Energy-Aware Opportunistic Mobile Data Offloading Under Full and Limited Cooperation

Sylvia T. Kouyoumdjieva and Gunnar Karlsson
ACCESS Linnaeus Center, School of Electrical Engineering
KTH Royal Institute of Technology, Stockholm, Sweden
Email: {stkou, gk}@ee.kth.se

***Abstract*—Opportunistic networking (a.k.a. device-to-device communication) is considered a feasible means for offloading mobile data traffic. Since mobile nodes are battery-powered, opportunistic networks must be expected to satisfy the user demand without greatly affecting battery lifetime. To address this requirement, this work introduces *progressive selfishness*, an adaptive and scalable energy-aware algorithm for opportunistic networks used in the context of mobile data offloading. The paper evaluates the performance of progressive selfishness in terms of *both* application throughput and energy consumption via extensive trace-driven simulations of realistic pedestrian behavior. The evaluation considers two modes of nodal cooperation: full and limited, with respect to the percentage of nodes in the system that adopt progressive selfishness. The paper demonstrates that under full cooperation the proposed algorithm is robust against the distributions of node density and initial content availability. The results show that in certain scenarios progressive selfishness achieves up to 85% energy savings during opportunistic downloads while sacrificing less than 1% in application throughput. Furthermore, the study demonstrates that in terms of total energy consumption (by both cellular and opportunistic downloads) in dense environments the performance of progressive selfishness is comparable to downloading contents directly from a mobile network. Finally, the paper shows that progressive selfishness is robust against the presence of non-cooperative nodes in the system, and that in certain scenarios the system-level performance does not deteriorate significantly under limited cooperation even when 50% of the nodes in the system do not adhere to the specifics of the algorithm.**

***Index Terms*—mobile data offloading, selfishness, duty cycling, energy savings, opportunistic networking.**

I. Introduction

Due to the proliferation of mobile devices in recent years, mobile network operators are now expected to satisfy immense data traffic demands via the cellular network (i.e. more than 24 exabytes of monthly mobile data traffic by 2019 as predicted by Cisco [1]). Thus, mobile data offloading has been suggested as a complement. One promising approach for offloading network traffic is based on mobile opportunistic networks which allow contents to be shared directly among mobile devices when in proximity.

Previous studies on mobile data offloading via opportunistic communication [2], [3], [4] aim at maximizing the data delivered to mobile devices. However, they do not consider the limited battery capacity of the mobile device. The energy consumption is strongly affected by the 802.11 radio interface when turned on in ad-hoc mode [5].

Recent studies propose duty cycling as a viable strategy for decreasing energy consumption in opportunistic networks [6], [7]. Such solutions however assume that all nodes participating in the content exchange are altruistic and willing to share data with others throughout their lifetime. This may be an overly optimistic assumption, and nodes that have already obtained all contents of interest may prefer to opt out of distribution in order to save energy. By default, selfishness has always been considered harmful to the performance of opportunistic networks and different mechanisms have been suggested for providing incentives for nodes to behave altruistic [8], as well as for detection and avoidance of selfish nodes throughout the routing process [9]. However, they do not take into account the price in terms of energy consumption that a node pays for being altruistic.

In this work we aim at decreasing the energy consumption while *at the same time* retaining the application throughput of an opportunistic network. The work is motivated by our previous study [10] in which we argue that opportunistic mobile data offloading should take into account not only the requirements of mobile operators, but also those of mobile devices. The main contributions are:

- We revisit the concept of selfishness in opportunistic networks and demonstrate that in contrast to prior understanding, selfishness can decrease energy consumption while satisfying user demands.
- We propose an adaptive and scalable energy-aware algorithm for opportunistic networks, *progressive selfishness*, which combines the merits of two energy saving mechanisms that have always been considered mutually exclusive: duty cycling and selfishness.
- We perform extensive trace-driven simulation analysis using realistic pedestrian mobility and evaluate the performance of progressive selfishness under two modes of nodal cooperation: full and limited, with respect to the percentage of nodes in the system that adopt the algorithm.
- We show that under full cooperation progressive selfishness is robust against parameters such as node density and initial content availability. The results demonstrate that progressive selfishness can achieve up to 85% energy savings while losing as little as 1% in application throughput.
- We show that under limited cooperation progressive self-

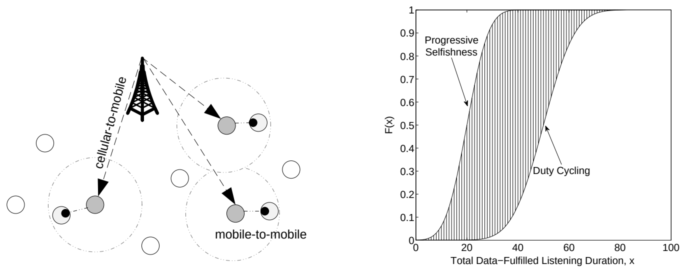

(a) (b)

Fig. 1. (a) Mobile data offloading: only a subset of nodes in the observed area download contents directly from the cellular network; the rest of the nodes obtain contents opportunistically. (b) Cumulative density function of the lower and upper bounds of the total listening duration achieved by progressive selfishness.

ishness is robust against the presence of non-cooperative nodes in the system. We also show that in sparsely populated areas deviating from the progressive selfishness algorithm deteriorates the performance of cooperative and non-cooperative nodes alike.

The rest of this paper is structured as follows. Section II revisits popular energy saving mechanisms, and introduces progressive selfishness. Section III outlines the evaluation scenario, and Sections IV and V present results from realistic pedestrian mobility scenarios under full and limited cooperation, respectively. In the context of our findings, Section VI discusses previous work. Finally, Section VII concludes the study.

II. Progressive Selfishness

We assume that users are pedestrians equipped with mobile devices moving in an urban area, e.g. a grid of streets in a city, and that a mobile operator wishes to disseminate data to all users in the observed area, software updates or special offers, for instance. Instead of downloading data separately to every user, the mobile operator relies on data offloading as illustrated in Fig. 1(a). Upon entering the area, some users are instructed to download parts of the data directly from the cellular network and store them in their caches. The rest of the users attempt to download content items opportunistically when in communication range with a node that already has the data. In this work we do not consider how the mobile operator determines the best candidate nodes to initially carry contents. Instead we focus on the performance of the opportunistic data dissemination.

Throughout its lifetime a node is either data-seeking, or data-fulfilled. A node is *data-seeking* if it is missing one or more content items of those provided in the observed area. A node is *data-fulfilled* if it already has downloaded all content items.

In the context of mobile data offloading, the main objective of a data-seeking node is to obtain as many content items of those provided by a mobile operator as possible. However, a data-seeking node needs to discover data at a low energy cost. In order to save energy while searching for contents, we allow data-seeking nodes to duty cycle (DC) within a cycling interval T_c, i.e. to iteratively turn their radio interfaces on and off. A node can only discover other peers and exchange data with them while its radio interface is turned on. In our previous work [6] we suggest that the time during which a radio interface is turned on should be chosen uniformly at random $d \sim \mathcal{U}(0, T_c)$ in the beginning of every cycling interval, and the radio interface should be consecutively turned off for the remaining $(T_c - d)$ time units. We adopt this strategy, since it was shown to decrease energy consumption roughly by half without incurring significant application throughput losses for an opportunistic content distribution system.

