EURECOM
Department of Mobile Communications
Campus Sophia Tech
Les Templiers
450 route des Chappes
B.P. 193
06410 Biot
FRANCE

Research Report RR-13-283

Understanding the Effects of Social Selfishness on the Performance of Heterogeneous Opportunistic Networks

July 20$^{\mathrm{th}}$, 2013
Last update July 20$^{\mathrm{th}}$, 2013

Pavlos Sermpezis and Thrasyvoulos Spyropoulos

[1]EURECOM's research is partially supported by its industrial members: BMW Group, Cisco, Monaco Telecom, Orange, SAP, SFR, STEricsson, Swisscom, Symantec.

Tel : (+33) 4 93 00 81 00
Fax : (+33) 4 93 00 82 00
Email : {pavlos.sermpezis,thrasyvoulos.spyropoulos}@eurecom.fr

Understanding the Effects of Social Selfishness on the Performance of Heterogeneous Opportunistic Networks

Pavlos Sermpezis and Thrasyvoulos Spyropoulos

Abstract

In Opportunistic Networks the majority of communication mechanisms make use of relay nodes for delivering the messages. Any possible unwillingness of the relay nodes to cooperate, can affect gravely the performance of routing protocols and message dissemination techniques. In this paper we propose a framework for analysing cases where the level of cooperation, or the *selfishness* of nodes, is related to the social ties of the nodes or their mobility patterns. We model selfishness in Heterogeneous Opportunistic Networks and investigate the effect of it in communication performance. Specifically, we analyse and derive expressions for important metrics, namely the message delivery delay, the average power consumption and the message delivery probability. We demonstrate the applicability of our results in various application scenarios and validate their accuracy with simulations on both synthetic and real-world networks.

Index Terms

Opportunistic Networks, Heterogeneous Mobility, Performance Modeling, Social Selfishness

Contents

1 Introduction **1**

2 Preliminaries **3**
2.1 Mobility Model 3
2.2 Selfishness Models 5

3 Message Delivery Delay **7**
3.1 Delay in Heterogeneous Networks 7
3.2 Effect of Selfishness 9
3.3 Case Studies 12
3.4 Validation 13
3.4.1 Synthetic Simulations 13
3.4.2 Real-world Traces 14

4 Performance vs Power Consumption **15**
4.1 Delivery Delay vs Power Consumption 16
4.1.1 Validation 19
4.2 Delivery Probability vs Power Consumption 20
4.2.1 Opportunistic Content Sharing 21
4.2.2 Evaluation 23

5 Related Work **24**

6 Conclusion **27**

A Proof of Lemma 3.1 **27**

B Proof of Lemma 3.2 **28**

C Proof of Result 5 **28**

List of Figures

1 Epidemic spreading over a homogeneous network with N nodes . 4
2 Epidemic spreading in a heterogeneous network with 4 nodes . . . 4
3 *Delivery Delay* in networks with $N = 100$ nodes and varying *Mobility* characteristics ($\mu_\lambda = 1$ and $cv_\lambda \in [0, 3]$) for three different selfishness policies and under epidemic routing. The theoretical values of Delivery Delay for two parameters (p_0) for each selfishness policy are denoted with dashed lines and the corresponding simulations' average delivery delays are denoted with dots. 13
4 Relative decrease of delay, $\frac{E[T_D]}{E[T_D^{max}]}$, of SnW routing, in scenarios with (a) Policy B ($p_1 = 0, p_2 = 1$ and variable p_0) selfishness and $L = 10$ copies, and (b) Policy D ($p_0 = 0.2$ and variable m) selfishness and $L = 20$ copies. 15
5 Relative decrease of delay, $\frac{E[T_D]}{E[T_D^{max}]}$, of SnW routing, in scenarios with (a) Policy B ($p_1 = 0, p_2 = 1$ and variable p_0) selfishness and $L = 20$ copies, and (b) Policy D ($p_0 = 0.2$ and variable m) selfishness and $L = 20$ copies. 16
6 Power consumption - message delivery delay trade off. Synthetic simulations with (a) uniform and (b) non-uniform selfishness policies. Simulations on the (c) Infocom and (d) Sigcomm real traces of both uniform and non-uniform selfishness policies scenarios. . . 19
7 (a) Probability Delivery Ratio of a content of policy A selfishness (blue) and policy C selfishness (black) for different power consumption levels. (b) Relative difference of the Probability Delivery Ratio between Policy C and Policy A selfishness, i.e. $\frac{PDR_C - PDR_A}{PDR_A}$. Number of content copies $M = 1$ and $T = \frac{10}{\mu_\lambda}$. Mobility characteristics: $\mu_\lambda = 1$, $cv_\lambda = 1$. 24
8 (a) Probability Delivery Ratio of a content of policy A selfishness (blue) and policy C selfishness (black) for different power consumption levels. (b) Relative difference of the Probability Delivery Ratio between Policy C and Policy A selfishness, i.e. $\frac{PDR_C - PDR_A}{PDR_A}$. Number of content copies $M = 5$ and $T = \frac{5}{\mu_\lambda}$. Mobility characteristics: $\mu_\lambda = 1$, $cv_\lambda = 2$. 25
9 Relative difference of the Probability Delivery Ratio between Policy C and Policy A selfishness, i.e. $\frac{PDR_C - PDR_A}{PDR_A}$, in the *Sigcomm trace* with (a) $M = 3$, (b) $M = 5$ number of copies and $T = 20/\mu_\lambda$, and in *Infocom trace* with (c) $M = 3$, (d) $M = 5$ number of copies and $T = 10/\mu_\lambda$. 25

1 Introduction

Opportunistic Networks and Delay Tolerant Networks (DTNs) are envisioned to support communication in challenging environments with total or partial absence of infrastructure (e.g. emergency situations after a disaster, mobile sensor networks) and/or to enhance existing networks (e.g. offloading of cellular networks, collaborative cloud computing). They are composed mainly of nodes moving in an area much larger than their transmission range. Data exchange between nodes (i.e. portable devices like smartphones, laptops, PDAs etc.) can take place only when they are within transmission range of each other, or, as it is also called, when they are *in contact*.

As a result, traditional end-to-end connectivity is very volatile, intermittent, or often absent. Message dissemination from a source to a destination node could be achieved by direct transmission [1], when source and destination come in contact. Yet, this leads to large delivery delays, low throughput and high packet losses. To improve the situation, other ("relay") nodes could be used that *carry* the message and *forward* it to the destination (or to other, better relays) when they contact with it.

Many routing techniques have been proposed for Opportunistic Networks. In social-oblivious ("random") routing protocols, e.g. epidemic [2], two-hop [3], spray and wait [4], replication is used as a diversity mechanism to improve performance. A small or large number of relays, chosen randomly, receive a copy of the message and can independently relay it. In utility-based or social-aware routing (see [5] for a survey) "good" relays are selected based on social (or other) characteristics. The main goal of these protocols is to achieve good performance (delay and delivery probability) with the minimum amount of overhead (e.g. number of transmissions or relays used per message), and many of them have been shown to achieve very good tradeoffs in different mobility settings.

However, the vast majority of works proposing, modeling, or optimizing protocols for Opportunistic Networks assume cooperation of nodes in relaying messages: when the protocol dictates that a relay node should receive or transmit a message (neither destined to nor originating from it), it does. In practice, a relay node might: (i) never be willing to carry traffic for 3rd parties, (ii) be willing to only perform some number of transmissions/receptions for relay traffic, or (iii) be more willing to receive or transmit traffic from nodes it has some (social) "ties" with. The reasons for this reluctance range from privacy concerns (e.g. not trusting an exchange with unknown nodes) to resource consumption (e.g. battery depletion). Such behaviors are natural, and could significantly degrade the predicted performance of the above protocols.

To this end, some recent works have used both simulations and analysis to study the effect of having some "selfish" nodes among "altruistic" nodes [6], or the effect of nodes reducing the transmission probability for all relay traffic (e.g. accepting or forwarding a packet with a probability $p < 1$) [7, 8]. Nevertheless, these works assume a mostly *uniform* behavior of relays when it comes to treating contacts with different nodes.

Contrary to the above approach, everyday experience suggests that people take into account the strength of their relation with a peer, when deciding whether to cooperate or not (*social selfishness*). As a result, a node A may be more willing to spend some energy (or take the risk) to forward a message of possible interest to an encountered node B, if A and B have strong *ties*, than if B is unknown to A. Furthermore, a long line of research has revealed that: (i) the strength of the "social" tie between two nodes (where "social" here may also be context-dependent) can often be reasonably predicted by the contact rate between them [9, 10], and (ii) the contact patterns and rates between mobile nodes exhibit significant amounts of heterogeneity [11, 12].

This opens up a very large space of possible cooperation policies, whose performance might be intimately related to the underlying mobility. E.g. a node might choose to forward (or accept) messages only to (from) nodes that it encounters frequently enough, or attempt to explore "weak ties" [13]. Alternatively, a node could instead "modulate" the forwarding probability as a function of the encounter rate with a given node. The following questions are then raised:

Q.1 Can we predict the performance of a routing mechanism, under a given cooperation policy, if we only know some basic statistics about the underlying heterogeneous mobility process?

Q.2 Can we improve performance by choosing the cooperation policy wisely, subject to a given constraint (e.g. power consumption rate for relay traffic)?

The former question is relevant, for example, when the policy is given (related to external, e.g. security factors). One then might like to know what kind of performance it should expect from the network, so as to choose the right protocol or protocol parameters, without knowing the global network topology, or to decide whether opportunistic networking is useful in this context or it is better to simply use the infrastructure. The latter question is relevant when we can assume that the average node is willing to contribute some fixed amount of resources (e.g. amount of power spent for relay traffic) towards participating in an opportunistic network, but we are interested in how to best use these resources to optimize network performance.

Our main contributions in this paper are

- We propose a generic model for *social selfishness* (or cooperation) related to mobility, which can capture a wide range of selfish behaviors and describe cooperation policies proposed in past literature (Section 2).

- Towards answering the first question, we use our model to provide closed-form expressions for the expected message delivery delay under a large class of mobility scenarios with heterogeneous contact rates; these expressions provide insights about the effect of the cooperation policy used and of the macroscopic mobility properties (mean value and variance of contact rates) (Section 3).

- Towards answering the second question, we examine the achievable performance-power consumption tradeoff *regions* under different cooperation policies. Specifically, we show that (i) when considering an interesting class of *Power-vs-Delay* tradeoffs, complex "social-based" policies cannot achieve better performance than the simple uniform policy, while (ii) when we consider *Power-vs-Delivery-Probability* tradeoffs, social cooperation policies can indeed be optimized (Section 4).

- Finally, we show that the intuition of our framework can be useful also in some real-world scenarios with significantly more complexity than the class of heterogeneous mobility models that we consider for our analysis.

2 Preliminaries

2.1 Mobility Model

We consider a network $\mathcal{N}$, with N nodes. We assume that the node transmission range is much smaller than the total network area, so that each pair of nodes can only communicate directly during the *contact events* of this pair (i.e. when the two nodes come into the transmission range of each other). We model this sequence of contact events for a pair of nodes $\{i, j\}$ by a *random point process*[1]. Specifically, we will perform our analysis assuming the following *class* of heterogeneous mobility models (or *Contact Networks*):

Definition 1 (Heterogeneous Contact Network)**.** *Inter-contact times between a given pair of nodes $\{i, j\}$ are exponentially distributed, with contact rate λ_{ij}, and independent of each other. The contact rates λ_{ij} for each pair $\{i, j\}$ are independently drawn from an arbitrary distribution $f_\lambda(\lambda)$, with finite mean μ_λ and variance σ_λ^2 (or coefficient of variation $cv_\lambda = \frac{\sigma_\lambda}{\mu_\lambda}$).*

With the choice of the above model we try to strike a tradeoff between realism and usability. We will now motivate our choices above in a bit more detail.

