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

Research Report RR-13-284

Modelling and Analysis of Communication Traffic Heterogeneity in Opportunistic Networks

July 20^{th}, 2013
Last update July 20^{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

Modelling and Analysis of Communication Traffic Heterogeneity in Opportunistic Networks

Pavlos Sermpezis and Thrasyvoulos Spyropoulos

Abstract

In *opportunistic networks*, direct communication between mobile devices is used to extend the set of services accessible through cellular or WiFi networks. Due to their key role, mobility patterns and their impact in such networks have been extensively studied. In contrast, homogeneous communication traffic between nodes is assumed in most studies. This assumption is not generally true, as the mobility and social characteristics of nodes might affect the traffic exchanged between them. In this paper, we consider heterogeneous traffic patterns, propose appropriate models, and analyse the joint effect of traffic and mobility heterogeneity on the performance of popular forwarding algorithms. For example, we show that an increasing amount of (traffic and/or mobility) heterogeneity renders simple schemes, like direct transmission, significantly more useful than normally considered and diminishes the added value of additional randomly sprayed copies. We further validate these findings on datasets collected from real networks.

Index Terms

Opportunistic Networks, Heterogeneous Mobility, Heterogeneous Communication Traffic, Performance Modeling, Dataset Analysis

Contents

List of Figures

1 Introduction

Opportunistic or Delay Tolerant Networks (DTNs) [1] were initially envisioned to support communication in challenging environments, where infrastructure is limited or absent (e.g. emergency situations after disasters, mobile sensor networks). Lately, it has been suggested that they could also support or enhance existing networking infrastructure, e.g. by offloading traffic from cellular networks, enabling novel social and location-based applications, or introducing peer-to-peer collaborative computing [2, 3].

Opportunistic networks consist of mobile nodes (e.g. smartphones, laptops) that exchange data directly when they are *in contact* (i.e. within transmission range). Due to the limited range of direct communication (e.g. Bluetooth), communication is not continuous, and maintaining end-to-end paths is problematic. If nodes are not willing to relay 3rd party traffic, a message can only be transferred from a source node to a destination node when they come in contact (*direct transmission routing* [4]). If other nodes are willing to collaborate, they could copy the message from the source (or another relay), *store* and *carry* it and, finally, *forward* it when they encounter the destination node. Such replication and relaying schemes could improve performance (*relay-assisted routing*, e.g. [5, 6]), albeit at increased complexity and resource overhead.

Since message exchanges take place only during contacts between nodes, mobility plays a major role both in the performance and the design of protocols and applications. As a result, sophisticated utility-based schemes have been proposed that select relays based on their mobility patterns and/or social characteristics [7]. Furthermore, a lot of effort has been made recently to capture the mobility patterns of real networks in simulations and theory [8, 9, 10]. These mobility patterns can often greatly affect the performance of different schemes.

Somewhat surprisingly, the communication traffic patterns used in studies of opportunistic networks have not received an equal amount of attention. It is usually assumed, implicitly or explicitly, that all traffic is uniform: each pair of nodes exchanges the same amount of messages. However, intuition suggests that traffic between nodes, just like contacts, cannot be expected to be homogeneous either. This is also supported by empirical studies on social networks [11, 12], where the frequency of message exchanges might widely vary among pairs of nodes. Further, nodes that have a social relation or reside/move in the same areas, often tend to exchange more messages than others. Therefore, a number of interesting questions arise: *How should one model the heterogeneity in communication traffic? Do heterogeneous traffic patterns affect the performance of information dissemination mechanisms and to what extent?*

Towards answering this question, in this paper we investigate *if*, *when* and *how* traffic patterns affect the communication performance in opportunistic networks. Specifically:

- We examine what characteristics of traffic heterogeneity affect performance, and propose an analytically tractable model that can describe a large range of non-uniform traffic patterns (Section 2).

- We derive analytical expressions for calculating the joint effect of traffic and mobility heterogeneity in the performance of basic forwarding mechanisms (Section 3).
- We use these expressions to show that the common understanding about these mechanisms, e.g. the gains from having additional replicas, might radically change when traffic is heterogeneous (Section 3.2).
- We validate our analytical findings through simulations (Section 4.1) and, by applying them to datasets of real-world networks that contain information about both the mobility and the communication patterns of participating nodes (Section 4.2).

2 Network Model

2.1 Mobility

We consider a network $\mathcal{N}$, where N nodes move in an area, much larger than their transmission range. Data packet exchanges between a pair of nodes can take place only when they are in proximity (*in contact*). Hence, the dissemination of a message is subject to nodes mobility and the resulting *contact events*. To model this sequence of contact events, we will assume the following *class* of heterogeneous contact models.

Definition 1 (Heterogeneous Contact Network)**.**

Assumption 1: Contact events between a pair of nodes $\{i, j\}$ follow an independent Poisson process with rate λ_{ij}, i.e. inter-contact times are independent and exponentially distributed with rate λ_{ij}.

Assumption 2: Contact rates λ_{ij} are independently drawn from an arbitrary distribution with probability density function $f_\lambda(x)$ (with finite mean μ_λ and variance σ_λ^2).

Assumption 3: Contact duration is negligible compared to the time between contacts events, though sufficient for all data transfers to take place.

The assumption of Poisson contacts is pretty standard, and allows us to use a Markovian framework for analyzing dissemination processes, similarly to the majority of previous studies in Opportunistic / Delay Tolerant Networks [8, 10, 13, 14], as well as in other fields involving networks and/or contacts [15]. Furthermore, analyses of real-world contact traces provide some support, suggesting that the observed inter-contact time distributions (or, at least, their *tails*) can often be approximated by exponential distributions [16, 17]. While this assumption can sometimes be relaxed, to our best knowledge this only allows to derive asymptotic results [18, 19].

The second assumption introduces some heterogeneity in the standard model, in an attempt to better align the model with the findings of real-world trace analyses showing that the contact rates (or frequencies) of different pairs are largely heterogeneous [9, 16]. Moreover, by allowing rates to be drawn from an arbitrary distribution f_λ, we can (i) emulate a very diverse set of contact (and thus mobility) scenarios, e.g. a symmetric f_λ (uniform, normal) implies a balanced number of high and low rates, while a right-skewed f_λ, like a Pareto distribution, describes a network with most pairs having large intercontact times, but few meeting very frequently, and (ii) fit this distribution to match the rates observed in a real trace.

The third assumption is equivalent to saying that there are no bandwidth concerns in our framework. This is not always true [20], and, clearly, transmission and storage capacity are important issues in DTNs [7]. Nevertheless, these are mostly orthogonal issues to the topic of our study, so we don't consider them in our framework.

Summarizing, our main motivation for this model is to maintain the analytical tractability properties of standard models, while also integrating some mobility heterogeneity, whose joint effect with traffic heterogeneity we want to investigate. To ensure that our assumptions do not confound the conclusions drawn from our analysis, we will validate our results against real measurement traces, where many of these assumptions are known to not hold.

2.2 Communication Traffic

In addition to who meets whom and how often, another major question that should be raised in opportunistic networks (but rarely is) is *who wants to communicate with whom* and *how much traffic do they exchange*?

Intuitively, every pair of nodes cannot exchange the same amount of traffic or with the same frequency. To support intuition, studies from fields related to technological and social networks [11, 12, 21] have demonstrated the existence of heterogeneous traffic patterns. This seems to argue for a more complex traffic model (compared to a homogeneous one). The same studies further suggest that this heterogeneity depends on the *spatial* and *social* characteristics of these networks. Since *location-based services* [22] and *social networking* [23] are considered among the major applications supported by opportunistic networks, such traffic dependencies on social and/or spatial factors are very probable to appear. What is more, mobility characteristics have also been found to depend on spatial and social characteristics [9, 21, 24]. We would thus expect that traffic and mobility in such networks would exhibit some clear correlations [11, 12]. For instance, the *Pearson correlation coefficient* between mobility (contact rates) and traffic (rates) in a subset (i.e. in three cities) of the dataset analysed in [11], is presented in Table 1, where it is evident that the correlation is significant.

Table 1: Mobility-Traffic Correlation

	Austin	Oslo	San Francisco
corr	0.91	0.65	0.92

Before we embark on choosing the right model features, one should consider the following questions: *What features of heterogeneous traffic, if any, have an impact on the performance of opportunistic communication? Does the existence of heterogeneous traffic* per se *suffice to affect performance?*

Towards obtaining some initial understanding, we thus decided to compare the performance of some well-known opportunistic protocols (direct transmission [4], spray and wait [5], and 2-hop routing [14]) through Monte Carlo simulations on two real traces (we will discuss these traces in more detail, later, in Section 4), for three traffic scenarios: (i) homogeneous traffic: every pair of nodes exchange messages with the same rate; (ii) heterogeneous traffic that is *mobility independent*: we assign randomly to each pair a different traffic rate (drawn from a uniform distribution in $[1, 1000]$); (iii) heterogeneous traffic that is *mobility dependent*: traffic between two nodes is proportional to their contact rate.