Algorithm 1 Progressive Selfishness (DC-PS)

```
Tc ← initial cycling interval
w ← 1, initial inactivity window
t ← 0, initial time
N ← set of all content items
M ⊆ N, set of items to obtain
while node has battery capacity do
    if time to turn radio interface on then
        UNIFORMON()
    end if
    if time to turn radio interface off then
        if M ≡ ∅ then
            PROGRESSIVEOFF()
        else
            UNIFORMOFF()
        end if
    end if
end while
procedure UNIFORMON()
    d ← U(0, Tc)
    t ← t + d
end procedure
procedure UNIFORMOFF()
    t ← t + (Tc − d)
end procedure
procedure PROGRESSIVEOFF()
    if node has shared one or more items while on then
        w ← 1
    end if
    t ← t + (wTc − d)
    w ← 2w
end procedure
```

Once a node obtains all content items of interest, its objective changes and the focus shifts from downloading data to saving energy. Ultimately, a node would save greatest amount of energy if it chooses to opt out of the data dissemination process at the moment it becomes data-fulfilled. We refer to such behavior as *strict selfishness*. However, if all data-fulfilled nodes choose to be strictly selfish, data-seeking nodes may be brought to starvation due to the lack of active content providers, and be forced to eventually download the contents directly from the cellular network.

Thus, we propose data-fulfilled nodes to behave in a *progressively selfish* manner. Just as with regular duty cycling, a progressively selfish node iterates between a state in which its radio interface is turned on, and a state in which the radio interface is turned off. While the duration of the on period is again chosen uniformly at random $d \sim \mathcal{U}(0, T_c)$, the duration of the off period depends on the demand for contents from peers in the vicinity. We define an *inactivity window*, w, as a parameter that increases exponentially every time a data-fulfilled node does *not* deliver contents to neighboring peers while its radio is turned on. If there is no demand for the data carried by a data-fulfilled node, it progressively increases the off duration by $(wT_c - d)$, thus saving energy. However, the inactivity window w shrinks to 1 if a data-fulfilled node provides data to a peer. Since nodes are mobile, throughout their lifetime they may traverse areas with different densities of data-seeking nodes. Downloading contents to a data-seeking peer may infer that a data-fulfilled node has entered an area where data needs to be disseminated. The inactivity window w can shrink to 1 also if a node is instructed to do so by the mobile operator, or if the energy spent for opportunistic communication surpasses a predefined threshold. In this work however we only allow the inactivity window to shrink when a node downloads contents to a data-seeking peer. The details of progressive selfishness are presented in Algorithm 1.

In order to evaluate the potential energy savings achieved by progressive selfishness, we examine the distribution of the *total data-fulfilled listening duration*, x, i.e. the total amount of time during which a node keeps its radio interface turned on *after* it has become data-fulfilled. (We note that total data-fulfilled listening duration and total energy consumed are proportional.) Since the listening duration in each cycling interval is chosen uniformly at random $d \sim \mathcal{U}(0, T_c)$, the probability distribution of the total data-fulfilled listening duration follows the Irwin-Hall's uniform sum distribution [11].

$$f_x(x,n) = \frac{1}{2(n-1)!} \sum_{k=0}^{n} (-1)^k \binom{n}{k} \left(\frac{x}{T_c} - k\right)^{n-1} \operatorname{sgn}\left(\frac{x}{T_c} - k\right) \tag{1}$$

where $x = \sum_{i=1}^{n} d_i$ and d_i is the listening duration in the ith cycling interval, and n is the total number of cycling intervals during the data-fulfilled phase. Note that the amount of energy savings via progressive selfishness for a node depends on the number of times the radio interface has been switched on, n, which in turn is defined by parameters of the environment such as node density, content availability and lifetime. Thus, evaluating the overall performance of the algorithm is nontrivial, and requires knowledge of the system state at any moment in time. However, evaluating the lower and upper bound of the total data-fulfilled listening duration achieved by progressive selfishness is possible. For a node with lifetime L the upper bound is reached if the node provides data to other peers in each cycling interval, and $n = \lceil \frac{L}{T_c} \rceil$; in this case progressive selfishness is reduced to a simple duty-cycling scheme as described in [6]. The lower bound is reached if a data-fulfilled node is never requested to provide any content items to neighboring nodes while listening, and the inactivity window w continuously increases throughout the node's lifetime. Fig. 1(b) illustrates the distribution of the lower and upper bound of total listening duration of a node with a lifetime of 100 time units, and a cycling interval $T_c = 10$ time units.

III. Evaluation Scenario

Here we present a sample publish/subscribe service provided by an opportunistic content distribution system which allows nodes to discover contents in other peers [5].

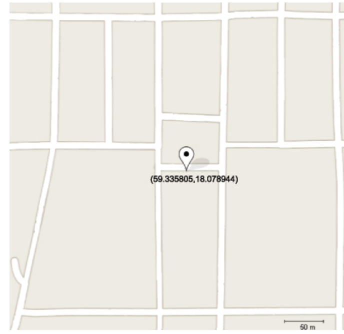

Fig. 2. Urban mobility scenario: a grid of streets representing a part of downtown Stockholm, Östermalm.

We assume that nodes are subscribed to feeds (e.g. topics) of interest through a publish/subscribe service provided by an opportunistic content distribution system. Each feed may hold one or more entries, and each entry can constitute of one or more chunks (which we call *content items* throughout this paper). Feeds, entries and chunks are announced by the mobile operator to every node entering the observed area. Whenever the radio interface of a node is turned on, it periodically broadcasts beacons to inform potential neighbors about its presence. (In the current study the beaconing interval is set to 0.5 s.) Upon an encounter with another node, the device has to determine 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 content items of interest.

We use an implementation of an opportunistic content distribution system in the OMNeT++ simulator [12] and we rely on the energy framework [13] for modeling battery consumption.

A. Mobility scenario

In order to realistically recreate pedestrian mobility, we use the Walkers traces [14] captured in Legion Studio [15], a commercial simulator initially developed for designing and dimensioning large-scale spaces via simulation of pedestrian behaviors. Its multi-agent pedestrian model is based on advanced analytical and empirical models which have been calibrated by measurement studies. Each simulation run results in a trace file, containing a snapshot of the positions of all nodes in the system every 0.6 s.

Fig. 2 presents the outdoor urban scenario considered in our evaluation. The Östermalm scenario consists of a grid of interconnected streets. Fourteen passages connect the observed area to the outside world. The active area is 5872 m^2. The nodes are constantly moving, hence the scenario can be characterized as a high mobility scenario. We note that it is not possible to capture all states of human mobility with a single setup, however the scenario is representative of typical day-time pedestrian mobility.

B. Content initialization

In our evaluation scenario we assume that all nodes carry devices and that all are interested in the contents provided.

We assume also that there are N available content items forming a single entry in a feed. The cache of a participating device may initially be filled with N_i randomly chosen content items ($N_i < N$). Whether the cache of a device is initialized with contents depends on the *injection probability* p_i. Thus, throughout its lifetime in the simulation, each participating node strives to obtain as many content items that belong to its subscription as possible. Content items have a mean size of 10 kB, and a standard deviation of 2 kB.

We note that the choice of initially filling the cache of nodes with only N_i items can be seen as an incentive for promoting opportunistic content dissemination. If we chose to inject all N content items into the cache of some nodes, they could directly opt out of dissemination, thus starving the offloading.