First, the assumption of independent and exponentially distributed inter-contact times (or equivalently Poisson contact processes) for each pair of nodes is needed to allow an exact analysis of performance metrics of interest using a Markovian framework. For this reason, it is a common assumption in most related works for epidemic spreading on Opportunistic Networks [3, 11, 14, 15, 16]. Furthermore, analyses of real-world traces, suggesting that the exponential distribution

[1]We will ignore the actual contact duration for simplicity and assume that contacts are instantaneous, since bandwidth concerns are orthogonal to the problem we consider here.

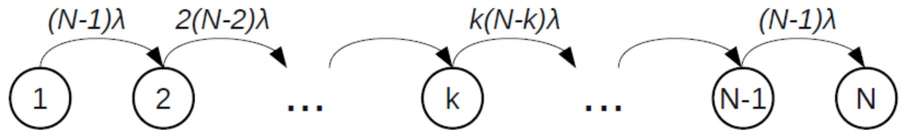

Figure 1: Epidemic spreading over a homogeneous network with N nodes

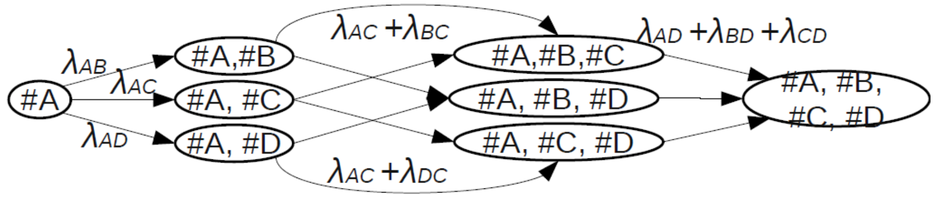

Figure 2: Epidemic spreading in a heterogeneous network with 4 nodes

can sometimes approximate the distribution of the inter-contact times [11, 12], or at least the tail of it [17, 18]. The assumption of independence (or even stationarity), while not that well supported by real traces, due to temporal or periodic characteristics in real mobility scenarios [19, 20, 21], is also necessary for analytical tractability (and any hope for closed form expressions). To our best knowledge, departing from the above assumptions (e.g. maintaining independence but allowing for pareto inter-contact time distributions [22, 23]), can only be used for asymptotic, convergence analysis about the message delivery delay of a routing protocol, i.e. if it achieves finite or infinite delay.

The second assumption is the heterogeneity of contact rates between different pairs of nodes. In previous analytical works, homogeneous mobility (i.e. $\lambda_{ij} = \lambda$ for each pair $\{i, j\}$) is often considered, because under this assumption closed-form results regarding the message delivery delay / probability can be found easily [3, 4]. For example, the message dissemination under epidemic routing can be modeled in the homogeneous case by a Markov Chain of N states as depicted in Fig. 1.

Unfortunately, study of real traces has provided strong evidence that contacts between different pairs of nodes are in fact largely heterogeneous, with some pairs never meeting each other and others meeting much more frequently [11, 12]. Similarly, Passarella et al. [10], shown, using data from real-world social networks, that (i) each person interacts and contacts its friends and acquaintances with higher rate as closer their relationship is, and (ii) the contact rates, between any individual and the other nodes, can be approximated by a distribution (which in our case is the distribution f_λ)[2].

This motivates us to depart from the homogeneous mobility model, and try to capture such heterogeneous rates. However, introducing different contact rates for

[2] Additionally, they found that the *Gamma* distribution matches well the observed distribution in real networks.

each pair of nodes complicates the problem. The message dissemination process depends on which certain set of nodes have the message at each state. Hence, the corresponding Markov Chain consists of 2^N states (see e.g. Fig. 2) and only numerical solutions (and only for N not large) [16] or upper bounds using rough spectral arguments [15], are allowed. Nevertheless, in [24], we have shown that analysis can also be extended for more complex cases, i.e. when contact rates are independently drawn from a distribution f_λ[3]. Although this mobility class cannot directly capture all types of macroscopic structure often observed in real-world networks (e.g. assortativity and community structure [25, 9]), we can use *any* valid probability density function f_λ to create an infinite range of random contact networks (in contrast to homogeneous models that correspond to only *one* function, i.e. $f_\lambda^{HOM}(\lambda) = \lambda_0 = const.$). Different functions lead to classes of contact processes with very different macroscopic characteristics. For example, an f_λ symmetric around μ_λ (e.g. uniform distribution) implies a balanced number of high and low rates, while a right-skewed f_λ (e.g. Pareto) describes a network with most pairs having large intercontact times, but few meeting very frequently. Furthermore, we can model networks with any level of sparsity by introducing in f_λ a finite probability mass at $\lambda = 0$ (i.e. with a delta function), modeling a chosen percentage of pairs that never meet.

For these reasons, we believe the above model strikes a good tradeoff, and as will see, allows us to explore the effect of different social-based cooperation policies and derive interesting insights, which is the main goal of this work. When possible, we will test these insights against real traces as well, to examine the extent to which departures from the above assumptions affect our conclusions.

2.2 Selfishness Models

The *store-carry-forward* mechanism requires from the relay nodes to (i) receive messages, (ii) store them, (iii) forward the messages they have to other relays and/or the destinations. As it is evident, this mechanism requires the cooperation of the relay nodes, and may put a heavy toll on their resources (bandwidth, storage space, battery life, etc.), dependent on the network traffic and protocol used. Furthermore, exchanging messages with unknown nodes may raise important security and privacy concerns. These considerations may render wireless nodes reasonably reluctant to relay traffic.

This unwillingness to cooperate might come in different flavors:

1. A node will not relay any traffic (*individual selfishness*).

2. A node will choose to relay each packet with some probability p. We will call this *uniform selfishness*.

3. A node will relay packets preferentially to other nodes it has a social relationship with (*social selfishness*).

---

[3]We will later provide some intuition about this result and we will modify it to consider node selfishness policies in Section 3.

The first case is an extreme case that could be handled with incentive or reputation mechanisms [26, 27, 28, 29, 30]. Such mechanisms are orthogonal (but possibly complementary) to our work. The second scenario has already been addressed in the past, with both theory and simulations, indicating that a low p can significantly hurt performance (e.g. [7]). The third case is closer, in our opinion, to human behavior. It is reasonable to assume that nodes are more willing to forward messages to or receive messages from nodes with whom they have a social tie. A social tie can be considered as a social relation in the real world (e.g. friendship), as a relation that originates from a routing mechanism (e.g. common interests in social-aware routing, SANE [31]), as a trust-relation that depends on how many times they have met in the past or they have collaborated (e.g. message exchanges or participation in a service composition) etc.

An important observation (for opportunistic networking) is that such social ties seem to be related with the mobility patterns. Studies from sociology [32] and social media [33] have shown that the stronger the social tie between two people is, the more they tend to meet or contact each other. Another study of Social Pervasive Networks [10], based on results from the anthropology field [34], shown that a relation between social ties and contact frequency (e.g. interaction on the respective social network) is supported in real networks. More recently, studies have directly suggested that the actual physical contact (related to mobility) can often serve as a good predictor for the strength of a social tie [35].

Combining the relations, we discussed above, between (i) *selfishness* and *social ties*, and (ii) *social ties* and *mobility patterns*, it is reasonable to assume a social selfishness model, where nodes decide to utilize a given contact opportunity with a probability $p_{ij} = p(\lambda_{ij})$, related to the contact frequency between the two nodes involved $\{i, j\}$. Such a model has been taken into account in a number of studies of routing protocols or message dissemination performance [13, 36, 37]. Some proposed strategies are, for example, to give more emphasis to ”strong ties” or ”weak ties” (i.e. large or small λ_{ij}): e.g. a node might decide to exchange messages only with the nodes it contacts more frequently, or the probability for message exchange to be linearly increasing with the contact rate of a pair of nodes, etc. To be able to capture most of the above selfishness behaviors (and more), in a simple and generic way, we choose to model this willingness to forward a message (essentially, the existence of related constraints affecting this willingness), in a *probabilistic* way.

Specifically, we propose two types of selfishness models, which correspond to typical behaviors that can appear in an Opportunistic Network scenario.

Definition 2. *[Selfishness: Type I] The probability for a message to be exchanged in a contact event between two nodes i and j, depends on their meeting rate λ_{ij} and is described by the relation:*

$$p_{ij} = p^{(I)}(\lambda_{ij}), \qquad p_{ij} \in [0, 1] \tag{1}$$

Definition 3. *[Selfishness: Type II] A pair of nodes i and j either can exchange messages in every contact event with probability p_{ij} or can never exchange messages with probability $1 - p_{ij}$. The probability p_{ij} depends on the meeting rate between these nodes, i.e. λ_{ij}, and is described by the relation:*

$$p_{ij} = p^{(II)}(\lambda_{ij}), \qquad p_{ij} \in [0, 1] \tag{2}$$

The probabilities for message exchange may depend, as described earlier, on various factors, e.g. willingness of the nodes, routing protocol mechanism, battery constraints, duration of the contact. The above two models allows to capture a number of such concerns. Furthermore, Type II selfishness is useful to capture situations where nodes decide a priori whether they will interact with a given node or not (e.g. due to security concerns), while Type I selfishness models situations where the contact probability might be modulated according, for example, to current battery level, content sensitivity, desire to control relay traffic, etc.

3 Message Delivery Delay

Having defined the types of node mobility and the types of node selfishness that we consider, we can now commence our analysis. Our goal is twofold:
1) To capture the combined effect of all nodes applying a given "selfishness" policy (or cooperation policy, to be less negative) on the performance of basic opportunistic routing protocols (e.g. epidemic routing, spray and wait, etc.).
2) To compare different cooperation behaviors and understand the impact of mobility properties on absolute and relative performance.

We state upfront that an *exact* analysis of random opportunistic routing protocols is already very challenging for heterogeneous mobility models (of the class of Def. 1), as explained earlier, and it becomes significantly more complex when social-selfishness policies are considered. While, in some cases, numerical analysis could be applied for small networks, it does not offer the kind of insights we are interested above. For this reason, we try instead to derive useful closed form approximations, that can be directly used for performance predictions as well as policy optimization.

3.1 Delay in Heterogeneous Networks

Before examining specific selfishness policies, we need first to consider how mobility heterogeneity affects the performance of simple opportunistic routing protocols. Let assume a Heterogeneous Contact Network $\mathcal{N}$ (Def. 1). Let also assume that the spreading follows the rules of a random routing protocol $\mathcal{P}$ (e.g. epidemic, direct transmission, 2-hop, Spray and Wait).