Results for the mean message delivery delay are shown in Fig. 1. As is evident from these results, when traffic heterogeneity is independent of mobility, the average delay is practically the same to the homogeneous case, *for all protocols, and across all scenarios* (including additional ones we've tried). While this result might seem somewhat surprising, it should not be. Since contacts between nodes are the only way to move messages around the network, unless the traffic pattern is *correlated* to mobility, this should not make a difference in the performance observed. In contrast, Fig. 1 shows a clear difference in average delay for all scenarios and protocols, compared to the homogeneous case, when traffic is heterogeneous *and* correlated with the contact rates.

These results provide an initial answer to the above questions: *it is not traffic heterogeneity itself that affects performance, but rather the joint effect of mobility and traffic (heterogeneity).* This observation, together with the initial insight coming from real datasets, motivates us to propose the following simple model for the amount of traffic between node pairs, that we will use through our analysis. While relatively simple, this model introduces the key feature, correlation between traffic and mobility, and allows us to model a number of interesting traffic patterns and amounts of (positive or even negative) correlations.

Definition 2 (Heterogeneous Communication Traffic)**.** *The end-to-end traffic demand (per time unit) between a pair of nodes $\{i, j\}$, is a random variable τ_{ij}, such that $E[\tau_{ij}] = \tau(\lambda_{ij})$, where $\tau(\cdot)$ is a continuous function from $\mathbb{R}^+$ to $\mathbb{R}^+$.*

Hence, traffic demand between node pairs can differ and is *on average* correlated with the nodes' contact rate. However, τ_{ij} itself is still random, allowing some node pairs to have little traffic demand even if they meet often (e.g. "familiar

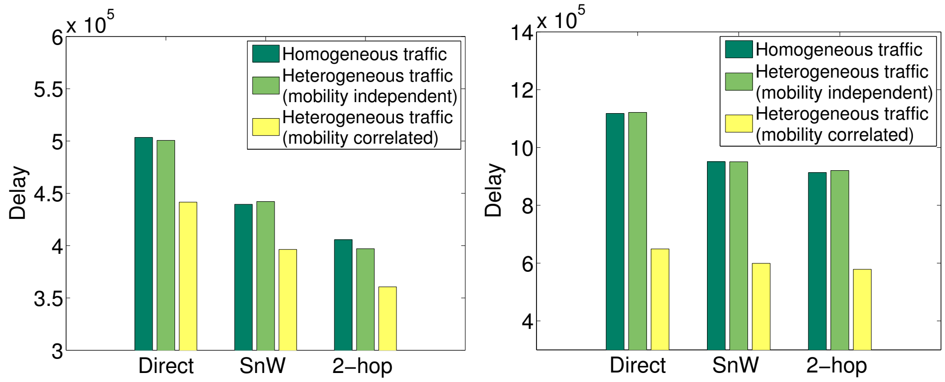

(a) Gowalla dataset (b) Strathclyde dataset

Figure 1: Mean message delivery delay of three routing protocols, namely Direct Transmission, Spray and Wait (*SnW*) and 2-hop, on the (a) Gowalla and (b) Strathclyde dataset.

strangers"). Furthermore, through the function $\tau(\cdot)$ one can introduce a number of different types and amounts of (positive or negative) correlations between traffic and mobility. While real mobility and traffic patterns are clearly expected to have a number of additional nuances and details, not captured by the models of Def. 1 and Def. 2, it turns out that these abstractions are still "rich" enough to allow us to draw useful conclusions.

3 Analysis

Consider now an opportunistic network with mobility and traffic according to the definitions of Section 2. To calculate a performance metric for this network, e.g. the expected delay, one would consider a large number of messages generated between various source-destination pairs. However, when end-to-end traffic is heterogeneous (as in Def. 2) the *effective* contact rate between sources and destinations of these message will be different than f_λ (Def. 1), the contact rate between a randomly chosen pair of nodes.

Proposition 1. *The probability density function f_τ of the contact rate between the source and the destination $\{s, d\}$ of a random message, in a network following Definitions 1 and 2, converges in probability as follows:*

$$f_\tau(x) \xrightarrow{p} \frac{1}{C} \cdot \tau(x) \cdot f_\lambda(x) \tag{1}$$

where $C = E[\tau(\lambda)] = \int_0^\infty \tau(x) f_\lambda(x) dx$

Proof. Consider a network $\mathcal{N}$ with N nodes. Let $d\lambda = O\left(\frac{1}{N}\right)$, and define the set of nodes with contact rate $\lambda_{ij} \in [\lambda, \lambda + d\lambda)$:

$$\mathcal{N}(\lambda) = \{\{i, j\} : i, j \in \mathcal{N}, \lambda \leq \lambda_{ij} < \lambda + d\lambda\},$$

The total number of messages generated per time unit between pairs $\in \mathcal{N}(\lambda)$ is equal to

$$T(\lambda) = \sum_{\{i,j\} \in \mathcal{N}(\lambda)} \tau_{ij} \quad (2)$$

where τ_{ij} in the sum are i.i.d. random variables with mean $\tau(\lambda)$. Then, the probability that the contact rate λ_{sd}, between the source and the destination of a randomly selected message, is in $[\lambda, \lambda + d\lambda)$, is given by

$$P\{\lambda \leq \lambda_{sd} < \lambda + d\lambda\} = \frac{T(\lambda)}{\sum_i \sum_j \tau_{ij}} = \frac{\sum_{\{i,j\} \in \mathcal{N}(\lambda)} \tau_{ij}}{\sum_i \sum_j \tau_{ij}} \quad (3)$$

We can express Eq. (3) as following:

$$P\{\lambda \leq \lambda_{sd} < \lambda + d\lambda\} = \frac{T(\lambda)}{\|\mathcal{N}(\lambda)\|} \cdot \frac{\|\mathcal{N}(\lambda)\|}{N(N-1)/2} \cdot \frac{N(N-1)/2}{\sum_i \sum_i \tau_{ij}}$$

where $\|\cdot\|$ denotes the cardinality of a set and $\frac{N(N-1)}{2}$ is the total number of node pairs in a network with N nodes. Let us further denote:

$$X_1 = \frac{T(\lambda)}{\|\mathcal{N}(\lambda)\|},\ X_2 = \frac{\|\mathcal{N}(\lambda)\|}{N(N-1)/2},\ X_3 = \frac{\sum_i \sum_i \tau_{ij}}{N(N-1)/2}$$

Applying the weak law of large numbers [25], it holds that for a *large network*[1]

$$X_1 \xrightarrow{p} \tau(\lambda) \quad \text{and} \quad X_2 \xrightarrow{p} f_\lambda(\lambda) \quad (4)$$

where $\xrightarrow{p}$ denotes convergence in probability.

Also, X_3 corresponds to the sample average of τ_{ij} over all disjoints sets $\mathcal{N}(\lambda)$. Thus, applying Cramér's theorem (*Theorem 6.5* in [25])[2] and using the convergence expressions of Eq. (4), we can get

$$X_3 \xrightarrow{p} \int_0^\infty \tau(y) f_\lambda(y) dy = E[\tau(\lambda)] = C$$

Similarly, using Cramér's theorem, it can be shown that the expression $X_1 \cdot X_2 \cdot \frac{1}{X_3}$ converges too, i.e.

$$X_1 \cdot X_2 \cdot \frac{1}{X_3} \xrightarrow{p} \tau(\lambda) \cdot f_\lambda(\lambda) \cdot \frac{1}{C}$$

Finally, denoting the probability density function of the source-destination contact rate λ_{sd} as $f_\tau(\lambda)$, i.e. $P\{\lambda \leq \lambda_{sd} < \lambda + d\lambda\} = f_\tau(\lambda) d\lambda$ gives us the desired result. $\square$

[1] When $N \to \infty$, then $d\lambda = O\left(\frac{1}{N}\right) \to 0$, and $\|\mathcal{N}(\lambda)\| = O\left(\frac{N(N-1)}{2} d\lambda\right) = O(N) \to \infty$.

[2] Equivalently, one could use here the *Continuous Mapping Theorem*.

3.1 End-to-end Delivery Performance

An opportunistic routing protocol tries to deliver the end-to-end traffic demand τ_{ij}, and we would like to consider the effects of different contact patterns f_λ and traffic patterns $\tau(\lambda)$ on its performance. There exists a very large abundance of proposed schemes [7] and it would not be possible, nor would it provide any intuition, to analyze the effect of heterogeneity on each and every one. Instead, we focus here on some basic mechanisms to gain intuition.