C. Full vs. Limited Cooperation

Since offloading is triggered by the network operator, cooperation is often assumed to be inherent to the system. In this work we evaluate the performance of progressive selfishness for two operational modes: *full* and *limited* node cooperation. Full cooperation implies that all nodes in the system are *cooperative*, i.e. they follow the progressive selfishness algorithm; under limited cooperation a subset of the nodes may be willing to deviate from the algorithm. A *non-cooperative* node is assumed to be entirely egocentric: during the data-seeking phase it is not willing to share content items with others, but instead only seeks to obtain items of interest from cooperative nodes, i.e. those operating in a progressively selfish manner; once a non-cooperative node obtains all contents, it becomes strictly selfish to save energy. In other words, non-cooperative nodes only consume resources of other peers in the network without contributing to the data dissemination. There may be different reasons as to why a node chooses to behave in a non-cooperative manner. For instance, a node may be tampered or it may simply have limited battery resources and therefore prefer to opt out of sharing contents for further reduction its own energy consumption.

D. Performance metrics and configurations

We focus on two performance metrics: *goodput* (i.e. application throughput) and *energy consumption* 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 by each node, B_i, 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 contents it can provide.

We differentiate between energy consumption by the opportunistic system, and energy consumption by downloads directly from the cellular network. While the first corresponds to the energy used for *opportunistic* data offloading, E_{opp}, the latter constitutes the energy for *initial* downloads, $E_{cell}^{(1)}$ (i.e. for filling caches with data in order to bootstrap the content exchange in the observed area), and the energy consumed for *follow-up* downloads, $E_{cell}^{(2)}$ (i.e. the energy needed for downloading contents that nodes were not able to obtain opportunistically). Thus, the total energy a node spends is then $E_{tot} = E_{cell}^{(1)} + E_{opp} + E_{cell}^{(2)} [J]$. The energy consumption of the opportunistic system is simply the energy consumed by the radio interface while it is turned on. An energy model for downloading x bytes of data over a 3G network is derived in [16]:

$$E_{cell} = \underbrace{0.025(x) + 3.5}_{\text{transfer energy}} + \underbrace{12.5 \times 0.62}_{\text{tail energy}} + \underbrace{0.02(x/C)}_{\text{maintenance energy}} \quad (2)$$

where C is the download data rate from the cellular network.

We also examine individual energy consumption patterns of cooperative and non-cooperative nodes. The individual energy consumption is defined as the energy used for opportunistic data offloading by a single node normalized with respect to the node's sojourn time in the simulation, $E_{ind} = E_{opp}^{(i)} / t_i$.

We compare the performance of a system in which the radio interface of data-seeking and data-fulfilled nodes is always turned on and no energy saving mechanisms are applied (**ON**), with three types of duty-cycling systems with different levels of selfishness, namely:

1) **DC** — both data-seeking and data-fulfilled nodes duty cycle as defined in [6]; for the purpose of this study the cycling interval is set to 10 s.
2) **DC-SS** — strict selfishness is applied as additional energy saving strategy: nodes opt out of the data dissemination process once they become data-fulfilled.
3) **DC-PS** — progressive selfishness is applied as additional energy saving strategy: data-fulfilled nodes follow the progressive selfishness algorithm outlined earlier.

IV. Performance under Full Cooperation

In the following subsections we present simulation results for the Östermalm scenario when all nodes in the system are fully cooperative and adhere to the progressive selfishness algorithm outlined earlier. We release the assumption of full cooperation in Section V.

If not stated otherwise, goodput is *normalized* with respect to the total amount of data each node would receive if it were to download contents directly via the cellular network. Energy consumption is evaluated both with respect to opportunistic downloads E_{opp}, as well as total energy E_{tot}. E_{opp} is normalized with respect to the energy spent by nodes that do not use energy saving mechanisms; E_{tot} is normalized with respect to the energy that nodes would spend if they were to download all content items directly from the cellular network. In order to provide a fair comparison between E_{opp} and E_{tot} we set the download data rate from the cellular network to equal the download data rate from a neighboring peer, $C = 2$ Mbps. Such data rates are typical for current 3G mobile networks with low mobility users moving at less than 10 km/h [17] as well as for wireless ad-hoc networks using IEEE 802.11b. Although such data rates may be considered low with the rise of gigabit wireless standards like WiGig, promising up to 7 Gbps [18], and the speculation that 5G networks would provide data rates up to 10 Gbps [19], we note that since the

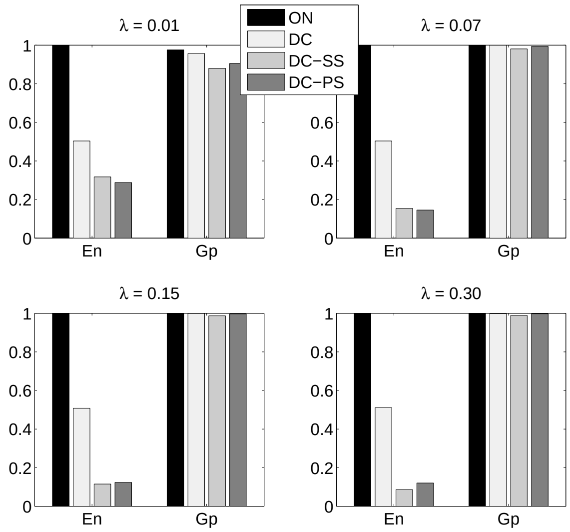

Fig. 3. Comparison of gains in energy consumption (En) and goodput performance (Gp) among configurations for the Östermalm scenario under different arrival rates.

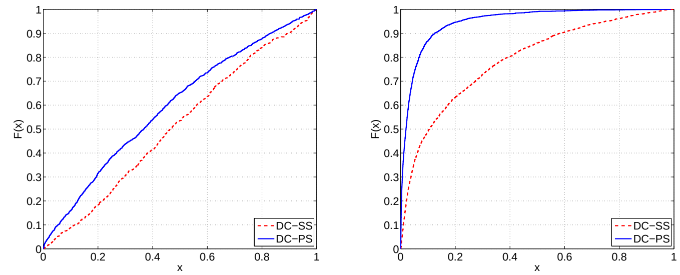

Fig. 4. Cumulative density function of the time (in percentage of lifetime) when a node becomes data-fulfilled under strict and progressive selfishness for an arrival rate of (a) λ = 0.01 nodes/s and (b) λ = 0.30 nodes/s for the Östermalm scenario.

data rate only affects the speed at which data is transmitted, the comparative analysis provided below would still hold under the assumption that networks have similar data rates.

A. Effect of node density

We first evaluate the performance of progressive selfishness under different arrival rates for the Östermalm scenario. Nodes arrive into the observed area according to a Poisson process with rates $\lambda = \{0.01, 0.07, 0.15, 0.30\}$ nodes/s from each of the fourteen entry points. The cache of each node is initially populated with $N_i = 5$ content items chosen uniformly at random from a total of $N = 10$ content items that a mobile operator aims to deliver to all nodes. The reader may assume that all content items form a single piece of contents. The results are presented in Fig. 3.

In sparsely populated scenarios (λ = 0.01 nodes/s), the opportunistic network manages to deliver approximately 97% of the contents when nodes do not use energy saving (ON). Adopting a simple duty cycling scheme delivers similar content volumes but at half the cost in terms of energy consumption. We also see that strictly selfish behavior decreases goodput since data availability is reduced when data-fulfilled nodes opt out of the dissemination process. What is interesting however is that in sparse scenarios progressive selfishness delivers *slightly higher* goodput at a *lower* energy cost in comparison to strict selfishness: Since data-fulfilled nodes continue to participate in the data dissemination, the time data-seeking nodes spend searching for content items decreases, thus the overall energy consumption is also reduced, Fig. 4(a). We note that progressive selfishness reduces energy consumption by 75%, and this reduction comes at only 6% goodput loss.