When the spreading process is at state k (i.e. k nodes have the message), we can use $\mathbf{C}_k$ to denote the set of nodes with a copy of the message[4]. To proceed to the next state $k+1$, a message needs to be exchanged between a node in $\mathbf{C}_k$ and a node among the ones without a copy. The *eligible* pairs $\{i,j\}$ of nodes whose contact will take the process to the next step depends on the protocol $\mathcal{P}$. Let us denote the set of these pairs as $\mathbf{S}_k$ and give two examples:

- Epidemic Routing: $\mathbf{S}_k = \{\{i,j\} \in \mathcal{N} : i \in \mathbf{C}_k, j \notin \mathbf{C}_k\}$ and the cardinality $\|\mathbf{S}_k\| = k(N-k)$.
- Spray and Wait (*SnW*) with k copies, during its wait phase: $\mathbf{S}_k = \{\{i,j\} \in \mathcal{N} : i \in \mathbf{C}_k, j = \text{destination}\}$ and the number of eligible pairs is $\|\mathbf{S}_k\| = k$.

Since (i) each inter-contact times of pairs $\{i,j\} \in \mathbf{S}_k$ are exponential and independent of others (i.e. Poisson processes), and (ii) the delay to move to the next state, $T_{k,k+1}$, is the time until the *first* pair $\{i,j\} \in \mathbf{S}_k$ comes in contact, it can be easily shown that $T_{k,k+1}$ will be also exponentially distributed with rate $\sum_{\{i,j\} \in \mathbf{S}_k} \lambda_{ij}$. Hence, its expectation will be:

$$E[T_{k,k+1}|\mathbf{C}_k] = \frac{1}{\sum_{\{i,j\} \in \mathbf{S}_k} \lambda_{ij}} \tag{3}$$

and using the conditional expectation properties, we get:

$$E[T_{k,k+1}] = \sum_{\mathbf{C}_k} \frac{1}{\sum_{\{i,j\} \in \mathbf{S}_k} \lambda_{ij}} \cdot P\{\mathbf{C}_k\} \tag{4}$$

It is evident that even the exact derivation of this delay is often not possible, as it involves complex combinatorics (i.e. calculating $P\{\mathbf{C}_k\}$). However, in [24] we have shown that the delay can be accurately approximated for the broad mobility class of Def. 1, as[5]:

Result 1. *The expectation of the spreading delay from state k to state $k+1$ for a Heterogeneous Contact Network can be approximated with a series expansion as*

$$E[T_{k,k+1}] = \frac{1}{M \cdot \mu_\lambda} \cdot \left(1 + \frac{\left(\frac{\sigma_\lambda}{\mu_\lambda}\right)^2}{M} + R\right) \tag{5}$$

where $M = \|\boldsymbol{S}_k\|$, and $R = \mathcal{O}(\frac{1}{M^2})$ corresponds to the impact of higher order terms.

[4] For example, in Fig. 2, the set $\mathbf{C}_2$ (i.e. $k=2$) could be one of the sets $\{\#A, \#B\}$ or $\{\#A, \#C\}$ or $\{\#A, \#D\}$.

[5] This result becomes exact only as we increase network size N, but provides a very accurate approximation for moderate network sizes. We refer the reader to [24] for the convergence conditions and details, and state here the basic result for completeness.

The above result corresponds to the first two terms of an approximation with a series expansion (often referred to as the *Delta* method [38]). Depending on the network size N and the variability of f_λ, fewer (just the first) or more terms could be used. We have observed that for moderate network sizes ($N > 100$), the first term approximation already offers reasonable accuracy.

3.2 Effect of Selfishness

It is easy to see that, when we introduce (social) selfishness, not all contacts resulting from the mobility model are useful in the spreading process, as was the case above. For instance, a node pair $\{i, j\}$ that meets with rate λ_{ij}, may exchange messages, on average, only half of the times (due to a Type I policy). Then, the *effective* (i.e. useful) contact rate will be $\lambda_{ij}^{'} = 0.5 \cdot \lambda_{ij}$.

The following lemmas, whose proofs can be found in Appendix, give the mean value and variance of the *effective* contact rates in networks with contact rate probability function f_λ (with μ_λ and σ_λ^2) and selfishness of Type I (Lemma 3.1) or Type II (Lemma 3.2).

Lemma 3.1. *The mean value, $\mu_\lambda^{(I)}$, and the variance, $\sigma_\lambda^{2(I)}$, of the effective contact rates in a network with contact rate probability function f_λ (μ_λ, σ_λ^2) and selfishness of Type I, are given by*

$$\mu_\lambda^{(I)} = E\left[\lambda \cdot p^{(I)}(\lambda)\right] \tag{6}$$

$$\sigma_\lambda^{2(I)} = E\left[\lambda^2 \cdot \left(p^{(I)}(\lambda)\right)^2\right] - \left(E\left[\lambda \cdot p^{(I)}(\lambda)\right]\right)^2 \tag{7}$$

where the expectations are taken over the p.d.f. f_λ.

Lemma 3.2. *The mean value, $\mu_\lambda^{(II)}$, and the variance, $\sigma_\lambda^{2(II)}$, of the effective contact rates in a network with contact rate probability function f_λ (μ_λ, σ_λ^2) and selfishness of Type II, are given by*

$$\mu_\lambda^{(II)} = E[\lambda \cdot p^{(II)}(\lambda)] \tag{8}$$

$$\sigma_\lambda^{2(II)} = E[\lambda^2 \cdot p^{(II)}(\lambda)] - \left(E[\lambda \cdot p^{(II)}(\lambda)]\right)^2 \tag{9}$$

where the expectations are taken over the p.d.f. f_λ.

Thus, when the network is characterised by social selfishness, we can use the above expressions in Result 1 to calculate the delay $E[T_{k,k+1}]$.

As discussed earlier, one can use one, two or more terms of Result 1 (and the respective moments, e.g. $\mu_\lambda^{(I)}$, $\sigma_\lambda^{2(I)}$, etc.) to increase accuracy. However, by including many terms, expressions get complex and it might be difficult to be used for optimization or to provide insights. Thus, without loss of generality and in order to simplify our discussion, in the remainder we will use the simplest first order approximation, i.e. $E[T_{k,k+1}] = \frac{1}{M \cdot \mu_\lambda^{(I)}}$ for Type I selfishness (similarly for Type II).

Having computed the delay $E[T_{k,k+1}]$, we can now use the linearity of expectation rule to calculate the expected message *delivery* delay under different random routing protocols.

Result 2. *The expected message delivery delay in an Heterogeneous Contact Network can be approximated by*

$$E[T_D] = \frac{c(N,L)}{\mu_\lambda^{eff.}}, \tag{10}$$

where $\mu_\lambda^{eff.}$ *is given by Eq.* (6) *or Eq.* (8) *for selfishness of Type I or Type II, respectively, and* $c(N,L)$ *is a constant dependent on the size of the network,* N*, the routing protocol* $\mathcal{P}$ *and the number of message copies,* L*. Values of* $c(N,L)$ *are given in Table 1 for three well-known routing protocols.*

In other words, as a first order approximation, the message delivery delay under random routing protocols is inversely proportional to the mean value of the effective contact rates in the network. Furthermore, the effect of Type I and Type II policies, with the same function $p(\lambda)$, turns out to be equal. We will thus not differentiate between the two policies, in the remainder, and simply refer to the mean effective contact rate as μ_λ^{eff}.

Finally, it is interesting to note that the effect of the mobility heterogeneity, in this first order approximation, when nodes are *not* selfish, affects performance only through its mean and not its variance (We have confirmed this to be the case for large N and non-heavy-tailed f_λ). In contrast, as we will show in the following sections, this is not the case when we introduce social selfishness in the spreading process.

Table 1: The values of $c(N,L)$ for three routing protocols.

Epidemic	$c(N,L) \approx \frac{\ln(N)}{N}$
2-hop	$c(N,L) = \sum_{k=1}^{N-1} \frac{k^2 \cdot (N-1)!}{(N-1)^{k+2} \cdot (N-k-1)!}$
SnW	$c(N,L) \leq \sum_{k=1}^{L-1} \frac{k^2 \cdot (N-1)!}{(N-1)^{k+2} \cdot (N-k-1)!} + \left(\frac{L}{N-1} + \frac{1}{L}\right) \frac{(N-1)!}{(N-1)^L \cdot (N-L-1)!}$

Table 2: Mean effective contact rate, $\mu_\lambda^{eff.} = E[\lambda \cdot p(\lambda)]$.

Policy A	Policy B	Policy C	Policy D
Gamma: $f_\lambda(x) = \frac{\beta^a}{\Gamma(\alpha)} \cdot x^{\alpha-1} \cdot e^{-\beta x}, \quad \alpha, \beta, x > 0$		$\left(\text{where} \quad \alpha = \frac{1}{cv_\lambda^2}, \; \beta = \frac{1}{\mu_\lambda \cdot cv_\lambda^2}\right)$	
$\mu_\lambda \cdot p_0$	$\mu_\lambda \cdot \left(p_2 + (p_1 - p_2)\frac{\gamma\left(1 + \frac{1}{cv_\lambda^2}, \frac{\lambda_0}{\mu_\lambda cv_\lambda^2}\right)}{\Gamma\left(1 + \frac{1}{cv_\lambda^2}\right)}\right)$	$\mu_\lambda \cdot p_1 \cdot \frac{\gamma\left(1 + \frac{1}{cv_\lambda^2}, \frac{\lambda_0}{\mu_\lambda cv_\lambda^2}\right)}{\Gamma\left(1 + \frac{1}{cv_\lambda^2}\right)} + p_2 \cdot \lambda_0 \cdot p_0$	$\mu_\lambda \cdot p_0 \cdot \left(1 - \frac{1}{(1 + m \cdot \mu_\lambda cv_\lambda^2)^{1+\frac{1}{cv_\lambda^2}}}\right)$
	where $\quad \gamma\left(\frac{1}{cv_\lambda^2}, \frac{\lambda_0}{\mu_\lambda cv_\lambda^2}\right) = \Gamma\left(\frac{1}{cv_\lambda^2}\right) \cdot (1 - p_0)$		
Exponential: $f_\lambda(x) = \frac{1}{\mu_\lambda} \cdot e^{-\frac{x}{\mu_\lambda}}, \quad \mu_\lambda, x > 0$			
$\mu_\lambda \cdot p_0$	$\mu_\lambda \cdot \left(p_1 + (p_2 - p_1) \cdot p_0 \cdot [1 - \ln(p_0)]\right)$	$\mu_\lambda \cdot \left(p_1 - p_0 \cdot \left(p_1 - (p_2 - p_1) \cdot \ln(p_0)\right)\right)$	$\mu_\lambda \cdot p_0 \cdot \left(1 - \frac{1}{(1 + m \cdot \mu_\lambda)^2}\right)$
Pareto: $\overline{F}_\lambda(x) = \left(\frac{\beta}{x + \beta}\right)^\alpha, \quad \alpha, \beta, x > 0$		$\left(\text{where} \quad \alpha = \frac{2 \cdot cv_\lambda^2}{cv_\lambda^2 - 1}, \; \beta = \mu_\lambda \cdot \frac{cv_\lambda^2 + 1}{cv_\lambda^2 - 1}\right)$	
$\mu_\lambda \cdot p_0$	$\mu_\lambda \cdot \left(p_1 + (p_2 - p_1) \cdot p_0 \cdot \left(1 - \alpha + \alpha \frac{1}{p_0^{\frac{1}{\alpha}}}\right)\right)$	$\mu_\lambda \cdot \left(p_1 \cdot \left(1 - p_0 \cdot \left(1 - \alpha + \alpha \frac{1}{p_0^{\frac{1}{\alpha}}}\right)\right) + p_2 \cdot p_0 \cdot \beta \cdot \left(\frac{1}{p_0^{\frac{1}{\alpha}}} - 1\right)\right)$	N.A.