The approach with the minimum overhead and complexity is *Direct Transmission ("DT")*: nodes wishing to exchange data or information with each other, may do so, only when they are in direct contact, without involving any relays. Furthermore, this is the only feasible approach if nodes do not have incentives to relay traffic they are not personally interested in, e.g. due to privacy or resource-related concerns [26]. Nevertheless, Direct Transmission is known to suffer from long delays and low throughput [27].

To improve the performance of direct transmission, *replication* or *relay-assisted* schemes can be used. Extra copies can be handed over to encountered nodes, and the destination can receive the message from either the source or any of the relays, reducing thus the expected delivery delay. Taken to the extreme, schemes like epidemic routing [28] forward the message at every possible encounter (deterministically, probabilistically, or based on some utility-function). Yet these do not usually scale well beyond networks with few tens of nodes, due to large resource usage.

In networks with homogeneous mobility and traffic, it is known that using just a few extra copies leads to significant performance gains. For example, in a network of 1000 nodes, simply distributing 10 extra copies to the first 10 nodes encountered provides an almost 10-fold improvement in delay compared to direct transmission [5]. Although this also comes with a 10-fold increase in the amount of (storage and bandwidth) resources needed, it presents a very useful tradeoff to DTN protocol designers.

However, when it comes to heterogeneous mobility and traffic, Proposition 1 suggests that, unlike the above example, the source is no longer equivalent with other random relays, in terms of their probability of contacting an intended destination soon. It is thus of particular interest to examine whether the above trade-off still holds, if one considers the joint effect of realistic mobility and communication traffic patterns.

We thus consider, in the following, *Relay-assisted* routing, which is a simple abstraction of schemes that use extra randomly chosen relays[3]. We derive expressions for the *maximum achievable performance gain* compared to Direct Transmission, in terms of delivery delay and delivery probability, the two main metrics considered in related work. These results provide bounds for the maximum achievable improvement by Relay-assisted routing, as a function of the number of extra copies L.

[3] We will briefly consider utility-based schemes in Section 5.

Result 1. *Let R denote the ratio of the expected delivery delay of Relay-Assisted routing, $E[T_R]$, over the expected delivery delay of Direct Transmission routing, $E[T_{DT}]$. When L extra copies are used, then*

$$R = \frac{E[T_R]}{E[T_{DT}]} \quad \geq \quad \frac{1}{E\left[\frac{\tau(\lambda)}{\lambda}\right]} \cdot E\left[\frac{\tau(\lambda)}{\lambda + L \cdot \mu_\lambda}\right] \tag{5}$$

where the expectations are taken over f_λ.

Proof. Let $I_{sd}(t)$ be an indicator random variable that is equal to 1 if nodes s and d are within transmission range at time t, and 0 otherwise. Let further T_{sd} denote the random inter-contact time between node pair $\{s, d\}$:

$$T_{sd} = \inf\{t > 0 : I_{sd}(0) = 1, I_{sd}(0^+) = 0, I_{sd}(t) = 1\}.$$

Since we have assumed that contact duration is negligible (Assumption 3 of Def.1), the contact process is essentially a point process, and the above could be simplified to $T_{sd} = \inf\{t > 0 : I_{sd}(0) = 1, I_{sd}(t) = 1\}$.

Assume now that end-to-end messages between $\{s, d\}$ are generated at random times and *independently* from the contact process. If T_{DT} denotes the delay of directly transmitting a message from s to d, and the contact rate between s and d is $\lambda_{sd} = x$, then one can use renewal-reward theory [29] to show that

$$E[T_{DT}|\lambda_{sd} = x] = E[T_{sd}^{(e)}|\lambda_{sd} = x] = \frac{1}{x}.$$

That is, the expected delay of direct transmission is equal to the mean of the *residual* (or *excess*) inter-contact time $T_{sd}^{(e)}$, which is an exponential variable with the same rate x.

Using the property of conditional expectation and the distribution of λ_{sd} (Proposition 1) we can get:

$$\begin{aligned} E[T_{DT}] &= \int_0^\infty E[T_{DT}|\lambda_{sd} = x] f_\tau(x) dx = \int_0^\infty \frac{1}{x} f_\tau(x) dx \\ &= \frac{1}{C} \int_0^\infty \frac{\tau(x)}{x} f_\lambda(x) dx = \frac{1}{E[\tau(\lambda)]} \cdot E\left[\frac{\tau(\lambda)}{\lambda}\right] \end{aligned} \tag{6}$$

Assume now that the same messages, between $\{s, d\}$ are routed using Relay-Assisted routing, with L message copies given to L relays. Let T_R^* denote the *total* delay to deliver a message using Relay-Assisted routing, T_R the remaining delay after all L copies have been distributed, T_{fwd} the time to distribute the L copies to the L relays, and $p_{fwd} = P(T_R^* < T_{fwd})$ the probability that the message is delivered to the destination before L relay nodes have been found.

Since relays are selected randomly (e.g. [5]), $p_{fwd} = \frac{L}{N} \to 0$ for $L \ll N$. Similarly, if $L^2 \ll N$, $\frac{T_{fwd}}{T_R} \to 0$ [5]. We can thus focus only on T_R, the time after L relays have received a copy. Note that this assumption does not affect the validity of the results (i.e. they still are bounds).

Denote now with $\mathcal{L}$ the set of selected relays. Using a similar argument as in the direct transmission case, and Assumption 1 of Def.1 we can easily show that,

$$T_R \equiv \min_{i \in \mathcal{L} \cup \{s\}} T_{id} \sim exp(X_r)$$
$$X_r = \lambda_{sd} + \sum_{i \in \mathcal{L}} \lambda_{id} = \lambda_{sd} + X_R$$

where $X_R = \sum_{i \in \mathcal{L}} \lambda_{id}$, and the expected value of T_R will be

$$E[T_R] = \frac{1}{X_r} = \frac{1}{\lambda_{sd} + X_R}, \tag{7}$$

where $\lambda_{sd} \sim f_\tau$ (Proposition 1) and $X_R \sim f_R = f_\lambda^{(*L)}$, the L-fold convolution of f_λ.

Then, from Eq. (7) and using the property of conditional expectation, we find:

$$\begin{aligned} E[T_R] &= \int_0^\infty \int_0^\infty E\left[T_R | \lambda_{sd} = x, X_R = y\right] f_\tau(x) dx f_R(y) dy \\ &= \int_0^\infty \int_0^\infty \frac{1}{x+y} \cdot f_\tau(x) dx \cdot f_R(y) dy \\ &= \frac{1}{E[\tau(\lambda)]} \int_0^\infty \int_0^\infty \frac{\tau(x)}{x+y} \cdot f_\lambda(x) dx \cdot f_R(y) dy \end{aligned} \tag{8}$$

where in the last equality we substituted the expression for f_τ from Proposition 1. Now Eq. (8) can also be written as

$$E[T_R] = \frac{1}{E[\tau(\lambda)]} \cdot \int_0^\infty E_R\left[\frac{1}{x+y}\right] \cdot \tau(x) \cdot f_\lambda(x) dx \tag{9}$$

where the expectation $E_R[\cdot]$ is taken over f_R. Using Jensen's inequality[4] for the function $h(y) = \frac{1}{x+y}$, we get:

$$E_R\left[\frac{1}{x+y}\right] \geq \frac{1}{x + E_R[y]} \tag{10}$$

where $E_R[y]$ is given by (as the expectation of a sum of L i.i.d. random variables with expectation μ_λ) [29]:

$$E_R[y] = E[X_R] = E\left[\sum_{i \in \mathcal{L}} \lambda_{id}\right] = L \cdot \mu_\lambda \tag{11}$$

Hence, using Eq. (10) and Eq. (11) in Eq. (9), we get

$$E[T_R] \geq \frac{\int_0^\infty \frac{\tau(x)}{x + L \cdot \mu_\lambda} \cdot f_\lambda(x) dx}{E[\tau(\lambda)]} = \frac{E\left[\frac{\tau(\lambda)}{\lambda + L \cdot \mu_\lambda}\right]}{E[\tau(\lambda)]} \tag{12}$$

Finally, dividing Eq. (12) with Eq. (6) gives Eq. (5). $\square$

[4]Jensen's inequality for a convex function $h(x)$: $E[h(x)] \geq h\left(E[x]\right)$.

Result 2. *Denote with $P_{(src.)}$ the probability that a message is delivered to the destination by the source node, rather than by any of the relays. Then,*

$$P_{(src.)} \;\; \geq \;\; \frac{1}{E[\tau(\lambda)]} \cdot E\left[\frac{\lambda \cdot \tau(\lambda)}{\lambda + L \cdot \mu_\lambda}\right] \tag{13}$$

Proof. Using similar arguments and notation as Result 1, the above event is equivalent to the destination contacting the source before any other relay.