As the density increases, the goodput performance of progressive selfishness becomes comparable to the goodput performance of a system in which nodes do not adopt energy saving mechanism, while energy consumption is decreased by approximately 85%. An interesting observation is that at high densities (λ = 0.30 nodes/s) the energy consumed by strictly selfish nodes becomes *less* than the energy consumed by progressively selfish nodes. This may seem counter-intuitive. Fig. 4(b) provides some insight. In the case of strict selfishness, 50% of the nodes become data-fulfilled in the first 10% of their lifetime while in the case of progressive selfishness, this number increases to almost 90%. Due to the high initial data availability, only 12% of the data-fulfilled nodes consecutively disseminate contents to other peers. When the mobile operator wishes to disseminate small amounts of data among users and the node density is high (λ = 0.30 nodes/s), the energy spent by nodes that are in a data-seeking state becomes negligible in comparison to the energy spent by nodes that are in a data-fulfilled state. To address this, we have previously proposed a simple enhancement of the progressive selfishness algorithm [20], however we do not discuss it further in the current work.

Table I presents the mean number of contact opportunities a node experiences throughout its lifetime under different arrival rates. (Here a contact is accounted for when a node receives a beacon from another peer in range.) Due to the small amount of data that needs to be disseminated among nodes, as the arrival rate increases, the adaptive nature of progressive selfishness achieves more than 90% reduction in the number of peers a node encounters, filtering out contacts with peers that cannot contribute to the node (i.e. data-seeking nodes that have the same content items as the node, or other data-fulfilled nodes if the node is already data-fulfilled).

We believe that scenarios in which small amounts of contents have to be disseminated to a large crowd of participants will not constitute the majority of use-cases when it comes to mobile data offloading via opportunistic communication. In the following subsection we explore the performance of progressive selfishness with respect to the amount of data which needs to be disseminated opportunistically in the observed area in order to evaluate the scalability of the algorithm.

B. Effect of content availability

We showed that progressive selfishness achieves high data dissemination rates and up to 85% energy savings independent of the node density. In this subsection we increase the traffic, and assume that a mobile operator wishes to disseminate information of size 1 MB to all nodes entering the observed

TABLE I
MEAN NUMBER OF CONTACT OPPORTUNITIES UNDER DIFFERENT ARRIVAL RATES FOR THE ÖSTERMALM SCENARIO.

Arrival rate nodes/s	Mean density nodes/m^2	Mean number of contact opportunities			
		ON	DC	DC-SS	DC-PS
λ = 0.01	0.006	5.56	5.49	2.23	2.19
λ = 0.07	0.043	40.17	38.87	3.43	4.71
λ = 0.15	0.085	85.17	73.24	3.51	6.91
λ = 0.30	0.129	140.47	115.39	3.96	10.91

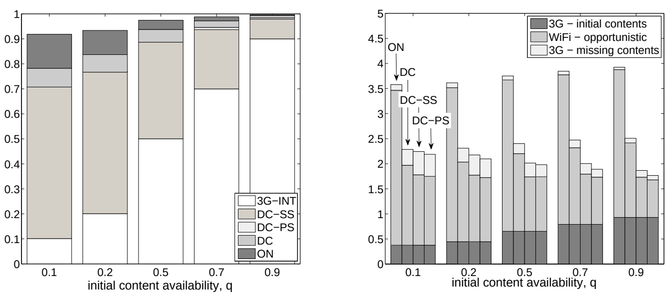

(a) (b)

Fig. 5. Effect of initial content availability for the Östermalm scenario with λ = 0.01 nodes/s: (a) goodput performance, and (b) total energy consumption.

area, a short video clip, for instance. We further divide this data into entries of mean size 10 kB as before, and we define $q = N_i/N$ to be the initial content availability, i.e. how many content items are fetched from the cellular network by nodes entering the observed area. We then vary the initial content availability q = {0.1, 0.2, 0.5, 0.7, 0.9}; $q = 0.1$ represents a scenario in which contents should be mostly offloaded opportunistically, and $q = 0.9$ corresponds to a scenario in which most of the contents is initially pre-loaded into the nodes.

Fig. 10 presents the results for a sparsely populated Östermalm scenario, i.e. λ = 0.01 nodes/s. Observe that in Fig. 10(a) the 3G-INT part represents data downloaded directly from the mobile operator to bootstrap the offloading process. The rest of the bars are stacked *behind* one another and illustrate in an *incremental* manner the amount of data that is offloaded across different configurations. In other words, DC-SS represents how much data strict selfishness delivers on top of the initial downloads from the cellular network, DC-PS shows how much *more* goodput progressive selfishness delivers on top of strict selfishness, DC illustrates how much goodput duty cycling delivers on top of progressive selfishness, and ON shows the amount of *additional* data delivered on top of duty cycling. This representation is adopted throughout the rest of the paper.

Even when the initial content availability is low, $q = 0.1$, progressive selfishness offloads approximately 60% of the data, Fig. 10(a) and Table II. If nodes instead kept their radio interfaces turned on, the opportunistic network would offload more than 80% of the total data for the same initial content availability. We note that on one hand this increase of 20% in goodput requires approximately 55% more energy for offloading via opportunistic communication (comparing the bars in Fig. 10(c) representing the energy consumed by opportunistic downloads). On the other hand, if this data is not disseminated opportunistically, nodes would have to download it from the

TABLE II
OFFLOADED TRAFFIC FOR THE ÖSTERMALM SCENARIO UNDER DIFFERENT CONFIGURATIONS AND DIFFERENT ARRIVAL RATES.

Configuration		$q = 0.1$	0.2	0.5	0.7	0.9
λ = 0.01	ON	81.7%	73.8%	47.3%	28.8%	9.7%
	DC	68.1%	64.1%	43.7%	27.2%	9.4%
	DC-SS	60.5%	56.3%	38.6%	23.6%	7.8%
	DC-PS	60.5%	56.6%	38.6%	24.5%	8.4%
λ = 0.15	ON	89.4%	79.5%	49.8%	30.0%	10.0%
	DC	88.6%	79.0%	49.6%	29.9%	10.0%
	DC-SS	86.9%	77.6%	48.8%	29.5%	9.8%
	DC-PS	87.7%	78.4%	49.2%	29.9%	10.0%

cellular network upon exiting the observed area. The reader may then ask whether it is more profitable (in terms of energy consumption) to spend energy during the opportunistic offload or to spend energy for follow-up downloads from the cellular network.

Fig. 10(c) illustrates the mean energy consumption in each of the three energy consumption states: *initial* download for bootstrapping the content exchange in the observed area, *opportunistic* download for offloading mobile data and *follow-up* download for obtaining missing content items from the cellular network before leaving. Observe that all values in Fig. 10(c) are normalized with respect to the estimated energy a node would consume if it were to download all data directly from the cellular network, given in Eq. (2). In sparse areas, progressive selfishness decreases both the opportunistic and the total energy consumption across all values of the initial content availability. Compared to the performance of a duty-cycling system the reduction in energy consumption is insignificant for small values of the initial content availability and increases with the increase of q. Due to the sparsity of the scenario, and the amounts of data that nodes need to obtain opportunistically, very few nodes can utilize the benefits of progressive selfishness, since their inactivity window is constantly shrunk to 1. We also note that across all values of the initial content availability the total energy spent under progressive selfishness is doubled as compared to the energy that nodes would spend if they were to download contents directly from the cellular network. Although there is a reduction in the energy consumption with the increase of the initial content availability, Fig. 10(c), in the context of mobile data offloading operators would prefer to keep the value of q as low as possible. Thus, in this sparse scenario with high data loads there is a clear trade-off between the requirements of mobile nodes (low energy consumption), and those of mobile operators (high offload ratio). Progressive selfishness satisfies both parties by decreasing the total energy consumption in mobile nodes without greatly penalizing the load reduction in the cellular network.