3.3 Case Studies

With the basic performance result now in hand, we can go ahead and consider specific mobility processes, f_λ, and selfishness policies, $p(\lambda)$. To this end, we have analysed four policies (Table 3), which can represent a wide (and diverse) set of common behaviors for social selfishness and/or have been proposed before [13]. We will describe these policies only as Type I selfishness, but the analysis holds for the respective Type II policies as well.

Policy A *Uniform*: Each pair of nodes exchanges messages with probability p_0 every time they contact. The selfishness is not related with the contact rates between nodes.

Policy B *Strong / Weak ties*: Each pair of nodes exchanges messages with probability p_1 if they contact with rate less than λ_0 and with probability p_2 otherwise. The values of p_1 and p_2 determine the level of selfishness between pairs with strong and weak ties, respectively, while the value of λ_0 corresponds to the percentage of pairs that have strong (or weak) ties.

Policy C *Limit - Rates*: Each pair of nodes exchanges messages with probability p_1 if they contact with rate lower than λ_0, and adjust the message exchange probability if they contact with higher rate. Hence, for all pairs $\{i, j\}$ with $\lambda_{ij} > \lambda_0$, it will hold that $p(\lambda_{ij}) \cdot \lambda_{ij} = p_2 \cdot \lambda_0 = const.$.

Policy D *Exponential*: Each pair of nodes $\{i, j\}$ exchanges messages with probability $p_0 \cdot (1 - e^{-m \cdot \lambda_{ij}})$, where λ_{ij} is their meeting rate and $p_0 < 1$ and m are positive constants. The message exchange probability is higher for node pairs that meet more frequently.

Table 3: Selfishness policies.

Policy A	$p(\lambda) = p_0$	
Policy B	$p(\lambda) = \begin{cases} p_1 & : \lambda \leq \lambda_0 \\ p_2 & : \lambda > \lambda_0 \end{cases}$	$\overline{F}_\lambda(\lambda_0) = p_0$
Policy C	$p(\lambda) = \begin{cases} p_1 & : \lambda \leq \lambda_0 \\ p_2 \cdot \dfrac{\lambda_0}{\lambda} & : \lambda > \lambda_0 \end{cases}$	$\overline{F}_\lambda(\lambda_0) = p_0$
Policy D	$p(\lambda) = p_0 \cdot (1 - e^{-m \cdot \lambda})$	

To find the expected message delivery delay, for a certain network size and a certain routing protocol, only the computation of the effective contact rates' mean value ($\mu_\lambda^{eff.}$) is needed (Result 2). In Table 2, we present the closed form expressions for the $\mu_\lambda^{eff.}$ for these selfishness policies, under different mobility patterns. Specifically, we considered three cases for the contact rates distribution f_λ: (i) Gamma, (ii) Exponential[6], and (iii) Pareto distribution. We chose to analyze these

[6]The Exponential distribution can be defined also as a Gamma distribution with parameters $\alpha = 1$ and $\beta = \mu_\lambda^{-1}$. However, for clarity, we present the results separately.

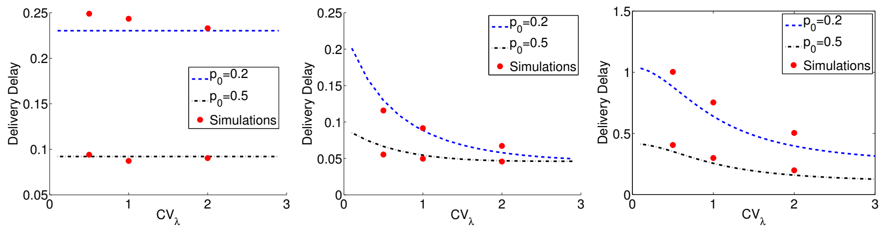

(a) Policy A (b) Policy B ($p_1 = 0, p_2 = 1$) (c) Policy D ($m = 0.25$)

Figure 3: *Delivery Delay* in networks with $N = 100$ nodes and varying *Mobility* characteristics ($\mu_\lambda = 1$ and $cv_\lambda \in [0, 3]$) for three different selfishness policies and under epidemic routing. The theoretical values of Delivery Delay for two parameters (p_0) for each selfishness policy are denoted with dashed lines and the corresponding simulations' average delivery delays are denoted with dots.

distributions, because they capture a large range of contact variabilities, and (especially Gamma) were shown to match well the observed contact rates distributions in real social networks [10].

Similar closed form expressions of $\mu_\lambda^{eff.}$, which depend only on the selfishness policy's parameters, $p(\lambda)$, and the first moments of the contact rates distributions, f_λ, can be found as well for other cases of $p(\lambda)$ and f_λ.

3.4 Validation

The results derived so far provide us with closed-form predictions for the performance of various protocols and selfishness behaviors under a broad class of mobility models. In Section 3.4.1, we first validate their accuracy in (synthetic) scenarios belonging to this mobility class, in order to isolate the effects of the various analytical approximations we have performed towards obtaining the expressions for these otherwise very complex problems. Then, in Section 3.4.2 we further consider trace-driven scenarios, where in addition to approximation errors, departures from many, if not most, of the model assumptions are expected to introduce further inaccuracies.

3.4.1 Synthetic Simulations

We developed a simulator that generates synthetic networks with mobility conforming to the mobility class of Def. 1: In each scenario, we assign to each pair $\{i, j\}$ a contact rate λ_{ij}, which we draw randomly from f_λ[7] and create a sequence of contact events (according to a Poisson process with rate λ_{ij}). We also assign to $\{i, j\}$ a probability p_{ij} according to the function $p(\lambda)$. Then, we simulate a large

[7]In the results we present, the contact rates are drawn from a Gamma distribution, $f_\lambda \sim Gamma$, with variable parameters μ_λ and cv_λ (see Fig. 3).

number of message exchanges, by choosing randomly for each message the source-destination pair, and calculate the mean simulated delivery delay by averaging the results.

In Fig. 3 we present, for networks with $N = 100$ nodes, how the mobility heterogeneity (i.e. $cv_\lambda = \frac{\sigma_\lambda}{\mu_\lambda}$) affects the message delivery delay, under different selfishness policies.

The theoretical results (dashed lines) show (Fig. 3(a)) that in the case of uniform selfishness (Policy A), the mobility heterogeneity (i.e. cv_λ) level does not affect the message delivery delay. For the same parameters of $p(\lambda)$ (i.e. p_0), the expected message delivery delay is equal for different mobility heterogeneity scenarios. However, for the non-uniform selfishness policies (Policies B and D), where the selfishness depends on the pairs' contact rates, mobility heterogeneity highly affects the message delivery delay (Fig. 3(b) and Fig. 3(c)). For the same parameters of $p(\lambda)$, the expected delivery delay decreases as the mobility heterogeneity level increases.

In all cases presented in Fig. 3, the synthetic simulations results (dots), are very close to our theoretical predictions, despite the various assumptions and approximations we used in our theoretical analysis. We have also performed simulations for larger networks (i.e. 300 and 1000 nodes), with similar findings.

3.4.2 Real-world Traces

In this section, we conduct simulations on the following sets of real mobility traces[8]:

Cabspotting: GPS coordinates from 536 taxi cabs collected over 30 days in San Francisco [39].

Infocom: Bluetooth sightings of 98 mobile and static nodes (*iMotes*) collected during Infocom 2006 [40].

Sigcomm: Bluetooth sightings of 76 mobile users of the MobiClique application at Sigcomm 2009 [41].

In Fig. 4 and Fig. 5 we show, for the *Cabspotting* and the *Infocom* trace, respectively, how the delivery delay of SnW routing decreases as the cooperation between nodes increases. Specifically, we present the relative delay decrease[9], $\frac{E[T_D]}{E[T_D^{max}]}$, i.e. the ratio of the average delivery delay in each scenario ($E[T_D]$) over the delay of the scenario with the highest level of selfishness ($E[T_D^{max}]$).

In Fig. 4(a) we simulated scenarios where nodes apply a Policy B selfishness (Table 3) with parameters $p_1 = 0, p_2 = 1$ (i.e. only "strong" ties). In each scenario different values of p_0 (i.e. percentage of pairs that cooperate) are selected; higher values of p_0 correspond to scenarios with less selfishness. Results of scenarios where nodes apply a Policy D selfishness are presented in Fig. 4(b). It can be

[8]Due to space limitations, we present here results only on the first two traces and we test our predictions on the Sigcomm trace in following sections.

[9]We present relative values in order to allow a direct comparison between the two traces, whose characteristics (network size, mobility statistics, etc.) differ significantly.

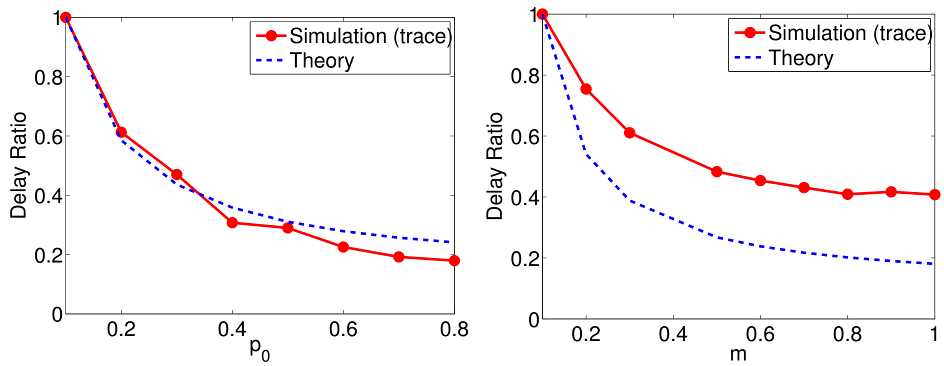

(a) Cabspotting, Policy B (b) Cabspotting, Policy D

Figure 4: Relative decrease of delay, $\frac{E[T_D]}{E[T_D^{max}]}$, of SnW routing, in scenarios with (a) Policy B ($p_1 = 0, p_2 = 1$ and variable p_0) selfishness and $L = 10$ copies, and (b) Policy D ($p_0 = 0.2$ and variable m) selfishness and $L = 20$ copies.

seen that for Policy B, the accuracy is significant, while for Policy D, the average simulated delivery delay (red line) decreases slower than predicted (dashed blue line). However, for both policies, the simulation results and theoretical predictions agree qualitatively, even if not always quantitatively.

In the *Infocom* trace (Fig. 5), the theoretical predictions are less accurate than in *Cabspotting*. The main reason for this, is that the mobility patterns of the Infocom trace deviate from the assumptions of our mobility model more than the mobility patterns of the Cabspotting trace. In particular, we observed higher community structure and temporal characteristics that cannot be captured by a Poisson contact process (i.e. during night, there are almost no contacts).

4 Performance vs Power Consumption

We have so far considered the effect of different selfishness policies on performance, assuming that the actual policy is *given* (e.g. user preferences, security or privacy concerns, etc.). However, it might be the case that a node's reluctance to always relay 3^{rd} party traffic stems from resource-related concerns (e.g. spending energy). In this case, the selfishness policy could be seen as a way for the node to control the amount of resources (e.g. transmission power) contributed to participate in the network.