Then, $P_{(src.)} \equiv P\{T_{sd} < T_{r-d}\}$ (where $T_{r-d} = \min_{i\in\mathcal{L}}\{T_{id}\}$), will be given by the ratio $\frac{\lambda_{sd}}{\lambda_{sd}+X_R}$ [29]. Conditioning on the the rates λ_{sd} and X_R, we can write

$$\begin{aligned}
&P_{(src.)} \equiv P\{T_{sd} < T_{r-d}\} = \\
&= \int_0^\infty \int_0^\infty P\{T_{sd} < T_{r-d} | \lambda_{sd} = x, X_R = y\} f_\tau(x)dx f_R(y)dy \\
&= \int_0^\infty \int_0^\infty \frac{x}{x+y} \cdot f_\tau(x)dx \cdot f_R(y)dy \\
&= \int_0^\infty \int_0^\infty E_R\left[\frac{x}{x+y}\right] \cdot f_\tau(x)dx
\end{aligned} \tag{14}$$

and applying Jensen's inequality as in Eq. (10), we get

$$\begin{aligned}
P_{(src.)} &\geq \int_0^\infty \frac{x}{x + L\cdot\mu_\lambda} \cdot f_\tau(x)dx \\
&= \frac{1}{E[\tau(\lambda)]} \cdot \int_0^\infty \frac{x\cdot\tau(x)}{x + L\cdot\mu_\lambda} \cdot f_\lambda(x)dx
\end{aligned} \tag{15}$$

which is equal to Eq. (13) and proves the result. □

3.2 Insights for Real Opportunistic Networks

The expressions we derived in Result 1 are generic and can be used under any mobility and traffic pattern (i.e. for any f_λ and $\tau(\cdot)$). To obtain some further insights, in this section, we consider specific classes of mobility and traffic patterns that capture commonly observed characteristics of real networks.

Mobility

We will assume the contact rates to be *gamma distributed* (i.e. $f_\lambda \sim \Gamma(\alpha,\beta)$).

Our choice is initially motivated by the findings of Passarella et al. [9], who have shown, through statistical analysis of pervasive social networks' datasets, that the *Gamma distribution* matches well the observed contact rates. In addition, the analytical findings of [9], further suggest that the choice of a Gamma distribution can be supported in real opportunistic networks and can explain many of the observed properties (e.g. distribution of *aggregate* inter-contact times). Finally, by selecting appropriately the parameters α and β, we can assign *any* desired value to the mean value μ_λ and the variance σ_λ^2 of the contact rates. This allows us to describe (or fit up to the first two moments) a large range of scenarios with different mobility heterogeneities captured by $CV_\lambda = \frac{\sigma_\lambda}{\mu_\lambda}$.

Traffic

We further describe the traffic using a *polynomial function* of the form $\tau(x) = c{\cdot}x^k$, $c > 0$.

As in the case of mobility, the reasons for our choice are as following. Observations of real networks have shown that the nodes with high contact frequencies tend to exchange more traffic [11, 12], which is consistent with the above choice when $k > 0$. In general, the value of k captures the amount of traffic heterogeneity. Furthermore, by choosing $0 < k < 1$ (or $k > 1$) one can emulate concave (or convex) functions and, thus, approximate different traffic patterns. Finally, one can also consider negative correlations, by choosing $k < 0$. Although less common, these could arise, for example, in applications where users want to communicate more when they do not meet frequently (e.g. messaging).

Under the above assumptions, the following result for the relative performance of the information dissemination mechanisms we consider in this paper, holds.

Result 3. *In a Heterogeneous Contact Network where $f_\lambda \sim \Gamma(\alpha, \beta)$ with mean value μ_λ and variance σ_λ^2 (coefficient of variation $CV_\lambda = \frac{\sigma_\lambda}{\mu_\lambda}$) and $\tau(x) = c \cdot x^k$, it holds:*

$$1 \geq \ \ R = \frac{E[T_R]}{E[T_{DT}]} \ \ \geq R_{min} = \frac{1 + (k-1) \cdot CV_\lambda^2}{1 + k \cdot CV_\lambda^2 + L} \tag{16}$$

for $k > 1 - \frac{1}{CV_\lambda^2}$, and

$$1 \geq \ \ P_{(src.)} \ \ \geq \ \ P_{min} = \frac{1 + k \cdot CV_\lambda^2}{1 + (k+1) \cdot CV_\lambda^2 + L} \tag{17}$$

for $k \geq -\frac{1}{CV_\lambda^2}$.

Proof. Eq. (5), for $\tau(x) = c \cdot x^k$, is written as

$$R \geq \frac{1}{E[\lambda^{k-1}]} \cdot E\left[\frac{\lambda^k}{\lambda + L \cdot \mu_\lambda}\right] \tag{18}$$

The expectations in Eq. (18) are taken over the contact rates'(*Gamma*) distribution, whose general form is [29]

$$f_\lambda(x) = \tfrac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x}$$

where $\alpha > 0$ is the *shape* parameter, $\beta > 0$ the *rate* parameter. Its mean value and variance are given by $\mu_\lambda = \frac{\alpha}{\beta}$ and $\sigma_\lambda^2 = \frac{\alpha}{\beta^2}$, respectively, and, equivalently, we can write

$$\alpha = 1/CV_\lambda^2, \quad \beta = 1/\left(\mu_\lambda \cdot CV_\lambda^2\right) \tag{19}$$

To calculate Eq. (18), first we find an expression for $E[\lambda^{k-1}]$:

$$
\begin{aligned}
E[\lambda^{k-1}] &= \int_0^\infty x^{k-1} f_\lambda(x)dx = \int_0^\infty x^{k-1} \frac{\beta^\alpha}{\Gamma(\alpha)} x^{\alpha-1} e^{-\beta x} dx \\
&= \frac{\Gamma(k-1+\alpha)}{\Gamma(\alpha)} \frac{1}{\beta^{k-1}} \int_0^\infty \frac{\beta^{k-1+\alpha}}{\Gamma(k-1+\alpha)} x^{(k-1+\alpha)-1} e^{-\beta x} dx \\
&= \frac{\Gamma(k-1+\alpha)}{\Gamma(\alpha)} \frac{1}{\beta^{k-1}} \qquad (20)
\end{aligned}
$$

where the integral in the second line is equal to 1 because the integrated function is the pdf of a Gamma distribution with parameters $\alpha^{'} = k-1+\alpha$ (it must hold that $\alpha^{'} > 0$, which means that $k > 1-\alpha = 1 - \frac{1}{CV_\lambda^2}$) and $\beta^{'} = \beta$.

Similarly to the derivation of Eq. (20), it can be shown that

$$
\begin{aligned}
& E\left[\frac{\lambda^k}{\lambda + L\cdot\mu_\lambda}\right] = \\
&\quad = \frac{\Gamma(k+\alpha)}{\Gamma(\alpha)} \frac{1}{\beta^k} \cdot \int_0^\infty \frac{1}{x+L\cdot\mu_\lambda} \frac{\beta^{k+\alpha}}{\Gamma(k+\alpha)} x^{k+\alpha-1} e^{-\beta x} dx \\
&\quad = \frac{\Gamma(k+\alpha)}{\Gamma(\alpha)} \frac{1}{\beta^k} \cdot E_{\lambda'}\left[\frac{1}{\lambda' + L\cdot\mu_\lambda}\right] \qquad (21)
\end{aligned}
$$

where λ' follows a Gamma distribution with parameters $\alpha' = k+\alpha$ and $\beta' = \beta$. Since the function $g(x) = \frac{1}{x+c}$ is convex, we can apply Jensen's inequality to Eq. (21) and get

$$
\begin{aligned}
E\left[\frac{\lambda^k}{\lambda + L\cdot\mu_\lambda}\right] &\geq \frac{\Gamma(k+\alpha)}{\Gamma(\alpha)} \frac{1}{\beta^k} \cdot \frac{1}{E[\lambda'] + L\cdot\mu_\lambda} \\
&= \frac{\Gamma(k+\alpha)}{\Gamma(\alpha)} \frac{1}{\beta^k} \cdot \frac{1}{\frac{k+\alpha}{\beta} + L\cdot\mu_\lambda} \qquad (22)
\end{aligned}
$$

where we substituted $E[\lambda'] = \frac{\alpha'}{\beta'} = \frac{k+\alpha}{\beta}$.

Thus, from Eq. (20) and Eq. (22), it holds for R (Eq. (18)):

$$
R \geq \tfrac{\Gamma(k+\alpha)}{\Gamma(k-1+\alpha)} \cdot \tfrac{1}{\beta} \cdot \tfrac{1}{\frac{k+\alpha}{\beta} + L\cdot\mu_\lambda}
$$

and because of the Gamma function's property $\Gamma(z+1) = z\cdot\Gamma(z)$, we can write

$$
R \geq \tfrac{k-1+\alpha}{\beta} \cdot \tfrac{1}{\frac{k+\alpha}{\beta} + L\cdot\mu_\lambda} \qquad (23)
$$

and Eq. (16) follows easily by substituting α and β from Eq. (19) to Eq. (23).