Fig. 6 shows the goodput performance and energy consumption for the Östermalm scenario when the density is increased (λ = 0.15 nodes/s). Due to the larger amount of participants in the opportunistic content dissemination process, even at low initial content availability (q = 0.1), progressive selfishness offloads approximately 87% of the data, Fig. 11(a) and Table II. Leaving the radio interface turned on contributes just a couple of percent more to the goodput for the same

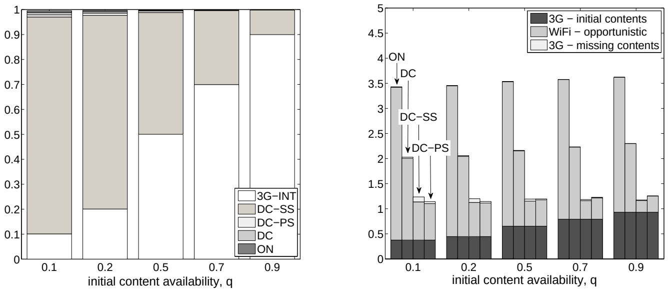

(a) (b)

Fig. 6. Effect of initial content availability for the Östermalm scenario with λ = 0.15 nodes/s: (a) goodput performance, and (b) total energy consumption.

value of q. In terms of energy savings, progressive selfishness not only decreases the energy consumption across all values of the initial content availability but the total energy spent for disseminating the information in the observed area is comparable to the energy a node would spend when downloading all data directly from the cellular network.

To summarize, independent of the node density in the area, progressive selfishness offloads mobile data at a lower energy cost than other solutions across all values of the initial content availability. In the context of mobile data offloading, it is recommended that the initial content availability is kept low.

C. *Effect of injection probability*

In the previous sections we assumed that each node downloads some contents directly from the cellular network upon entrance in order to participate in the dissemination process. In reality, however, only a subset of nodes, often referred to as a target set, may be chosen to carry data initially, while all other nodes attempt to download contents opportunistically from them. Here we do not discuss the optimal choice of the target set but instead choose nodes uniformly at random according to an injection probability p_i. (For details on the optimal choice of the initial target set we refer the reader to [2].)

Fig. 7 presents results for the Östermalm scenario with λ = 0.15 for three different values of the injection probability p_i = {0.01, 0.05, 0.1}, a total of N = 10 content items of mean size 10 kB each, and an initial content availability q = 0.5.

Even at p_i = 0.01, progressive selfishness decreases the energy consumption by 80%, Fig. 7(c), while satisfying approximately 95% of the overall traffic demand, Fig. 7(a). Observe that the bars representing energy consumption in Fig. 7(c) are normalized with respect to the energy spent by ON, and are stacked behind one another to illustrate the incremental energy savings across configurations. Comparing to a strictly selfish behavior which is usually discussed in literature, progressive selfishness decreases by half the energy cost, while providing higher goodput. To explain this phenomenon, we define *user satisfaction* as the percentage of nodes that by the end of their lifetime have obtained the whole set of content items, and do not need to perform follow-up downloads upon exiting. The results are illustrated in Fig. 7(b). We see that although strict selfishness delivers similar goodput performance as progressive selfishness, the user satisfaction is only 12% as compared to 80% for progressive selfishness with p_i = 0.01.

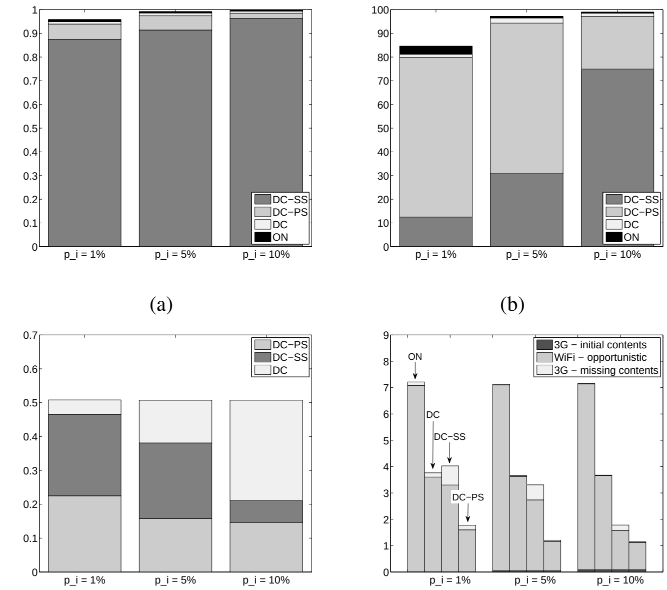

(a) (b) (c) (d)

Fig. 7. Effect of injection probability for Östermalm scenario with λ = 0.15 nodes/s: (a) goodput performance, (b) user satisfaction, (c) energy savings during opportunistic downloads, and (d) total energy savings.

The total energy consumption under different injection probabilities is presented in Fig. 7(d). When p_i = 0.01 progressive selfishness decreases the total energy consumption by a factor of 5 as compared to a configuration in which nodes have their radio interfaces turned on, and the main reduction does come from the energy savings during opportunistic offloading. An interesting observation is that strictly selfish behavior requires higher energy consumption for follow-up downloads due to the low user satisfaction factor. As the injection probability increases, the total energy consumed by progressively selfish nodes is further decreased, and for p_i = 0.1 it becomes comparable to the energy that nodes would spend if they were to download everything from the cellular network.

V. Performance under Limited Cooperation

In this section we present results for the Östermalm scenario under limited cooperation. We allow a subset of all nodes in the system to deviate from the progressive selfishness algorithm and behave in a non-cooperative manner as defined in Section III-C. Upon entering the area, each node chooses uniformly at random whether to adhere to the progressive selfishness algorithm while searching for the contents announced by the mobile operator or to deviate from it.

A. *Effect of node density*

We first evaluate the performance of the opportunistic network operating under progressive selfishness with limited cooperation in sparse (λ = 0.01 nodes/s) and dense (λ = 0.15 nodes/s) environments, Fig. 8 and 9 respectively. We vary the percentage of non-cooperative nodes in the system $k \in [5, 50]$ and we compare the performance with respect to

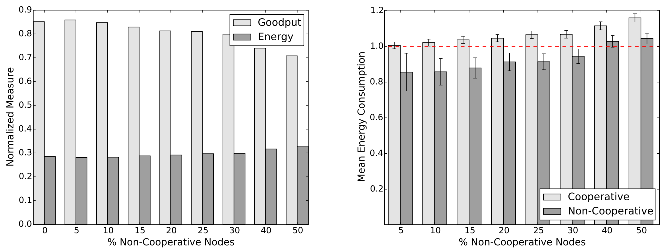

(a) (b)

Fig. 8. Effect of node density under limited cooperation for the Östermalm scenario with λ = 0.01 nodes/s: (a) normalized goodput and energy consumption on system level, and (b) normalized mean energy consumption for cooperative vs. non-cooperative nodes; dotted line represents the energy consumed by DC-PS under full cooperation.