Moreover, nodes would not object to use a different policy, e.g. one that improves the network-wide, and thus average node performance, if it wouldn't result in a higher expected resource consumption for them. For instance, if with a policy x and a policy y, a node consumes the same energy, but the message delivery probability achieved by policy x is higher than this of y, i.e. $P_x > P_y$, then it could choose to apply policy x in order to improve the overall network performance.

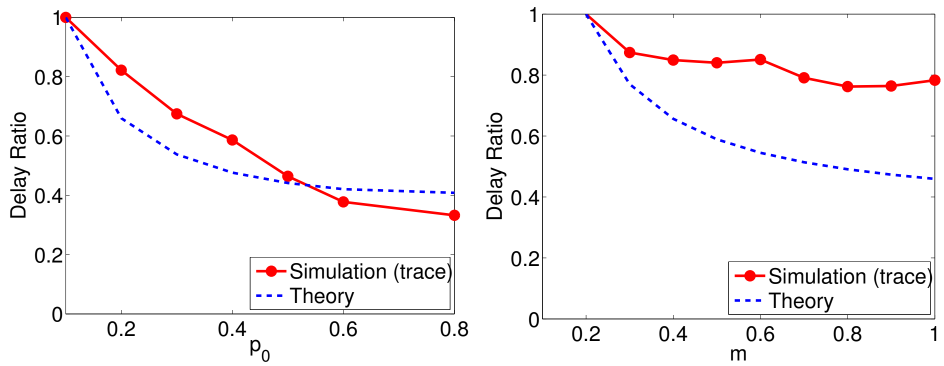

(a) Infocom, Policy B (b) Infocom, Policy D

Figure 5: Relative decrease of delay, $\frac{E[T_D]}{E[T_D^{max}]}$, of SnW routing, in scenarios with (a) Policy B ($p_1 = 0, p_2 = 1$ and variable p_0) selfishness and $L = 20$ copies, and (b) Policy D ($p_0 = 0.2$ and variable m) selfishness and $L = 20$ copies.

Table 4: Notation for the Communication Traffic Model

TTL	Message lifetime
N_f	Nb of flows
M	Window length (in nb of messages)
$E[N_t^{msg}]$	Avg nb of message transmissions per message and per node
T_∞	Observation time
N_m	Nb of generated messages during time interval $[0, T_\infty]$
$E[N_t]$	Avg nb of transmissions per node in the time interval $[0, T_\infty]$
E_t	Avg energy consumption of a message transmission
P	Avg power consumption

To this end, in this section, we examine the extent to which nodes could achieve different tradeoffs between resource consumption and network performance, using different policies, in two generic communication scenarios. At first, using a simple communication traffic injection model, we investigate the tradeoff between *Delivery Delay* and *Power Consumption* (Section 4.1). In the second case, we turn our attention to the possible *Delivery Probability - Power Consumption* tradeoffs that can be achieved by an opportunistic content sharing mechanism (Section 4.2).

4.1 Delivery Delay vs Power Consumption

As mobile devices rely on their batteries, whose energy capacity is limited, power consumption becomes a crucial issue. Nodes might prefer saving energy resources than consuming a significant amount of them for network operations (i.e. storing and relaying messages). Nevertheless, the total power consumption for relay traffic does not only depend on the policy choice, but also on the total message load in the network, and the protocol used. In order for a node to be able to estimate the expected power overhead of a given policy, we need to "level the ground", in a sense, and define a simple traffic model (see Table 4 for notation) that will allow us to compare directly the power overhead of different policies.

Let us assume that there are (on average) N_f number of flows in the network, i.e. N_f number of source-destination pairs that exchange messages (we assume that sources are "backlogged", i.e. always have messages to transmit). In order to ensure that source nodes do not insert new messages (input rate) faster than the network can deliver (output rate), some flow control mechanism is needed.

Some works suggest the use of an "out-of-band" channel (e.g. cellular network) for acknowledgements [42]. In this case, each source node could be forced, e.g. to not send a new message before the previous message is ACKed. In fact, we could also assume a window of M messages per flow that can go unacknowledged before a new message is send. Thus, if $E[T_D]$ is the expected delay of a message, the total load per flow is M messages per $E[T_D]$ time units (assuming an instant acknowledgement). If on the other hand, a slower "in-band" flow control is used, the RTT could also be expressed as $c \cdot E[T_D]$, $c > 1$.

Alternatively, each message can be assigned a *message lifetime* value, i.e. a TTL, after which the message cannot be forwarded or delivered to the destination, and nodes can drop any copies of it, in order to release valuable storage space[10]. To achieve a high message delivery probability (i.e. $PDR \approx 1$), the message lifetime must be set as

$$TTL = c_{TTL} \cdot E[T_D] \tag{11}$$

and c_{TTL} is large enough, such that the probability that the TTL expires before the message is delivered to the destination is small (this is necessary since we are interested in this section on the message delay).

Regardless of the exact flow control policy used (not of interest to this work), the above discussion suggests that a reasonable model for (stable) traffic loads is to assume that each source injects on average M new packets for each time interval $c \cdot E[T_D]$ (with c dependent on the flow control policy). In following, without loss of generality, we will assume a TTL flow control mechanism.

Under the condition of large $c = c_{TTL}$ (i.e. value of TTL such as $PDR \approx 1$), we can easily show that the average number of transmissions per message a node has to perform, $E[N_t^{msg}]$, depends only on the routing protocol, $\mathcal{P}$, and the network size, N, and is independent of the message delivery delay, i.e.:

$$E[N_t^{msg}] = c_t \tag{12}$$

where c_t is a constant dependent on N and $\mathcal{P}$. As an example, in *Spray and Wait* routing with L copies, as there are in total approximately L messages transmissions

[10]The lack or volatility of end-to-end paths in opportunistic networks, implies that the implementation of a transport protocol with feedback per packet (e.g. as ACK messages in TCP), as described above, might be either inefficient or infeasible. As a result, the TTL can often be used as an implicit flow control, allowing up to M new packets per TTL for each flow.

($L-1$ to relay nodes and 1 to destination node) before the expiry of the TTL, the average number of message transmissions per message and per node is given by $E[N_t^{msg}] = \frac{L}{N}$.

Hence, if we observe our network with communication traffic as described above, for a long time period T_∞, it follows easily that the expected number of generated messages is

$$N_m = N_f \cdot M \cdot \frac{T_\infty}{c \cdot E[T_D]}. \tag{13}$$

As the number of transmissions per generated message a nodes does is $E[N_t^{msg}]$, the total number of transmissions a node does in the time interval $[0, T_\infty]$ is

$$E[N_t] = N_m \cdot E[N_t^{msg}] = N_f \cdot M \cdot \frac{T_\infty}{c \cdot E[T_D]} \cdot c_t \tag{14}$$

where we substituted from Eq. (12) and Eq. (13)

Then, the power consumption rate can be calculated as

$$P = \frac{\text{Total Energy Consumption in } [0, T_\infty]}{T_\infty} = \frac{E_t \cdot E[N_t]}{T_\infty} \tag{15}$$

where E_t is the average energy for a single message transmission. The following result follows after substituting Eq. (14) in Eq. (15).

Result 3. *The average node power consumption is inversely proportional to the average message delivery delay and is given by*

$$P = c_p \cdot \frac{1}{E[T_D]} \tag{16}$$

where $c_p = \frac{E_t \cdot N_f \cdot M \cdot c_t}{c}$.

In Result 3, c_p is a constant that depends on the (i) network size N, (ii) the protocol used $\mathcal{P}$, (iii) the message size (E_t) and (iv) the traffic intensity (N_f, M). However, c_p is independent of the *selfishness policy* and the *mobility* of the nodes. Therefore, the main implication that comes of Result 3, is that:

Corollary 4.1. *In a Heterogeneous Contact Network, no matter how simple or sophisticated the selfishness policy used, the achievable power-delay operating regimes are exactly the same; In other words, whatever power-delay tradeoff can be achieved by some socially selfish policy, can also be achieved by the simple uniform policy.*

The above conclusion is somewhat surprising at first, given the range of strategies available under our social selfishness definition. However, we will try to shed some light on this counter intuitive result: Let assume a relay node i with some messages in its buffer. At the next contact event, i will forward each of the messages, e.g. to node j, with some probability, which depends on the protocol and the state of j (i.e. if j has the message or is the destination, etc.)[11]. It, then, follows that the more (effective) contact events a node has, the more messages it will transmit (i.e. power consumption). Since all nodes apply the same policy, the average number of contact events per time unit (and thus the power consumption) is the same for every node and $\propto E[p(\lambda) \cdot \lambda] \equiv \mu_{\lambda}^{eff.}$. Now, considering the discussion and results in Section 3, which show that the delivery delay is inversely proportional to $\mu_{\lambda}^{eff.}$, the relation suggested by Result 3 becomes evident.

To this direction, we can derive the following result (by simply combining Results 2 and 3) that relates the power consumption with the selfishness policy and mobility characteristics:

Result 4. *The average node power consumption in an Heterogeneous Contact Network is approximately given by*

$$P = \frac{c_p}{c(N,L)} \cdot \mu_{\lambda}^{eff.} \tag{17}$$

where $c(N,L)$ and c_p are defined in Results 2 and 3, respectively.

Thus, the expressions in Table 2 can be used to compute the average node power consumption, under the selfishness policies of Table 3.

4.1.1 Validation

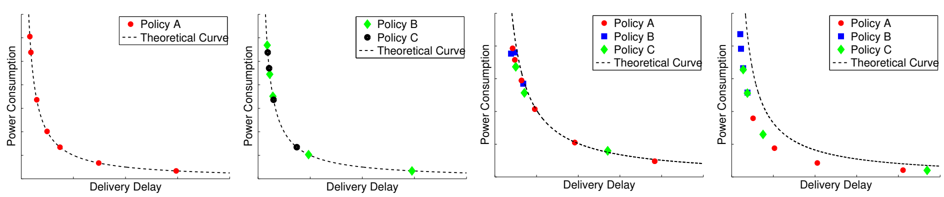

(a) Synthetic: uniform policy (b) Synthetic: non-uniform policies (c) Infocom (d) Sigcomm

Figure 6: Power consumption - message delivery delay trade off. Synthetic simulations with (a) uniform and (b) non-uniform selfishness policies. Simulations on the (c) Infocom and (d) Sigcomm real traces of both uniform and non-uniform selfishness policies scenarios.

From Result 3 we can see that the relation between power consumption and message delivery delay can be described by a reciprocal function or by a curve of the form $y = \frac{a}{x}$.

To investigate how accurate this prediction is, we first consider a heterogeneous mobility scenario ($f_{\lambda} \sim Gamma$, $\mu_{\lambda} = 1$, $cv_{\lambda} = 1$), consisting of 100 nodes. We

[11] Similarly, i will receive a message from j with some probability. Since we assume backlogged sources, the number of messages in the buffer of each node will be *on average* the same.

generate communication traffic between node pairs, according to the rules of the traffic model described in Section 4.1, and select SnW with $L = 10$ copies as the routing protocol. We perform Monte Carlo simulations. At first, we simulate scenarios with the uniform selfishness policy (Policy A) and choose values for the selfishness intensity (i.e. p_0) spanning the range $(0, 1)$, i.e. for minimum to maximum power consumption. Fig. 6(a) shows the simulation results for some sample values of p_0. It can be seen there that these exactly match our theoretical predictions.