Now, to find $P_{(src.)}$, at first, we set $\tau(x) = c\cdot x^k$ in Eq. (13)

$$
P_{(src.)} \geq \frac{1}{E[\lambda^k]} \cdot E\left[\frac{\lambda^{k+1}}{\lambda + L\cdot\mu_\lambda}\right] \qquad (24)
$$

Using Eq. (20) - Eq. (22) (where instead of k we consider $k+1$), the result for $P_{(src.)}$ follows similarly as before. $\square$

The expressions of Result 3 depend only on 3 parameters (CV_λ, k, L) and, thus, could be used to tune Relay-Assisted schemes: At first, since mobility (CV_λ) and traffic (k) parameters are characteristics of the network, they either remain constant or change slowly over a long time period. Hence, we can assume that nodes know their values, or can estimate them (e.g. with a distributed mechanism, locally, etc.). Then, the required number of relays L to achieve a certain expected delay, could be easily estimated.

Practical Example: If the measured network characteristics are $CV_\lambda = 2$ and $k = 2$, then from Result 3 we get $R = \frac{5}{9+L}$. Therefore, to achieve delivery delay two times faster than Direct Transmission, one extra copy should be used ($L = 1 \rightarrow R = 0.5$), while to achieve 4 times faster delivery, $L = 11$ relay nodes are needed. In the latter case, if traffic/mobility heterogeneity has not been taken into account [5], the prediction would be $L = 3$ and this would lead only to 2.5 (instead of 4) times faster delivery (i.e. $R = \frac{5}{12}$).

3.3 Discussion and Implications

It is evident from the above example that traffic heterogeneity can have a major impact on performance and thus protocol design. Table 2 formalizes this impact, by considering how R_{min} and P_{min} (Eq. (16) and Eq. (17)) behave: the *middle column* shows their monotonicity as traffic heterogeneity (k), mobility heterogeneity (CV_λ), and amount of extra copies (L) increase; the *right column* gives their values in the limit for large/small k or CV_λ.

Some important observations that follow from Table 2 are:

O.1. A strong positive correlation (large k) between traffic and mobility reduces the added value of extra copies (i.e. R_{min}, P_{min} increase), while a negative (or weak positive) correlation increases the added value.

O.2. For high heterogeneity (traffic and mobility), our results imply that a unicast message is likely to arrive to its destination at the time the source and destination come in contact (i.e. within physical proximity). This raises questions about the usefulness of opportunistic networking for unicast applications in which end-to-end traffic is highly correlated with contact frequency (e.g. Facebook messaging). In fact, as positive correlations and high mobility heterogeneity is what has mostly been observed in measurement analyses [11, 12], the take home message might be rather grim for unicast applications over opportunistic networks.

O.3. On the other hand, our results suggest that potential unicast applications with an end-to-end traffic demand between nodes with non-frequent meetings (e.g. social peers residing in different communities) could benefit a lot (more than normally assumed) by simple opportunistic networking solutions (e.g. replication-based routing).

O.4. Finally, a positive message can be inferred for the feasibility of content-centric applications (e.g. file sharing, service composition [2]), which normally rely on direct encounters with nodes storing content of interest. This is due to the increased relative performance of Direct Transmission, compared to the uniform traffic case.

Table 2: R_{min}, P_{min}: Monotonicity and Asymptotic Limits

Parameter x	Monotonicity as $x \nearrow$	Limits fpr $min\{x\}$	 $max\{x\}$
mobility heterogeneity: $CV_\lambda \in [0, \infty)$	R_{min} $\nearrow$ if $k > 1 + \frac{1}{L}$ $\searrow$ otherwise	$\frac{1}{1+L}$	$1 - \frac{1}{k}$
	P_{min} $\nearrow$ if $k > \frac{1}{L}$ $\searrow$ otherwise	$\frac{1}{1+L}$	$1 - \frac{1}{k+1}$
traffic heterogeneity: $k \in [k_{min}, \infty)$	$R_{min}, P_{min} \nearrow$	0	1
extra copies: $L \quad (L \ll N)$	$R_{min}, P_{min} \searrow$	-	

However, a further treatment of content-centric traffic is beyond the scope of this paper.

4 Model Validation

To validate our model and analysis, in this section we compare the theoretical results against Monte Carlo simulations, on various synthetic scenarios and on datasets of real networks.

4.1 Synthetic Simulations

We generate synthetic networks, conforming to the mobility and traffic models of Section 2, as following: We assign to each pair $\{i, j\}$ a contact rate λ_{ij}, which we draw randomly from f_λ, and create a sequence of contact events (Poisson process with rate λ_{ij}). Since $E[\tau_{ij}] = \tau(\lambda_{ij})$ (from Def. 2), we draw the traffic rate for each pair $\{i, j\}$ as $\tau_{ij} \sim Uniform[0, 2\tau(\lambda_{ij})]$. Then, we simulate a large number of message exchanges, choosing randomly for each message the source-destination pair according to the weights τ_{ij}.

We created different scenarios ($N, f_\lambda, \tau(\cdot)$) to verify our analysis under various network parameters.

To keep up with the analysis of Section 3.2, we used the *Gamma distribution* as the contact rates distribution f_λ and traffic functions of a polynomial form, $\tau(x) = c \cdot x^k$.

Fig. 2 shows the simulation results for the ratio of delivery delays R (Fig. 2(a)) and the probability of delivery by the source $P_{(src.)}$ (Fig. 2(b)) along with the theoretical bounds R_{min} and P_{min} (dashed lines). The results correspond to two different traffic scenarios, namely (i) $\tau(x) = c \cdot x^2$ and (ii) $\tau(x) = c \cdot x^4$. Similarly, in

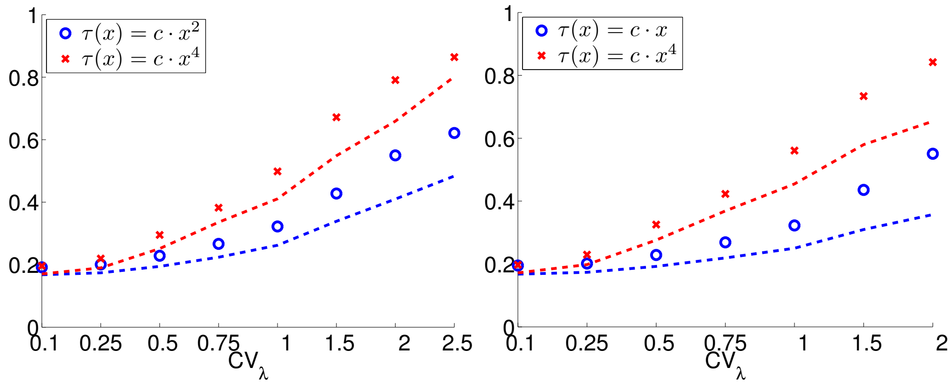

(a) $R = \frac{E[T_R]}{E[T_{DT}]}$ (b) $P_{(src.)}$

Figure 2: Comparison of Direct Transmission and Relay Assisted routing in two traffic scenarios under varying mobility heterogeneity: (a) The ratio R and the theoretical bound $Rmin$, (b) The probability $P_{(src.)}$ and the theoretical bound $Pmin$. Simulation results are denoted with markers and theoretical predictions (lower bounds) with lines.

Fig. 3 we present the results for two different mobility scenarios with (i) light heterogeneity $CV_\lambda = 0.5$ and (ii) heavy heterogeneity $CV_\lambda = 2$, under varying traffic heterogeneity. As expected, our predictions are always lower than the simulation results and the bounds are tight, except for some cases with heavy heterogeneity both in mobility and traffic. Also, as our results predict (Table 2), for increasing mobility and traffic heterogeneity, the lower bounds (and the actual simulated values as well) increase, demonstrating that the gain of the extra copies diminishes under such conditions. For example, from Fig. 2(a) we can see that, while for almost homogeneous mobility ($CV_\lambda = 0.1$) the expected delay of Relay-Assisted routing ($L = 5$) is 5 times lower than this of Direct Transmission, i.e. $R = 0.2$, when mobility becomes very heterogeneous ($CV_\lambda = 2.5$), delay is improved less than 1.25 times, i.e. $R \geq 0.8$.

In Fig. 4 and 5 we present more simulation results for R and $P_{(src.)}$ under different scenarios. In all the scenarios the results are consistent with our theory. We can see that as Table 2 predicts, for very high values of mobility and traffic heterogeneity Direct Transmission can achieve comparable performance with Relay-Assisted routing, i.e. R and $P_{(src.)}$ become almost 1.