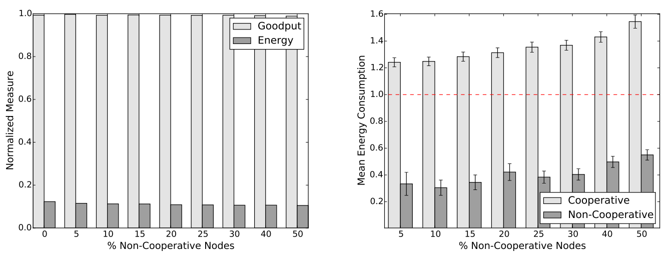

(a) (b)

Fig. 9. Effect of node density under limited cooperation for the Östermalm scenario with λ = 0.15 nodes/s: (a) normalized goodput and energy consumption on system level, and (b) normalized mean energy consumption for cooperative vs. non-cooperative nodes; dotted line represents the energy consumed by DC-PS under full cooperation.

the performance of a system with full cooperation, $k = 0$. Furthermore, we assume that the amount of data the mobile operator wishes to disseminate to the participating nodes is small, a total of $N = 10$ content items.

In sparsely populated scenarios, Fig. 8(a), the amount of data delivered by the opportunistic network decreases with the increase of the non-cooperative population. Simultaneously the energy consumed for opportunistic content distribution increases since nodes spend on average more time in the data-seeking phase searching for contents. (Observe that the goodput and the energy consumption are normalized with respect to the goodput and energy consumption in ON, i.e. when all nodes do not apply any energy saving mechanisms and are entirely altruistic.) Fig. 8(a) shows that even at low densities an opportunistic network operating with progressive selfishness can support a small population of non-cooperative nodes ($k = 10\%$) without loss in goodput. If we examine the individual energy consumed by cooperative nodes (i.e. nodes that follow the progressive selfishness algorithm) and non-cooperative nodes, Fig. 8(b), we see that non-cooperative nodes are able to decrease their energy consumption by 15% for small values of k however this decrease comes at a price of increased energy consumption for nodes that adhere to the progressive selfishness algorithm. Interestingly, as the percentage of non-cooperative nodes increases it becomes non-beneficial for nodes to deviate from the DC-PS scheme; when $k > 40\%$ both cooperative and non-cooperative nodes spend more energy than they would have spent if they were all following the DC-PS algorithm (the dotted line in Fig. 8(b)).

As the density increases the progressive selfishness algorithm is capable of supporting the system performance both in terms of goodput and energy consumption even when 50% of the population behaves in a non-cooperative manner, Fig. 9(a). Thus, it is impossible to detect the presence of non-cooperative nodes in the system by simply examining the means of the performance metrics. Fig. 9(b) however reveals that differentiating between cooperative and non-cooperative nodes is possible when examining the individual energy consumption. On average cooperative nodes experience higher energy consumption (with respect to DC-PS under full cooperation, dotted line) even at small values of k. Furthermore, non-cooperative nodes experience approximately 50% of energy savings with respect to DC-PS even at $k = 50\%$. One way to interpret this result is to discuss the need of introducing proper incentives for nodes to adhere to the DC-PS algorithm. However, an alternative view point would be that *even without* incentives, the performance of progressive selfishness does not decrease significantly. As mentioned earlier, a node may choose to be non-cooperative simply because it currently does not have enough battery resources; such a scenario is highly probable in urban environments since mobile devices often have different battery resources at any point in time. The proposed algorithm is thus able to achieve high offload ratio under limited cooperation.

B. Effect of content availability

In Section IV-B we showed that progressive selfishness performs well at low values of the initial content availability, $q = 0.1$, even when the mobile operator wishes to offload contents of a larger size. We here explore whether this statement holds in the presence of limited cooperation.

In sparse scenarios the larger the contents to be disseminated and the lower the initial content availability, the more prominent the reduction in goodput obtained via opportunistic contacts, Fig. 10(a). On the one hand, at $k = 50\%$ the offload ratio is decreased by more than 50% with respect to the goodput achieved by ON. On the other hand, at low densities and small values of q non-cooperative nodes achieve little to no energy savings by deviating from the progressive selfishness algorithm, Fig. 10(c). Nodes are forced to spend most of their lifetime in the system searching for content items due to the low initial content availability and the larger content size. However, the lack of full cooperation requires more nodes to rely on follow-up downloads from the mobile network thus increasing the total energy consumption, Fig. 10(b), although on average the energy consumed by opportunistic downloads does not change significantly.

In dense scenarios, the larger node population again obscures the effect of limited cooperation. The goodput achieved by opportunistic downloads deteriorates only at high levels of non-cooperation ($k > 40\%$), Fig. 11(a), which results into increased energy for follow-up downloads, Fig. 11(b). However, even for values of $k > 40\%$ it is still beneficial for nodes to behave in a non-cooperative manner due to the individual energy savings they achieve, Fig. 11(c).

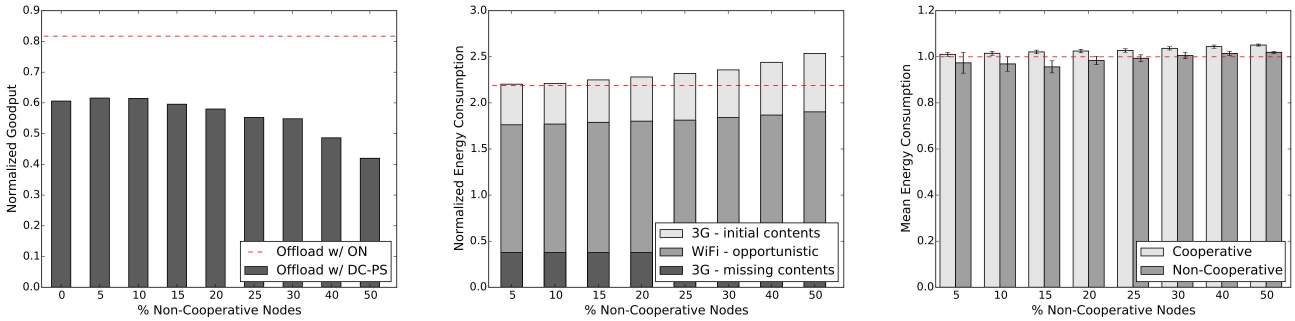

(a) (b) (c)

Fig. 10. Effect of initial content availability $q = 0.1$ for Östermalm scenario with λ = 0.01 nodes/s and large size contents: (a) goodput via opportunistic downloads, (b) total energy consumption, and (c) energy consumption for cooperative vs. non-cooperative nodes; dotted line in plots (b) and (c) represents the energy consumed by DC-PS under full cooperation.

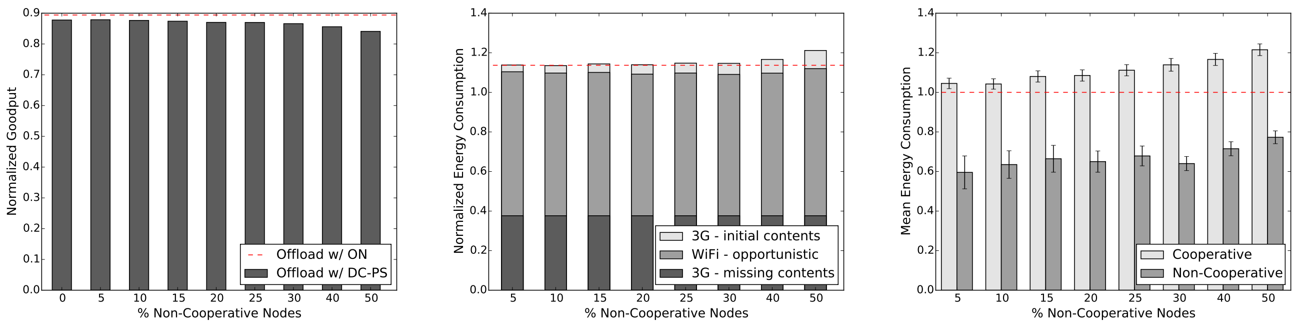

(a) (b) (c)

Fig. 11. Effect of initial content availability $q = 0.1$ for Östermalm scenario with λ = 0.15 nodes/s and large size contents: (a) goodput via opportunistic downloads, (b) total energy consumption, and (c) energy consumption for cooperative vs. non-cooperative nodes; dotted line in plots (b) and (c) represents the energy consumed by DC-PS under full cooperation.