We then simulate scenarios with different, non-uniform selfishness policies, in order to examine whether the delay-power curve is indeed the same or not. As is evident by Fig. 6(b), the simulated results for both non-uniform policies considered also coincide with the theoretical curve, which is also the delay-power curve for the uniform policy. In other words, by changing selfishness policies and their parameters, one can only achieve a shift on the theoretical curve.

To further examine the validity of this interesting finding, we test our predictions also in two real-world scenarios, the *Infocom* and *Sigcomm* traces. In Fig. 6(c), we use SnW routing with $L = 5$ copies in the Infocom trace and we create traffic conditions as described earlier. We measured the delivery delay of the messages and the power consumption of the nodes and plot the achievable delay-power tradeoff points for different policies. As it can be seen, our qualitative finding also holds here (i.e. all policies seem to have the same achievable region), and experimental values are quite close to the theoretically predicted curve. Similar observations can be made for the results of simulations on the Sigcomm trace (Fig. 6(d)). In this trace, although the theoretical curve seems to be a slightly displaced, it is clear that all policies also lie on the same tradeoff curve, as predicted.

4.2 Delivery Probability vs Power Consumption

In the previous section, we showed that the region of possible tradeoffs between *Delivery Delay* and *Power Consumption* is not affected by the selfishness policy. A key question arising then is: *is there not a way to achieve better performance-power tradeoff regions, e.g. compared to the uniform policy, by intelligently choosing the selfishness policy?*

In order to further explore this question, we turn our attention to another metric of high importance, namely the *delivery probability* of a message (or Probability Delivery Ratio, *PDR*). Thus, in this section, we investigate the *PDR - Power Consumption* tradeoffs using another example application, namely *content sharing* in opportunistic networks. The rationale behind this choice is twofold: first because content-centric applications have attracted increasing attention in both wired and wireless networks, and second to demonstrate the applicability of our framework to non end-to-end communication scenarios.

4.2.1 Opportunistic Content Sharing

In content sharing scenarios, new messages might be useful only for some fixed amount of time (e.g. related to the content nature), and interested nodes would like to *access* such messages before this time. We assume that there are $N_A(\leq N)$ nodes, in the network, that hold a content A for which another node i is interested in. This content can be data (e.g. a map, news, video, etc.) or even a service that these nodes can provide (e.g. Internet access or a computing service [43]). We also assume that this content can be only delivered directly when node i contacts any of the N_A nodes with the content, and not through relay nodes (this assumption might related to protocol complexity, but often comes very natural, as for example, when the content is an actual computing service the N_A providers can offer)[12].

The following result, whose proof can be found in C, gives the probability for a node i to successfully access content A by some time T.

Result 5. *In a Heterogeneous Contact Network with selfishness policy $p(\lambda)$, if N_A nodes hold a content A, then the probability for another node to access the content by a time T, is given by*

$$P_A\{T\} = 1 - \left(E\left[e^{-\lambda \cdot p(\lambda) \cdot T}\right]\right)^{N_A} \tag{18}$$

where the expectation is taken over f_λ.

Closed form expressions for the probability $P_A\{T\}$ under different selfishness policies (Table 3) and mobility patterns (f_λ) can be found in Table 5.

[12]Note that the selfishness policy applies even in this direct case, since e.g. content providers might not be equally willing to service or forward to any interested node.

Table 5: Probability a node to access the content by time T, $P_A\{T\}$.

Policy A	Policy B	Policy C
Gamma: $f_\lambda(x) = \frac{\beta^a}{\Gamma(\alpha)} \cdot x^{\alpha-1} \cdot e^{-\beta x}, \quad \alpha, \beta, x > 0$	$\left(\text{where} \quad \alpha = \frac{1}{cv_\lambda^2}, \;\; \beta = \frac{1}{\mu_\lambda \cdot cv_\lambda^2}\right)$	
$1 - \left(1 + p_0 \cdot T \cdot \mu_\lambda \cdot cv_\lambda^2\right)^{-\frac{N_A}{cv_\lambda^2}}$	$1 - \left(\frac{\gamma\left(\frac{1}{cv_\lambda^2}, \left(p_1 \cdot T + \frac{1}{\mu_\lambda cv_\lambda^2}\right) \cdot \lambda_0\right)}{\Gamma(\frac{1}{cv_\lambda^2}) \cdot \left(1 + p_1 \cdot T \cdot \mu_\lambda \cdot cv_\lambda^2\right)^{\frac{1}{cv_\lambda^2}}} + \frac{\Gamma(\frac{1}{cv_\lambda^2}) - \gamma\left(\frac{1}{cv_\lambda^2}, \left(p_2 \cdot T + \frac{1}{\mu_\lambda cv_\lambda^2}\right) \cdot \lambda_0\right)}{\Gamma(\frac{1}{cv_\lambda^2}) \cdot \left(1 + p_2 \cdot T \cdot \mu_\lambda \cdot cv_\lambda^2\right)^{\frac{1}{cv_\lambda^2}}} \right)^{N_A}$	$1 - \left(\frac{\gamma\left(\frac{1}{cv_\lambda^2}, \left(p_1 \cdot T + \frac{1}{\mu_\lambda cv_\lambda^2}\right) \cdot \lambda_0\right)}{\Gamma(\frac{1}{cv_\lambda^2}) \cdot \left(1 + p_1 \cdot T \cdot \mu_\lambda \cdot cv_\lambda^2\right)^{\frac{1}{cv_\lambda^2}}} + e^{-p_2 \lambda_0 \cdot T} \cdot p_0 \right)^{N_A}$
	where $\quad \gamma\left(\frac{1}{cv_\lambda^2}, \frac{\lambda_0}{\mu_\lambda cv_\lambda^2}\right) = \Gamma\left(\frac{1}{cv_\lambda^2}\right) \cdot (1 - p_0)$	
Exponential: $f_\lambda(x) = \frac{1}{\mu_\lambda} \cdot e^{-\frac{x}{\mu_\lambda}}, \quad \mu_\lambda, x > 0$		
$1 - (1 + p_0 \cdot T \cdot \mu_\lambda)^{-N_A}$	$1 - \left(\frac{1 - p_0^{1 + p_1 \cdot T \cdot \mu_\lambda}}{1 + p_1 \cdot T \cdot \mu_\lambda} + \frac{p_0^{1 + p_2 \cdot T \cdot \mu_\lambda}}{1 + p_2 \cdot T \cdot \mu_\lambda} \right)^{N_A}$	$1 - \left(\frac{1 - p_0^{1 + p_1 \cdot T \cdot \mu_\lambda}}{1 + p_1 \cdot T \cdot \mu_\lambda} + p_0^{1 + p_2 \cdot T \cdot \mu_\lambda} \right)^{N_A}$

We know from Result 4 that the average power consumption is proportional to $\mu_\lambda^{eff.}$. However, the expression for the content delivery probability (Result 5) relates to the mobility pattern and the selfishness policy in a non-linear way, that is also more complex than the case of delay. The first observation is that it's not easy to deduce a simple relation between the power consumption and the PDR, under generic mobility and selfishness characteristics, as was the case for power and delay (Result 3). The non-linearity also implies that it might now be possible indeed to change (and ultimately improve) the achievable power - performance (PDR) region.

4.2.2 Evaluation

To obtain some useful evidence, we will focus here on two selfishness policies, namely the *uniform* policy (Policy A) and the *limit-rates* policy (Policy C). Our choice for the specific non-uniform selfishness policy is based on the fact that it was proposed in [13] as a policy designed for a content dissemination application, which resembles the application we study.

From its definition (Table 3), we can see that Policy C limits the average number of effective contacts for pairs that contact more frequently than a certain threshold. The intuition behind this mechanism, is that a node i avoids communicating every time with the nodes j with whom it meets frequently, because (i) each effective contact incurs some energy consumption, and (ii) as they meet frequently, the probability node j to hold a content message in which node i is interested in and which did not exist in the memory of j their previous contact event, is small. Thus, limiting the effective contact events with frequently met nodes would result in a better PDR-power tradeoff.

Our theoretical predictions among with simulated results from two scenarios where we assign a content to random nodes and measured the delivery probability of it to a certain node, are presented in Figures 7 and 8. These confirm the intuition about the superiority of Policy C regarding content sharing applications. In Fig. 7(a) we present the PDR values in scenarios with uniform and rate-limit selfishness policies, where only one node holds the content message. As it can be seen, Policy C achieves always higher PDR than policy A for the same power consumption values. Specifically, Fig. 7(b) shows the improvement (i.e. the ratio $\frac{PDR_C - PDR_A}{PDR_A}$) in PDR we achieve with Policy C, which, for some values of power consumption, is almost 30%. In some other scenarios we simulated, this improvement was even up to 70%.

Fig. 8 present the comparison of the two policies, in a scenario with more heterogeneous mobility ($cv_\lambda = 2$) where $M = 5$ nodes hold the content. The observations about the performance of the two policies remain the same.

Finally, it is evident by both Fig. 7 and 8 that simulations results for the synthetic heterogeneous model (red dots) match our theoretical predictions very well.

As a final step, we test again the accuracy of our findings in two real-networks, the *Sigcomm* and *Infocom* traces. We simulated scenarios with different number

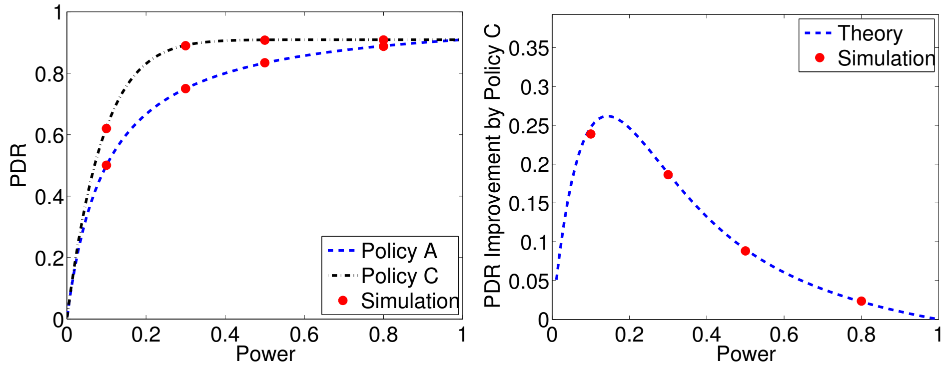

(a) (b)

Figure 7: (a) Probability Delivery Ratio of a content of policy A selfishness (blue) and policy C selfishness (black) for different power consumption levels. (b) Relative difference of the Probability Delivery Ratio between Policy C and Policy A selfishness, i.e. $\frac{PDR_C - PDR_A}{PDR_A}$. Number of content copies $M = 1$ and $T = \frac{10}{\mu_\lambda}$. Mobility characteristics: $\mu_\lambda = 1$, $cv_\lambda = 1$.

of content holders and for the same selfishness policies as before. The results are presented in Fig. 9 and compared to the theoretical prediction. As it can be observed, while the absolute values do not match exactly, Policy C again outperforms the uniform policy, and the relative performance improvement follow the shape of the theoretical curve quite well (this is very important when considering finding optimal operating points, using the theoretical curve).

Hence, we can conclude that our model can provide quite accurate predictions, even for real network scenarios. Finally, it is clear that, unlike the case of delay-power tradeoff, using social selfishness wisely can improve performance here, and our model could be used in order to predict the relative performance of different policies and, consequently, for policy optimization.