Till now, we presented the two metrics (R and $P_{(src.)}$) for the *relative* performance of Direct Transmission and Relay-Assisted routing. However, our analysis allows to calculate *absolute* values too (Eq. (6) and Eq. (12)). In Fig. 6 and Fig. 7 we present the message delivery delay under Direct Transmission, $E[T_{DT}]$, and Relay-Assisted routing (in our simulations we used the Spray and Wait routing protocol [5]), $E[T_R]$, under scenarios with varying mobility and traffic heterogeneity. From these plots, it becomes evident that the previously observed changes

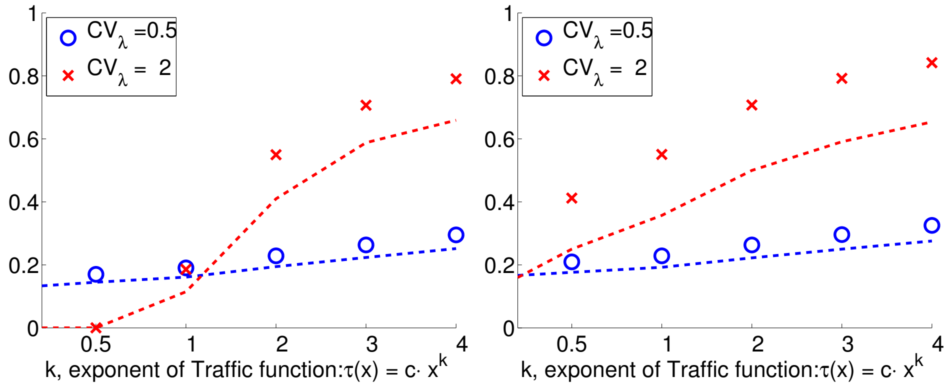

(a) $R = \frac{E[T_R]}{E[T_{DT}]}$ (b) $P_{(src.)}$

Figure 3: Comparison of Direct Transmission and Relay Assisted routing in two mobility scenarios and under varying traffic heterogeneity: (a) The ratio R and the theoretical bound $Rmin$, (b) The probability $P_{(src.)}$ and the theoretical bound $Pmin$. Simulation results are denoted with markers and theoretical predictions (lower bounds) with lines.

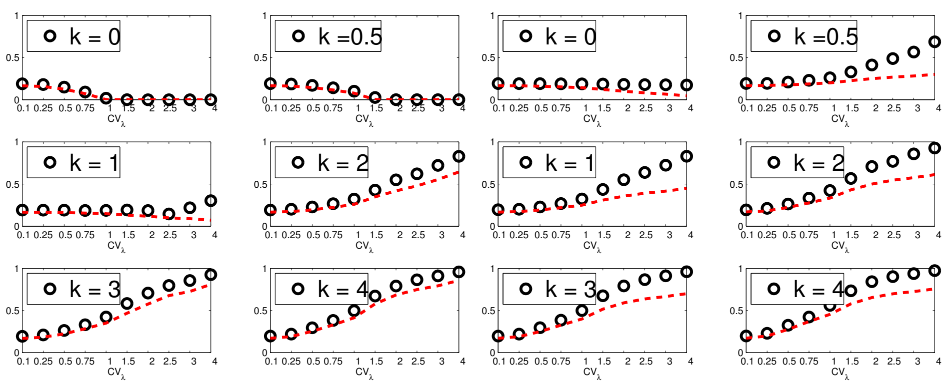

(a) $R = \frac{E[T_R]}{E[T_{DT}]}$ (b) $P_{(src.)}$

Figure 4: Comparison of Direct Transmission and Relay Assisted routing for different traffic scenarios ($\tau(x) = c \cdot x^k$) under varying mobility heterogeneity: (a) The ratio R and the theoretical bound $Rmin$, (b) The probability $P_{(src.)}$ and the theoretical bound $Pmin$. Simulation results are denoted with markers and theoretical predictions (lower bounds) with lines.

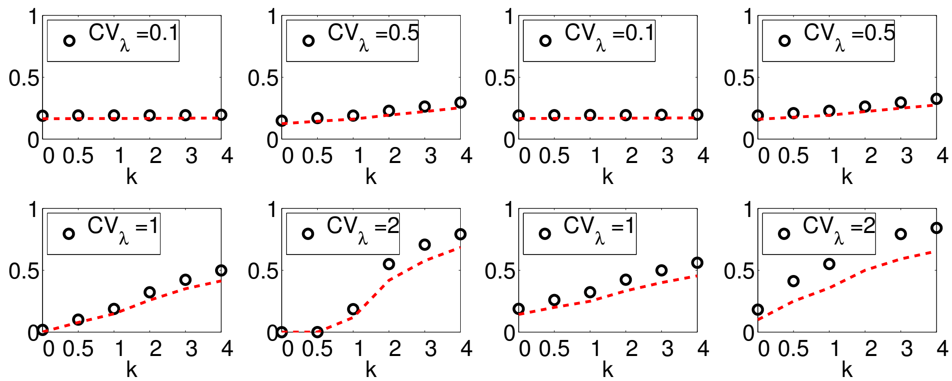

(a) $R = \frac{E[T_R]}{E[T_{DT}]}$ (b) $P_{(src.)}$

Figure 5: Comparison of Direct Transmission and Relay Assisted routing for different mobility scenarios (CV_λ) under varying mobility heterogeneity: (a) The ratio R and the theoretical bound $Rmin$, (b) The probability $P_{(src.)}$ and the theoretical bound $Pmin$. Simulation results are denoted with markers and theoretical predictions (lower bounds) with lines.

in relative performance (R and $P_{(src.)}$) are mainly due to the decrease of Direct Transmission delivery delay. In contrast, the delay of Relay-Assisted routing is less affected by heterogeneity.

An important observation is that for scenarios with low traffic heterogeneity(i.e. $k < 1$), as the mobility heterogeneity increases, Direct Transmission becomes inefficient and its expected delivery delay can become very large. This shows clearly the reasons why previous studies, which assumed homogeneous traffic patterns (i.e. $k = 0$), found Direct Transmission to be inefficient. However, this does not hold under any conditions. Under heterogeneous traffic, the perfomance of Direct Transmission improves significantly and this indicates that Direct Transmission can be an efficient solution under scenarios/applications where users that contact frequently tend to exchange more messages.

Moreover, considering the accuracy of our predictions (Eq. (6) and Eq. (12)), Table 3 contains the relative error between our prediction (Eq. (6)) and the simulation results for the exact delay of Direct Transmission, and the tightness of the predicted bound (Eq. (12)) for the Relay-Assisted routing delivery delay. We can see that the accuracy is significant and increases with the network size. This is because the condition $L \ll N$ we assumed in our analysis, holds better for larger networks.

4.2 Real-World Networks

After validating our results in synthetic scenarios, in order to investigate to what extent they can be applied in real-world networks, we conduct simulations on datasets

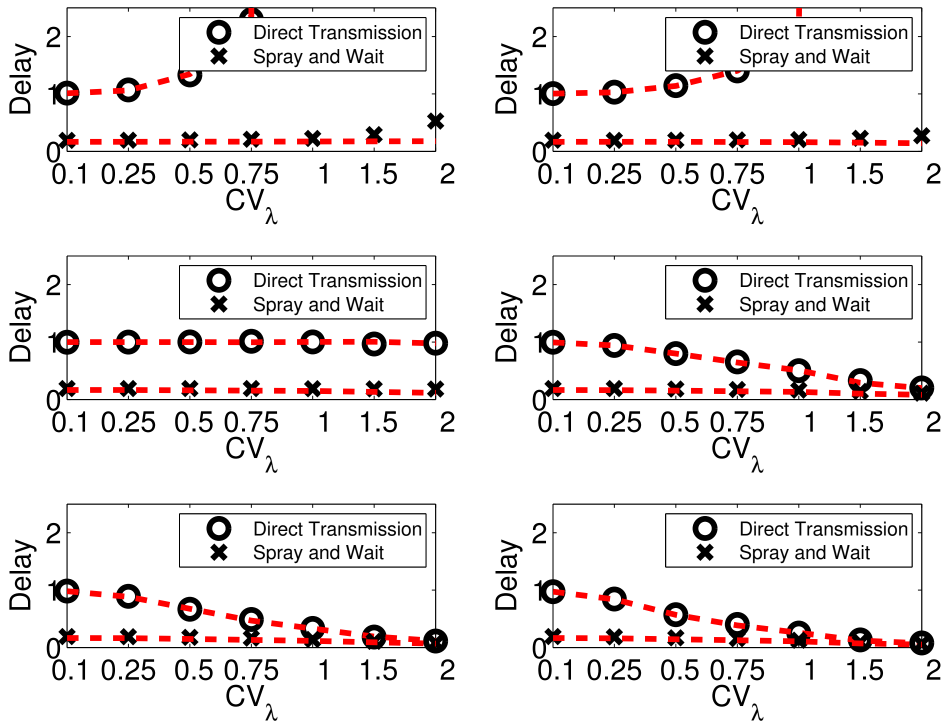

Figure 6: Expected message delivery delay of Direct Transmission and Relay-Assisted (Spray and Wait) routing. Scenarios with different traffic functions, i.e. $k = 0, 0.5, 1, 2, 3, 4$ (from left to right and up to down), and under varying mobility heterogeneity (CV_λ).