C. *Effect of injection probability*

Finally, we examine the effect of injection probability under limited cooperation. Again, we assume that only a small portion of nodes, $p_i = \{0.01, 0.05, 0.10\}$, entering the observed area initially obtain contents directly from the mobile operator, while the rest of the nodes attempt to obtain contents opportunistically. The total number of content items is $N = 10$, and the initial content availability is $q = 0.5$ for those nodes chosen to carry contents upon entering the area. Observe that we assume that the mobile operator does not differentiate between cooperative and non-cooperative nodes when it decides to inject contents in their caches upon arrival. In other words, a node that is initially injected with contents may be non-cooperative in nature. Thus, the amount of useful nodes which carry contents into the system may be less than p_i.

Fig. 12 shows the effect of injection probability on the individual energy consumption of cooperative and non-cooperative nodes for different values of p_i. It is not beneficial for non-cooperative nodes to deviate from the progressive selfishness algorithm only for high values of non-cooperation $k = 50\%$ and small values of the injection probability $p_i = 0.01$, Fig. 12(a). This is also notable in the reduction of goodput, Fig. 13(a), as well as the increase in the total energy consumption, Fig. 13(b). With the increase of the injection probability, the presence of non-cooperative nodes becomes less notable on a system level. Furthermore, it becomes beneficial (in terms of energy consumption) for nodes to behave in a non-cooperative manner since there are enough participants in the system to support the opportunistic data dissemination without severely deteriorating the offloaded goodput.

VI. Related Work

In this section we discuss previous work conducted in the field of mobile data offloading, and in the area of selfishness in opportunistic networks.

A. *Mobile data offloading*

Recent solutions for alleviating traffic load on cellular networks can be divided in two main categories: offloading to femtocells or existing WiFi networks [21], [22], and offloading through opportunistic communication. Although the large body of work produced in the area of offloading cellular traffic to femtocells and WiFi networks, such approach is limited to possible deployments and the availability of Internet access. Offloading mobile data through opportunistic communication does neither depend on available deployment, nor on Internet access, and has thus become a popular candidate for traffic offloading in recent years. Different studies attempt to optimize the traffic volumes delivered to end users through opportunistic communication. In [2] Han et al. study a target-set selection problem for choosing initial data carriers in order

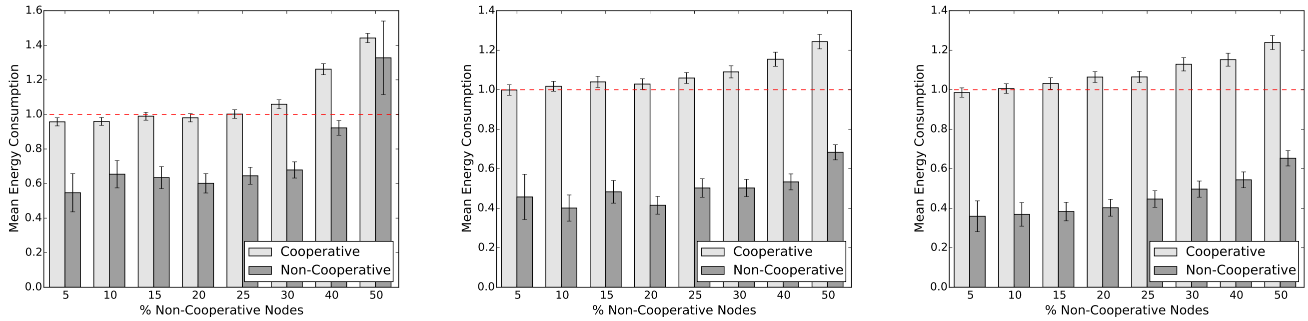

(a) (b) (c)

Fig. 12. Effect of injection probability on energy consumption for cooperative vs. non-cooperative nodes for the Östermalm scenario with λ = 0.15 nodes/s: (a) $p_i = 1\%$, (b) $p_i = 5\%$, and (c) $p_i = 10\%$.

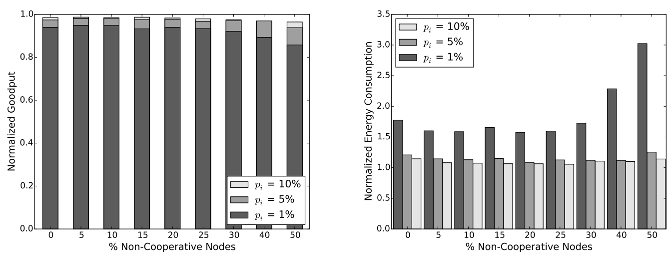

(a) (b)

Fig. 13. Effect of injection probability on (a) goodput and (b) total energy consumption for the Östermalm scenario with λ = 0.15 nodes/s.

to minimize the amounts of mobile data traffic. Lu et al. propose an opportunistic forwarding protocol for increasing the probability of data delivery [3]. However neither of the works consider the energy consumed by mobile devices in the process of offloading mobile data through opportunistic communication. In a recent work Ding et al. [23] recognize the importance of decreasing energy consumption during the offloading process, and study the energy savings that can be achieved when mobile data is offloaded to WiFi networks. In contrast, we investigate the energy savings achieved when mobile data is offloaded via opportunistic communication. Recently Mota et al. [24] introduced energy into their multi-criteria framework for opportunistic mobile data offloading however they only use energy consumption as a decision-making parameter for whether a node should participate in the offloading process. Instead, we profile the energy consumption of nodes that participate in the offloading process.

B. Selfishness in opportunistic networks

All work done in the area of mobile data offloading by default considers nodes to act altruistically throughout the mobile data offloading process via opportunistic communication. A separate body of work has been devoted to issues related to node selfishness in opportunistic networking. In [25] Hui et al. study the impact of different distributions of altruism on the throughput and delay of opportunistic networks, and demonstrate that opportunistic networks are robust towards the form of altruism distribution, however they fail to consider the impact of selfishness on the energy consumption in nodes. Lu et al. [26] and Zhao et al. [8] propose incentive schemes for promoting altruism in opportunistic networks, however neither of the works discusses the implications, in terms of energy consumption, of altruistic behavior in battery powered devices. Ciobanu et al. [9] devise a mechanism for detecting and avoiding selfish nodes throughout the routing process in opportunistic networks. In contrast to all mentioned studies, we promote selfishness in opportunistic networks used for mobile data offloading in order to decrease the energy consumption of participating devices.