5 Related Work

The feasibility of communication over an opportunistic network highly depends on the willingness of nodes to cooperate. To this end, many techniques and protocols were proposed in order to motivate nodes to act as relays for messages that are not generated by or destined to them [26, 27, 28, 29, 30]. In [26] a reputation mechanism is used to encourage nodes to cooperate in order (i) to be able to receive the messages destined to them, and (ii) the other nodes to offer them their services (i.e. relay their messages). Another approach, which results in growing incentives to nodes for acting as relays, is followed in [29], where each node i is willing to forward the messages of another node j, according to the number of messages node

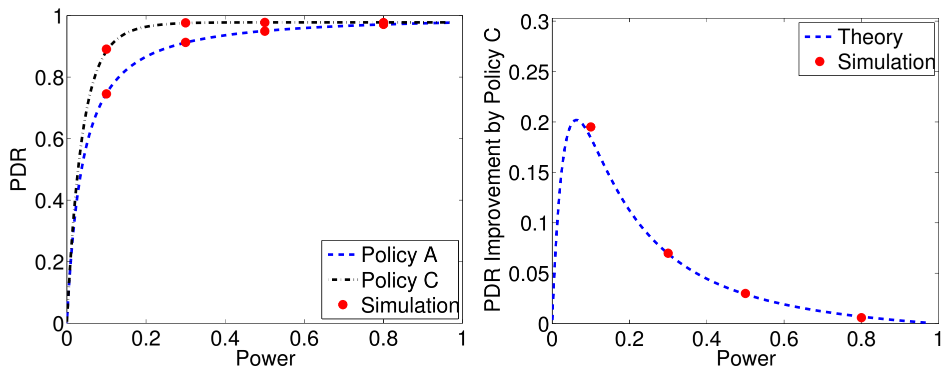

(a) (b)

Figure 8: (a) Probability Delivery Ratio of a content of policy A selfishness (blue) and policy C selfishness (black) for different power consumption levels. (b) Relative difference of the Probability Delivery Ratio between Policy C and Policy A selfishness, i.e. $\frac{PDR_C - PDR_A}{PDR_A}$. Number of content copies $M = 5$ and $T = \frac{5}{\mu_\lambda}$. Mobility characteristics: $\mu_\lambda = 1$, $cv_\lambda = 2$.

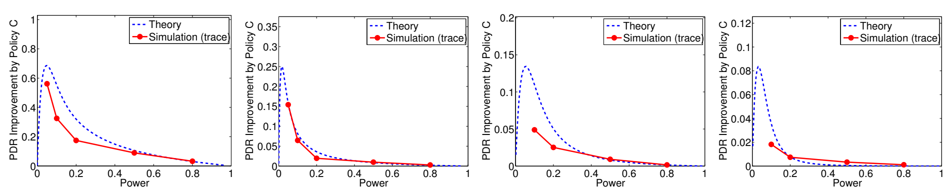

(a) Sigcomm, $M = 3$, $T = \frac{20}{\mu_\lambda}$ (b) Sigcomm, $M = 5$, $T = \frac{20}{\mu_\lambda}$ (c) Infocom, $M = 3$, $T = \frac{10}{\mu_\lambda}$ (d) Infocom, $M = 5$, $T = \frac{10}{\mu_\lambda}$

Figure 9: Relative difference of the Probability Delivery Ratio between Policy C and Policy A selfishness, i.e. $\frac{PDR_C - PDR_A}{PDR_A}$, in the *Sigcomm trace* with (a) $M = 3$, (b) $M = 5$ number of copies and $T = 20/\mu_\lambda$, and in *Infocom trace* with (c) $M = 3$, (d) $M = 5$ number of copies and $T = 10/\mu_\lambda$.

j has forwarded already for it. Finally, credit-based mechanisms are presented in [27] and [28], as well as barter-based mechanisms in [30].

Furthermore, many analytical and simulation-based studies investigate the effects of node selfishness on communication performance [7, 44, 45, 8, 6]. In [7] authors investigate, through simulations, how the performance of epidemic schemes is affected when the network comprises of non-cooperative nodes. They consider two kinds of selfishness: in each contact event either nodes are unwilling to copy a message with probability p_{nc} or they are unwilling to forward it with probability p_{nf}. For a similar scenario, Karaliopoulos [44] models probabilistic selfish behavior of nodes in homogeneous networks (i.e. constant contact rate λ for every pair of nodes). In [45] authors extend the work of [44] in terms of multicast applications. The authors of [8] model selfishness in a different scheme, where each node transmits its own message (operates as a source) with probability p and transmits one of the messages it has as a relay with probability 1-p. They assume that only one message can be exchanged per contact event and use only 2-hop routing. Another approach of selfishness is tackled in [6], where authors propose a selfishness model where each node is either selfish or altruist (modeled as a probability p_i for each node i) regarding all its contacts (i.e. i shows the same selfishness for every other node j, $p_{ij} = p_i$) and investigate through simulations the effect on the communication throughput.

The above protocols and studies, assume (under different models) that every node is either totally selfish or not. However, the assumption that users are selfish and are not willing to forward packets for anyone else, might not always hold. In this direction, the notion of *social selfishness* appears. In social selfishness, the nodes might be selfish only regarding some other nodes with which they have a weak (or even a strong) social relation ("tie") [36, 13, 37]. In [36] authors use a model of a network with two communities and introduce the notion of selfishness that depends on the contact rate between nodes. For nodes with high contact rate (e.g. within the same community) the selfishness is characterised by the probability p_i and for nodes with low contact rate another value p_o for selfishness is considered. They build a Markov Chain and investigate the effect on the performance through simulations. In [13] the authors investigate the role of the "weak ties" (i.e. pairs of nodes that contact infrequently, which in our case means the pairs of nodes with small contact rate λ_{ij}) in a content updating/dissemination scenario. Finally, in [37] a routing protocol, designed for networks where nodes have social selfishness behaviors, is proposed.

Our work, being the first to provide a theoretical framework and analytical closed-form results, complements previous studies on the effect of social selfishness on communication performance, which are limited to evaluation through simulations [13, 37] or analytical modeling of specific cases [36]. Moreover, not only the heterogeneous mobility model we consider can capture much wider range of scenarios than the models used in previous analytical studies [44, 8, 36], but also our results were shown to capture (either qualitatively or quantitatively) the much more complex characteristics of real-networks' mobility.

6 Conclusion

In this paper, we analysed the effect of social selfishness on opportunistic communications. At first, we proposed a framework for modeling heterogeneous mobility, which further allowed us to build a generic model that can describe a wide range of common social selfishness behaviors (related to privacy concerns, ressources consumption, etc.). Based on our mobility / selfishness framework, we derived closed form results for predicting the message delivery delay in a network with (socially) selfish nodes. Furthermore, we investigated how selfishness affects the performance - power consumption tradeoffs in a network, under two communication scenarios. We derived results that show if and when it is possible to optimize a selfishness policy in order to achieve better tradeoffs.

Due to the lack of existing solutions fighting social selfishness, we deem as essential to have an analytical framework for it and predict the performance degradation it causes on message dissemination, which as shown depends on various factors (selfishness behaviors and mobility). We believe that our work can be a useful tool for the design of novel protocols and applications for socially selfish environments.

A Proof of Lemma 3.1

Proof. As defined in Def. 1, the contact process for a pair $\{i, j\}$ is a Poisson process with rate λ_{ij}. Thus, if, according to Def. 2, in each of the contact events a message can be exchanged with probability p_{ij} (independently of what happened in the previous or following contact events), then the *effective contact events* are described by another Poisson process, which results after thinning the initial contact process. The rate of the new, thinned, Poisson process is then

$$\lambda_{ij}^{(I)} = \lambda_{ij} \cdot p_{ij} = \lambda_{ij} \cdot p^{(I)}(\lambda_{ij})$$

Hence, the mean value of the rate of the effective contact events, is given by

$$\begin{aligned}\mu_\lambda^{(I)} = E\left[\lambda^{(I)}\right] &= \int_0^\infty E\left[\lambda^{(I)}|\lambda_{ij} = x\right] \cdot f_\lambda(x) \cdot dx \\ &= \int_0^\infty \left(x \cdot p^{(I)}(x)\right) \cdot f_\lambda(x) \cdot dx \;\; = \;\; E\left[\lambda \cdot p^{(I)}(\lambda)\right]\end{aligned}$$

Similarly the second moment, is given by

$$\begin{aligned}E\left[\left(\lambda^{(I)}\right)^2\right] &= \int_0^\infty E\left[\left(\lambda^{(I)}\right)^2 |\lambda_{ij} = x\right] \cdot f_\lambda(x) \cdot dx \\ &= \int_0^\infty \left(x \cdot p^{(I)}(x)\right)^2 \cdot f_\lambda(x) \cdot dx \;\; = \;\; E\left[\lambda^2 \cdot \left(p^{(I)}(\lambda)\right)^2\right]\end{aligned}$$

and, finally, the variance can be computed as:

$$\sigma_\lambda^{2(I)} = E\left[\left(\lambda^{(I)}\right)^2\right] - \left(\mu_\lambda^{(I)}\right)^2 = E\left[\lambda^2 \cdot \left(p^{(I)}(\lambda)\right)^2\right] - \left(E\left[\lambda \cdot p^{(I)}(\lambda)\right]\right)^2$$

$\square$

B Proof of Lemma 3.2

Proof. According to Def. 3, a pair of nodes $\{i, j\}$ that contacts with rate λ_{ij}, either can always exchange a message during a contact event, with probability $p_{ij} = p^{(II)}(\lambda_{ij})$, or never exchanges messages during its contact events, with probability $1-p_{ij}$. The equivalent of this constraint mechanism, is a network where some pairs of nodes contact with their initial rate, i.e. $\lambda_{ij}^{(II)} = \lambda_{ij}$, and some never contact, i.e. $\lambda_{ij}^{(II)} = 0$.

Thus, we can compute the mean value of the *effective* contact events as following:

$$\begin{aligned}
\mu_{\lambda}^{(II)} &= \int_0^{\infty} E\left[\lambda^{(II)} | \lambda_{ij} = x\right] \cdot f_{\lambda}(x) \cdot dx \\
&= \int_0^{\infty} \left(x \cdot p^{(II)}(x) + 0 \cdot (1 - p^{(II)}(x))\right) \cdot f_{\lambda}(x) \cdot dx \\
&= \int_0^{\infty} \left(x \cdot p^{(II)}(x)\right) \cdot f_{\lambda}(x) \cdot dx = E[\lambda \cdot p^{(II)}(x)]
\end{aligned}$$

Similarly,

$$\begin{aligned}
E\left[\left(\lambda^{(II)}\right)^2\right] &= \int_0^{\infty} E\left[\left(\lambda^{(II)}\right)^2 | \lambda_{ij} = x\right] \cdot f_{\lambda}(x) \cdot dx \\
&= \int_0^{\infty} \left(x^2 \cdot p^{(II)}(x) + 0^2 \cdot (1 - p^{(II)}(x))\right) \cdot f_{\lambda}(x) \cdot dx \\
&= \int_0^{\infty} \left(x^2 \cdot p^{(II)}(x)\right) \cdot f_{\lambda}(x) \cdot dx = E[\lambda^2 \cdot p^{(II)}(x)]
\end{aligned}$$

and finally

$$\sigma_{\lambda}^{2(II)} = E\left[\left(\lambda^{(II)}\right)^2\right] - \left(\mu_{\lambda}^{(II)}\right)^2 = E\left[\lambda^2 \cdot p^{(II)}(\lambda)\right] - \left(E\left[\lambda \cdot p^{(II)}(\lambda)\right]\right)^2$$

$\square$

C Proof of Result 5

Proof. Let us denote $P_a\{j, T\}$ the probability the node i to contact a node $j \in [1, ..., N_A]$ and exchange messages with it (i.e. effective contact event) before a certain time T. Obviously $P_a\{j, T\}$ (i) depends on the contact rate λ_{ij} and the selfishness policy $p(\lambda)$, and (ii) as the inter-contact intervals are exponentially distributed it is given by[13]

$$P_a\{j, T | \lambda_{ij}, p(\lambda_{ij})\} = 1 - e^{-\lambda_{ij} \cdot p(\lambda_{ij}) \cdot T}$$

[13] The CDF of an exponential distribution with rate λ is given by $F(x) = 1 - e^{-\lambda \cdot x}$.