Table 3: Delay Predictions: Average Relative Differences

	Synthetic Scenarios	
	$N = 1000$	$N = 100$
Direct Transmission	0.2%	0.3%
Spray and Wait	1.4%	13.9%

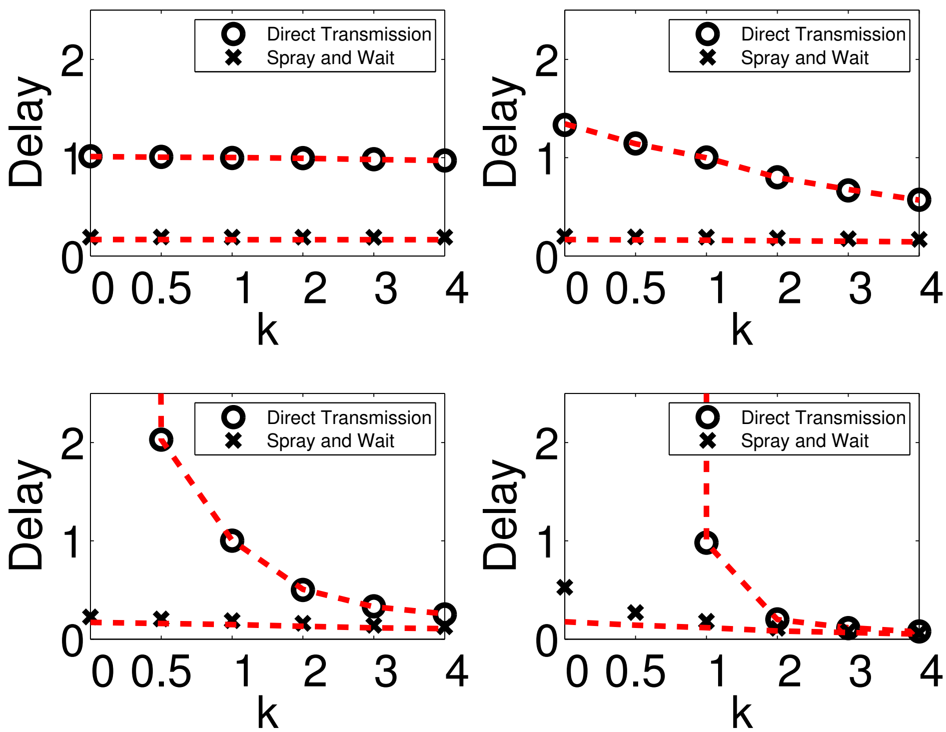

Figure 7: Expected message delivery delay of Direct Transmission and Relay-Assisted (Spray and Wait) routing. Scenarios with different mobility heterogeneity, i.e. $CV_\lambda = 0.1, 0.5, 1, 2$ (from left to right and up to down), and under varying traffic heterogeneity (k).

Table 4: Datasets Information

Dataset	Network Size	Contacts	Traffic
Gowalla / Twitter	(AU) 1004	Check-ins	Tweets
	(SF) 479		
Strathclyde	24	Bluetooth Proximity	Calls / SMS

collected from online social networks (*Gowalla / Twitter* dataset [11]) and a mobile phone usage experiment (*Strathclyde* dataset [30]). In following we decribe some main features of the datasets, which are also presented in Table 4.

Gowalla / Twitter dataset

Gowalla was[5] a location-based social network, where users were able to *check-in* at "spots" (bars, shops etc.) through their mobile phones. In addition, a user could connect her Gowalla account to her Twitter account. Hence, from this dataset, we could retrieve information related both to nodes' mobility (Gowalla *check-ins*) and communication traffic (*tweets*).

Mobility

Considering node mobility, our framework focuses on the contact events between nodes, because message exchanges can take place only during these events. Thus, in the dataset, we consider as a contact event the time when two users reside in the same "spot" simultaneously[6]. The contact rates λ_{ij} can be computed from the number of the contact events and the inter-contact time intervals. Then, to incorporate this information in our model, we fitted the contact rates distribution f_λ with a known probability distribution $\hat{f}_\lambda$. Specifically, in the two cities, Austin (AU) and San Francisco (SF), for which we have the most users' records (1004 and 479 nodes, respectively), the experimental CCDF (*complementary cumulative distribution function*) of the contact rates λ_{ij} can be approximated by a straight line on a log-log axes plot. This implies that f_λ can be approximated with a *Pareto* distribution (Fig 8(a) and Fig. 8(b)).

Communication Traffic

As an indication for the communication traffic that two nodes would exchange in an opportunistic network, we use the number of *tweets* ($\#tweets$) in which they are both involved. Hence, for each pair $\{i, j\}$ we set its traffic rate τ_{ij} equal to the number of tweets it exchanges, i.e. $\tau_{ij} = \#tweets_{ij}$. Then, we approximate

[5]It was launched in 2007 and closed in 2012.

[6]Since Gowalla users only check in and do not check out, we cannot infer directly this information. Therefore, following the methodology of [11], we assumed that each user remains at a spot he visited for 1 hour

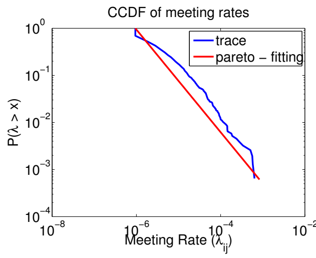

(a) Austin

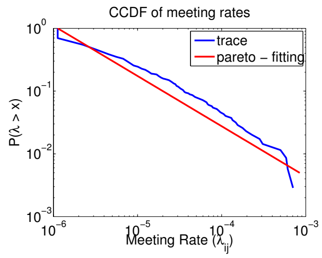

(b) San Francisco

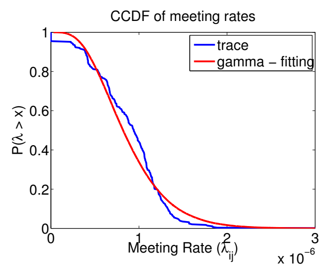

(c) Strathclyde

Figure 8: Fitting contact rates CCDF $\overline{F}_\lambda$ for the Gowalla (a) Austin, (b) San Francisco, and (c) Strathclyde dataset.

Table 5: Fitting traffic functions for the Gowalla dataset

Scenarios:	S1	S2	S3
τ_{ij}	$\sqrt{\#tweets_{ij}}$	$\#tweets_{ij}$	$(\#tweets_{ij})^2$
$\hat{\tau}(x)$ (AU)	$c \cdot x^{0.6}$	$c \cdot x^{0.83}$	$c \cdot x^{0.79}$
$\hat{\tau}(x)$ (SF)	$c \cdot x^{0.31}$	$c \cdot x^{0.35}$	$c \cdot x^{0.37}$

the observed relation between traffic and contact rates ($\tau_{ij} \sim \lambda_{ij}$) with a function $\hat{\tau}(x)$, in order to use it in our theoretical expressions. Furthermore, to investigate more possible correlations between the opportunistic traffic (τ_{ij}) and the Twitter traffic ($\#tweets$), we create scenarios where we set $\tau_{ij} = \sqrt{\#tweets_{ij}}$ and $\tau_{ij} = (\#tweets_{ij})^2$. The approximative functions $\hat{\tau}(x)$ for each scenario are presented in Table 5, where we can see that $\hat{\tau}(x)$ are of the form $c \cdot x^k$ with exponent $k < 1$.

Strathclyde dataset

The Strathclyde dataset was collected in an experiment, in which 27[7] high school students were selected and given modified smartphones, which recorded proximity events (through Bluetooth), calls and sms exchanged between the phone user and the other participants.

[7] Unfortunately, due to malfunctioning devices, only the data of 24 users are reliable.

Mobility

In this dataset the meeting events were already recorded and, thus, we did not have to preprocess the data as in the Gowalla dataset. We followed the same methodology to calculate the contact rates λ_{ij} and fit their distribution with a function $\hat{f}_\lambda$, where $\hat{f}_\lambda$ is here a *Gamma* distribution (Fig. 8(c)).