VII. Conclusion

In this paper we propose *progressive selfishness*, an adaptive and scalable energy-aware algorithm for improving energy-efficiency in mobile devices in the context of mobile data offloading via opportunistic communication. Previous work in the area of mobile data offloading via opportunistic communication focuses mainly on maximizing the data delivery to end users. However, since mobile nodes are battery-powered, opportunistic networking can only be considered a viable mechanism for offloading data *if* it delivers high content volumes at a low energy cost. Thus, we evaluated the performance of the proposed progressive selfishness algorithm *both* in terms of application throughput (i.e. goodput) and energy consumption via extensive trace-driven simulations of realistic pedestrian mobility. We introduced two modes of nodal cooperation: full and limited, with respect to the percentage of nodes in the system that adopt progressive selfishness. We showed that under full cooperation the algorithm decreases energy consumption of participating nodes during the opportunistic downloads with up to 85% across different node densities in the observed area without significantly compromising goodput. We then investigated the effects of initial content availability, and observed that the performance of progressive selfishness is robust to it across different node densities. We also showed that at higher densities the energy spent by nodes that adopt progressive selfishness becomes comparable to the energy nodes would spend for downloading the same data directly from the cellular network. Thus, progressive selfishness not only offloads up to 85% of the mobile data traffic, but does the offloading at a similar price in terms of energy. We also demonstrated that progressive selfishness scales with the injection probability of content carriers in the observed area, and we showed that the performance of the algorithm is independent of the area in which data is being offloaded. Finally, we investigated the performance of progressive selfishness under limited cooperation across different node densities, content availabilities and injection probabilities, and showed that progressive selfishness is robust against the presence of non-cooperative nodes in the system in terms of offloaded traffic volumes. We also demonstrated that in dense scenarios, progressive selfishness could tolerate up to 50% of non-cooperative nodes. Moreover, we showed that by simply evaluating average performance metrics it is impossible to determine the presence of non-cooperative nodes in the system, however examining individual energy consumption patterns may provide better insights.

References

[1] "Cisco visual networking index: Global mobile data traffic forecast update 2014 - 2019 white paper," 2015.

[2] 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.

[3] L. Xiaofeng, H. Pan, and P. Lio, "Offloading mobile data from cellular networks through peer-to-peer wifi communication: A subscribe-and-send architecture," *Communications, China*, vol. 10, no. 6, pp. 35–46, June 2013.

[4] Y. Li, M. Qian, D. Jin, P. Hui, Z. Wang, and S. Chen, "Multiple mobile data offloading through disruption tolerant networks," *Mobile Computing, IEEE Transactions on*, vol. 13, no. 7, pp. 1579–1596, July 2014.

[5] Ó. 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.

[6] S. T. Kouyoumdjieva and G. Karlsson, "Energy savings in opportunistic networks," in *Proc. IEEE/IFIP WONS*, Obergurgl, Austria, 2014.

[7] W. Gao and Q. Li, "Wakeup scheduling for energy-efficient communication in opportunistic mobile networks," in *Proc. INFOCOM*, 2013.

[8] G. Zhao, M. Chen, and X. Wei, "Ris: A reciprocal incentive scheme in selfish opportunistic networks," *Wireless Personal Communications*, vol. 70, no. 4, pp. 1711–1734, 2013.

[9] R.-I. Ciobanu, C. Dobre, M. Dascălu, S. Trăusan-Matu, and V. Cristea, "Sense: A collaborative selfish node detection and incentive mechanism for opportunistic networks," *Journal of Network and Computer Applications*, vol. 41, no. 0, pp. 240 – 249, 2014.

[10] S. T. Kouyoumdjieva and G. Karlsson, "The virtue of selfishness: Device perspective on mobile data offloading," in *Proc. IEEE WCNC*, March 2015.

[11] N. L. Johnson, S. Kotz, and N. Balakrishnan, *Continuous Univariate Distributions*. Wiley, 1995.

[12] Ó. R. Helgason and K. V. Jónsson, "Opportunistic networking in OMNeT++," in *Proc. SIMUTools, OMNeT++ workshop*, 2008.

[13] L. M. Feeney and D. Willkomm, "Energy framework: an extensible framework for simulating battery consumption in wireless networks," in *Proc. SIMUTools*, 2010.

[14] S. T. Kouyoumdjieva, Ó. R. Helgason, and G. Karlsson, "CRAWDAD data set kth/walkers (v. 2014-05-05)," Downloaded from http://crawdad.org/kth/walkers/, May 2014.

[15] "Legion Studio," http://www.legion.com/.

[16] N. Balasubramanian, A. Balasubramanian, and A. Venkataramani, "Energy consumption in mobile phones: a measurement study and implications for network applications," in *Proc. ACM IMC*, 2009, pp. 280–293.

[17] "Cellular Standards for Third Generation," http://www.itu.int/osg/spu/imt-2000/technology.html.

[18] "WiGig and the future of seamless connectivity," http://www.wi-fi.org/discover-wi-fi/wigig-certified.

[19] "Huawei: 5G, A Technology Vision," http://www.huawei.com/5gwhitepaper/.

[20] S. T. Kouyoumdjieva and G. Karlsson, "Energy-aware opportunistic mobile data offloading for users in urban environments," in *Proc. IFIP/TC6 Networking*, 2015.

[21] K. Lee, J. Lee, Y. Yi, I. Rhee, and S. Chong, "Mobile data offloading: How much can wifi deliver?" *Networking, IEEE/ACM Transactions on*, vol. 21, no. 2, pp. 536–550, April 2013.

[22] L. Gao, G. Iosifidis, J. Huang, and L. Tassiulas, "Economics of mobile data offloading," in *Proc. IEEE INFOCOM*, April 2013, pp. 3303–3308.

[23] A. Ding, B. Han, Y. Xiao, P. Hui, A. Srinivasan, M. Kojo, and S. Tarkoma, "Enabling energy-aware collaborative mobile data offloading for smartphones," in *Proc. IEEE SECON*, June 2013, pp. 487–495.

[24] V. F. Mota, D. F. Macedo, Y. Ghamri-Doudanez, and J. M. S. Nogueira, "Managing the decision-making process for opportunistic mobile data offloading," in *Proc. IEEE NOMS*, May 2014, pp. 1–8.

[25] P. Hui, K. Xu, V. Li, J. Crowcroft, V. Latora, and P. Lio', "Selfishness, altruism and message spreading in mobile social networks," in *Proc. IEEE INFOCOM Workshops*, April 2009, pp. 1–6.

[26] R. Lu, X. Lin, H. Zhu, X. Shen, and B. Preiss, "Pi: A practical incentive protocol for delay tolerant networks," *Wireless Communications, IEEE Transactions on*, vol. 9, no. 4, pp. 1483–1493, April 2010.

Sylvia Kouyoumdjieva is currently a PhD student at the Laboratory for Communication Networks at KTH, the Royal Institute of Technology. She received a BSc in Telecommunications from the Technical University of Sofia, Bulgaria (2006) and a MSc in Telecommunications from the Royal Institute of Technology, Stockholm, Sweden (2009). Her research focuses on opportunistic networking and mobile peer-to-peer systems.

Gunnar Karlsson is professor of KTH, the Royal Institute of Technology, since 1998; he is the director of the Laboratory for Communication Networks. He has previously worked for IBM Zurich Research Laboratory and the Swedish Institute of Computer Science (SICS). His Ph.D. is from Columbia University (1989), New York, and the M.Sc. from Chalmers University of Technology in Gothenburg, Sweden (1983). He has been visiting professor at EPFL, Switzerland, and the Helsinki University of Technology in Finland, and ETH Zurich in Switzerland. His current research relates to quality of service, wireless LAN developments and wireless content distribution.