings in mobile nodes utilizing a duty-cycling scheme in an opportunistic content dissemination system. We proposed an analytical framework for probabilistic evaluation of effective contact durations among duty-cycling nodes in opportunistic network and we further derived expressions for evaluating the mean energy consumption of nodes in the system. We verified the framework on realistic pedestrian mobility traces representing a simple street segment scenario. Furthermore, we used complex mobility scenarios such as a grid of streets and a subway station to evaluate the performance gains in terms of both energy consumption and application throughput (i.e. goodput) of a duty-cycling scheme in opportunistic context. We compared the performance of a duty-cycling scheme to previously proposed solutions for reducing energy consumption in opportunistic networks, and showed that duty-cycling decreases the energy consumption without significantly reducing the goodput. We also investigated the performance of a duty-cycling scheme when nodes have infinite versus finite battery life. We showed that when battery capacity is finite, duty-cycling not only reduces the energy consumption, but also virtually extends nodes' lifetime in the system allowing nodes to obtain more contents, thus increasing the goodput.

## REFERENCES

- [1] "VNI Forecast Highlights," [http://www.cisco.com/web/solutions/sp/vni/vni\\_forecast\\_highlights/index.html](http://www.cisco.com/web/solutions/sp/vni/vni_forecast_highlights/index.html).
- [2] C. Boldrini, M. Conti, and A. Passarella, "Modelling data dissemination in opportunistic networks," in *Proc. CHANTS*, 2008.
- [3] M. Papadopoulou and H. Schulzrinne, "Effects of power conservation, wireless coverage and cooperation on data dissemination among mobile devices," in *Proc. ACM MobiHoc*, Long Beach, CA, USA, 2001.
- [4] B. Han, P. Hui, V. Kumar, M. Marathe, J. Shao, and A. Srinivasan, "Mobile data offloading through opportunistic communications and social participation," *Mobile Computing, IEEE Transactions on*, vol. 11, no. 5, pp. 821–834, May 2012.
- [5] M. Stemm and R. H. Katz, "Measuring and reducing energy consumption of network interfaces in hand-held devices," *IEICE Transactions on Communications, special Issue on Mobile Computing*, 1997.
- [6] O. Helgason, E. A. Yavuz, S. Kouyoumdjieva, L. Pajevic, and G. Karlsson, "A mobile peer-to-peer system for opportunistic content-centric networking," in *Proc. ACM SIGCOMM MobiHeld workshop*, 2010.
- [7] S. T. Kouyoumdjieva, O. Helgason, E. A. Yavuz, and G. Karlsson, "Evaluating an energy-efficient radio architecture for opportunistic communication," in *Proc. of ICC E2Nets workshop*, 2012.
- [8] S. Guo, Y. Gu, B. Jiang, and T. He, "Opportunistic flooding in low-duty-cycle wireless sensor networks with unreliable links," in *Proc. MobiCom*, 2009.
- [9] S. Bader and B. Oelmann, "Adaptive synchronization for duty-cycling in environmental wireless sensor networks," in *Proc. ISSNIP*, 2009.
- [10] B. J. Choi and X. Shen, "Adaptive asynchronous sleep scheduling protocols for delay tolerant networks," *Mobile Computing, IEEE Transactions on*, vol. 10, no. 9, pp. 1283–1296, 2011.
- [11] Y. Xi, M. Chuah, and K. Chang, "Performance evaluation of a power management scheme for disruption tolerant network," *Mob. Netw. Appl.*, vol. 12, no. 5, Dec. 2007.
- [12] W. Gao and Q. Li, "Wakeup scheduling for energy-efficient communication in opportunistic mobile networks," in *Proc. of INFOCOM*, 2013.
- [13] W. Wang, V. Srinivasan, and M. Motani, "Adaptive contact probing mechanisms for delay tolerant applications," in *Proc. MobiCom*, 2007.
- [14] E. Shih, P. Bahl, and M. J. Sinclair, "Wake on wireless: an event driven energy saving strategy for battery operated devices," in *Proc. ACM MobiCom*, 2002.
- [15] T. Pering, Y. Agarwal, R. Gupta, and R. Want, "Coolspots: reducing the power consumption of wireless mobile devices with multiple radio interfaces," in *Proc. MobiSys*, 2006.
- [16] H. Jun, M. H. Ammar, M. D. Corner, and E. W. Zegura, "Hierarchical power management in disruption tolerant networks using traffic-aware optimization," *Comput. Commun.*, vol. 32, no. 16, pp. 1710–1723, 2009.
- [17] "Legion Studio," <http://www.legion.com/>.
- [18] O. Helgason, S. T. Kouyoumdjieva, and G. Karlsson, "Opportunistic communication and human mobility," *Mobile Computing, IEEE Transactions on*, 2013, to appear.
- [19] O. R. Helgason and K. V. Jónsson, "Opportunistic networking in OMNeT++," in *Proc. SIMUTools, OMNeT++ workshop*, 2008.
- [20] A. Nayebi, A. Khosravi, and H. Sarbazi-Azad, "On the link excess life in mobile wireless networks," in *Proc. ICCTA*, 2007.
- [21] F. Bai, N. Sadagopan, B. Krishnamachari, and A. Helmy, "Modeling path duration distributions in manets and their impact on reactive routing protocols," *Selected Areas in Communications, IEEE Journal on*, 2004.
- [22] N. L. Johnson, S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions*. Wiley, 1995.
- [23] L. M. Feeney and D. Willkomm, "Energy framework: an extensible framework for simulating battery consumption in wireless networks," in *Proc. SIMUTools*, 2010.