Since the probability a node to have the content is the same for all nodes, we can write

$$P_a\{j,T\} = \int_0^{\infty} P_a\{j,T|\lambda_{ij},p(\lambda_{ij})\} \cdot f_\lambda(x)dx$$
$$= \int_0^{\infty} \left(1-e^{-\lambda_{ij}\cdot p(\lambda_{ij})\cdot T}\right) \cdot f_\lambda(x)dx \;\; = \;\; 1-E\left[e^{-\lambda\cdot p(\lambda)\cdot T}\right]$$

where the expectation in is taken over f_λ.

Node i will not access the content by time T, only if it does not contact any of the N_A nodes. Hence, we can write for the probability that i will get the content by time T:

$$P_A\{T\} \; = \; 1-\overline{P}_A\{T\} \; = \; 1-\prod_{j=1}^{N_A}\overline{P}_a\{j,T\} \;\; = \;\; 1-\prod_{j=1}^{N_A}(1-P_a\{j,T\})$$

where $\overline{P}$ denotes the probability of the complementary event. Now, combining the above two equations and the fact that the nodes j with the content (and the respective contact rates λ_{ij}) are independent, it follows

$$P_A\{T\} = 1-\prod_{j=1}^{N_A}\left(1-\left(1-E\left[e^{-\lambda\cdot p(\lambda)\cdot T}\right]\right)\right)$$
$$= 1-\prod_{j=1}^{N_A}E\left[e^{-\lambda\cdot p(\lambda)\cdot T}\right] \;\; = \;\; 1-\left(E\left[e^{-\lambda\cdot p(\lambda)\cdot T}\right]\right)^{N_A}$$

□

References

[1] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: the single-copy case," *IEEE/ACM Trans. Netw.*, vol. 16, pp. 63–76, Feb. 2008.

[2] A. Vahdat and D. Becker, "Epidemic routing for partially connected ad hoc networks," Tech. Rep. CS-200006, Duke University, 2000.

[3] R. Groenevelt, P. Nain, and G. Koole, "The message delay in mobile ad hoc networks," *Performance Evaluation*, vol. 62, pp. 210–228, 2005.

[4] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Efficient routing in intermittently connected mobile networks: the multiple-copy case," *IEEE/ACM Trans. Netw.*, vol. 16, pp. 77–90, Feb. 2008.

[5] Y. Zhu, B. Xu, X. Shi, and Y. Wang, "A survey of social-based routing in delay tolerant networks: Positive and negative social effects," *Communications Surveys Tutorials, IEEE*, vol. PP, no. 99, pp. 1–15, 2012.

[6] 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*, 2009.

[7] A. Panagakis, A. Vaios, and L. Stavrakakis, "On the effects of cooperation in DTNs," in *Proc. IEEE COMSWARE*, IEEE, 2007.

[8] J. Liu, X. Jiang, H. Nishiyama, and N. Kato, "Performance modeling for two-hop relay with node selfishness in delay tolerant networks," in *Proc. IEEE GreenCom*, 2011.

[9] T. Hossmann, T. Spyropoulos, and F. Legendre, "Putting contacts into context: mobility modeling beyond inter-contact times," in *Proc. ACM MobiHoc*, 2011.

[10] A. Passarella, R. I. Dunbar, M. Conti, and F. Pezzoni, "Ego network models for future internet social networking environments," *Computer Communications*, vol. 35, no. 18, 2012.

[11] W. Gao, Q. Li, B. Zhao, and G. Cao, "Multicasting in delay tolerant networks: a social network perspective," in *Proc. of ACM MobiHoc*, 2009.

[12] V. Conan, J. Leguay, and T. Friedman, "Characterizing pairwise inter-contact patterns in delay tolerant networks," in *Proc. ACM Autonomics*, 2007.

[13] S. Ioannidis and A. Chaintreau, "On the Strength of Weak Ties in Mobile Social Networks," in *Proc. ACM SNS Workshop*, 2009.

[14] C.-H. Lee and D. Y. Eun, "On the forwarding performance under heterogeneous contact dynamics in mobile opportunistic networks," *IEEE Transactions on Mobile Computing*, vol. 12, no. 6, 2013.

[15] A. Picu, T. Spyropoulos, and T. Hossmann, "An analysis of the information spreading delay in heterogeneous mobility dtns," in *Proc. IEEE WoWMoM*, 2012.

[16] Y. Kim, K. Lee, N. B. Shroff, and I. Rhee, "Providing probabilistic guarantees on the time of information spread in opportunistic networks," in *IEEE INFOCOM*, 2013.

[17] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnović, "Power law and exponential decay of inter contact times between mobile devices," in *Proc. ACM MobiCom*, 2007.

[18] H. Cai and D. Y. Eun, "Crossing over the bounded domain: from exponential to power-law intermeeting time in mobile ad hoc networks," *IEEE/ACM Trans. Netw.*, vol. 17, pp. 1578–1591, Oct. 2009.

[19] A.-K. Pietilänen and C. Diot, "Dissemination in opportunistic social networks: the role of temporal communities," in *Proc. ACM MobiHoc*, 2012.

[20] W. Gao, G. Cao, T. La Porta, and J. Han, "On exploiting transient social contact patterns for data forwarding in delay-tolerant networks," *IEEE Trans. on Mob. Computing*, vol. 12, pp. 151–165, jan. 2013.

[21] G. Zyba, G. Voelker, S. Ioannidis, and C. Diot, "Dissemination in opportunistic mobile ad-hoc networks: The power of the crowd," in *Proc. IEEE INFOCOM*, 2011.

[22] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on the design of opportunistic forwarding algorithms," in *Proc. IEEE INFOCOM*, 2006.

[23] C. Boldrini, M. Conti, and A. Passarella, "Less is more: Long paths do not help the convergence of social-oblivious forwarding in opportunistic networks," in *Proc. ACM MobiOpp*, 2012.

[24] P. Sermpezis and T. Spyropoulos, "Spreading the word: Delay analysis of epidemic schemes in heterogeneous contact environments," tech. rep., Eurecom, 2013. `http://www.eurecom.fr/~sermpezi/TechRepSpreadWord.pdf`.

[25] M. Newman, *Networks: An Introduction*. New York, NY, USA: Oxford University Press, Inc., 2010.

[26] A. Mei and J. Stefa, "Give2get: Forwarding in social mobile wireless networks of selfish individuals," *IEEE Trans. on Dependable and Secure Computing,*, vol. 9, pp. 569–582, july-aug. 2012.

[27] H. Zhu, X. Lin, R. Lu, Y. Fan, and X. Shen, "Smart: A secure multilayer credit-based incentive scheme for delay-tolerant networks," *IEEE Trans. on Vehicular Technology*, vol. 58, pp. 4628–4639, oct. 2009.

[28] B. B. Chen and M. C. Chan, "Mobicent: a credit-based incentive system for disruption tolerant network," in *Proc. IEEE INFOCOM*, 2010.

[29] U. Shevade, H. H. Song, L. Qiu, and Y. Zhang, "Incentive-aware routing in dtns," in *Proc. IEEE ICNP*, 2008.

[30] A. Krifa, C. Barakat, and T. Spyropoulos, "Mobitrade: trading content in disruption tolerant networks," in *Proc. ACM CHANTS*, 2011.

[31] A. Mei, G. Morabito, P. Santi, and J. Stefa, "Social-aware stateless forwarding in pocket switched networks," in *Proc. IEEE INFOCOM*, 2011.

[32] M. S. Granovetter, "The Strength of Weak Ties," *The American Journal of Sociology*, vol. 78, no. 6, pp. 1360–1380, 1973.

[33] E. Gilbert and K. Karahalios, “Predicting tie strength with social media,” in *Proc. CHI*, 2009.

[34] W. X. Zhou, D. Sornette, R. A. Hill, and R. I. M. Dunbar, “Discrete hierarchical organization of social group sizes,” *Proc. Royal Society B: Biological Sciences*, vol. 272, pp. 439–444, Feb. 2005.

[35] T. Hossmann, G. Nomikos, T. Spyropoulos, and F. Legendre, “Collection and analysis of multi-dimensional network data for opportunistic networking research,” *Computer Communications*, vol. 35, no. 13, pp. 1613–1625, 2012.

[36] Y. Li, P. Hui, D. Jin, L. Su, and L. Zeng, “Evaluating the impact of social selfishness on the epidemic routing in delay tolerant networks,” *Comm. Letters.*, vol. 14, Nov. 2010.

[37] Q. Li, S. Zhu, and G. Cao, “Routing in socially selfish delay tolerant networks,” in *Proc. IEEE INFOCOM*, 2010.

[38] G. W. Oehlert, “A note on the delta method,” *The American Statistician*, vol. 46, no. 1, pp. 27–29, 1992.

[39] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, “CRAWDAD data set epfl/mobility (v. 2009-02-24).” Downloaded from http://crawdad.cs.dartmouth.edu/epfl/mobility, Feb. 2009.

[40] J. Scott, R. Gass, J. Crowcroft, P. Hui, C. Diot, and A. Chaintreau, “CRAWDAD data set cambridge/haggle (v. 2009-05-29).” Downloaded from http://crawdad.cs.dartmouth.edu/cambridge/haggle, May 2009.

[41] A.-K. Pietilainen and C. Diot, “CRAWDAD data set thlab/sigcomm2009 (v. 2012-07-15).” Downloaded from http://crawdad.cs.dartmouth.edu/thlab/sigcomm2009, July 2012.

[42] A. Balasubramanian, B. Levine, and A. Venkataramani, “Dtn routing as a resource allocation problem,” in *Proc. ACM SIGCOMM*, 2007.

[43] M. Conti, S. Giordano, M. May, and A. Passarella, “From opportunistic networks to opportunistic computing,” *Communications Magazine, IEEE*, vol. 48, pp. 126–139, sept. 2010.

[44] M. Karaliopoulos, “Assessing the vulnerability of DTN data relaying schemes to node selfishness,” *IEEE Communications Letters*, vol. 13, no. 12, pp. 923–925, 2009.

[45] Y. Li, G. Su, D. Wu, D. Jin, L. Su, and L. Zeng, “The impact of node selfishness on multicasting in delay tolerant networks,” *IEEE Trans. on Vehicular Technology,*, vol. 60, no. 5, pp. 2224–2238, 2011.