Communication traffic

We consider three scenarios, in each of which we use a different communication traffic metric: (i) total number of calls and sms, i.e. $\tau_{ij} = \#calls_{ij} + \#sms_{ij}$, (ii) total duration of calls, i.e. $\tau_{ij} = T_{ij}^{calls}$, and (iii) total length of sms (in characters), i.e. $\tau_{ij} = L_{ij}^{sms}$. For each scenario, we fitted a function $\hat{\tau}(x)$ as before, through the relation $\tau_{ij} \sim \lambda_{ij}$. The traffic functions are of the form $\hat{\tau}(x) = c \cdot x^k \cdot e^{-\beta \cdot x}$.

Simulations

In both datasets and for each traffic scenario, we generate 10000 messages at random time points, choosing each time the source - destination pair according to the weights τ_{ij}. We consider both Direct Transmission and Spray and Wait routing [5] with $L = 2, 5, 10, 20$ copies per message. In the analytical expressions we use the fitted functions $\hat{f}_\lambda(x)$ and $\hat{\tau}(x)$.

In Fig. 9 we present the ratio R, the probability $P_{(src.)}$ and the corresponding theoretical lower bounds R_{min} and P_{min}. We consider *homogeneous* and *heterogenous* (denoted with $*$) traffic scenarios in the Gowalla/Twitter (*AU* and *SF*) and Strathclyde (*St*) datasets. Two main observations that confirm (both quantitatively and qualitatively) the validity of our results are: (i) the theoretical bounds are always lower than simulation results, and (ii) the predictions follow the tendency (increasing with traffic heterogeneity) of real networks' behavior.

To demonstrate further to what extent our results can capture the effect of traffic heterogeneity in real scenarios, in Table 6 we compare the theoretical predictions (R_{min} and P_{min}) between pairs of *different* simulation scenarios for the Gowalla/Twitter dataset[8]. Specifically, in cases where the performance improves (or degrades), if the theoretical results also indicate a performance improvement (or degradation), the prediction is assumed to be correct. For example, in the scenarios *AU-S1* and *SF-S3* the ratios R are $R^{(AU-S1)} = 0.89$ and $R^{(SF-S3)} = 0.94$, i.e. $R^{(AU-S1)} < R^{(SF-S3)}$. For the theoretical predictions it holds also that $R_{min}^{(AU-S1)} = 0.64 < R_{min}^{(SF-S3)} = 0.68$ and, thus, the prediction is assumed to be correct. The "correct" predictions are denoted with ✓ and the "incorrect" with ×. The elements above the diagonal refer to the ratio R (and R_{min}), whether the lower triangular part refers to the probability $P_{(src.)}$ (and P_{min}) predictions.

[8] We denote with *S1*, *S2* and *S3* the corresponding scenarios presented in Table 5 and with *HOM* the scenarios with homogeneous traffic. In each city, we chose the scenarios with the highest difference in traffic heterogeneity, k.

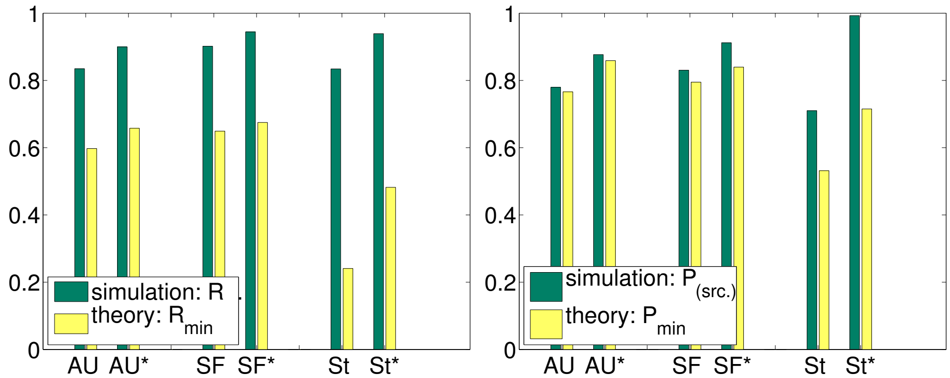

(a) $R = \frac{E[T_R]}{E[T_{DT}]}$ and R_{min} (b) $P_{(src.)}$ and P_{min}

Figure 9: Simulated and Theoretically predicted R and $P_{(src.)}$ for homogeneous and heterogeneous ($*$) traffic scenarios on the datasets

Table 6: Comparison of predictions for the metrics R and $P_{(src.)}$ between different scenarios on the Gowalla dataset

$*$	R	AU			SF		
$P_{(src.)}$	$*$	HOM	S1	S2	HOM	S1	S3
AU	HOM	$*$	✓	✓	✓	✓	✓
	S1	✓	$*$	✓	✓	✓	✓
	S2	✓	✓	$*$	×	✓	✓
SF	HOM	✓	✓	✓	$*$	✓	✓
	S1	✓	✓	×	✓	$*$	✓
	S3	✓	✓	×	✓	✓	$*$

Table 7: Comparison of predictions for the metrics R and $P_{(src.)}$ between different scenarios on the Strathclyde dataset

*	R	St		
$P_{(src.)}$	*	S1	S2	S3
	S1	*	✓	✓
St	S2	✓	*	✓
	S2	✓	✓	*

It is evident that in the majority of the cases we consider, the theoretical results can capture the relative changes in network performance, even between different environments (i.e. between *AU* and *SF*). Moreover, in the Starthclyde dataset, *all* the respective comparisons were found to be correct (✓), as it can be seen in Table 7.

Finally, Table 8 gives the relative difference between our theoretical predictions (Eq. (6) and Eq. (12)) and the simulation results for the *expected message delivery delay* of Direct Transmission and Spray and Wait routing. We consider all the scenarios presented in Table 5 for the city of Austin (*AU*) and San Francisco (*SF*) from the Gowalla/Twitter dataset. As it can be seen, the accuracy of the predictions for $E[T_d]$ (Direct Transmission) is significant for most of the scenarios. The errors for $E[T_r]$ (Spray and Wait) in most cases are higher than in Direct Transmission, though still satistying.

Table 8: Relative Difference of Expected Delivery Delay Predictions in the Gowalla dataset

	AU			SF		
	S1	S2	S3	S1	S2	S3
Direct Transmission	18%	13%	1%	8%	6%	8%
Spray and Wait	16%	36%	29%	22%	24%	23%

5 Extensions

We have tried to present our results in the context of simple schemes (e.g. unicast traffic, random relay selection), to keep analysis tractable and illustrate key principles. In this section, we discuss how our framework could be applied in some additional cases. Although far from complete, we believe this set of examples, further underlines the utility of our analysis.

Utility-based schemes are often used to select good relays for the intended replicas, rather than picking random ones, e.g. [6, 31, 32]. Although one could modify our Results 1, 2 and 3 to capture such schemes, e.g. by multiplying L

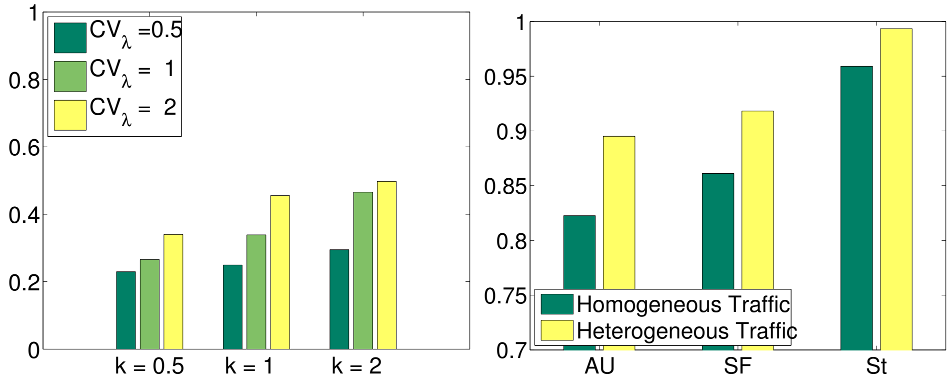

(a) Synthetic scenarios

(b) Real networks

Figure 10: $P_{(src.)}$ of utility-based routing in (a) synthetic scenarios with varying mobility (CV_λ) and traffic heterogeneity (k), and (b) real networks with homogeneous and heterogeneous traffic.

with an appropriate factor c_u (the actual improvement depending on the contact graph, and the protocol), it is beyond the scope of this paper to perform the detailed analysis. Instead, we demonstrate some preliminary simulation results suggesting that our conclusions hold also for utility-based routing.

We use a variant of the protocol presented in [32], in which the higher the contact rate λ_{id} between a node i and the destination d, the higher the probability that node i is selected as a relay. In Fig. 10(a) we present simulation results for the delivery metric $P_{(src.)}$ on synthetic scenarios with varying mobility (CV_λ) and traffic (k) heterogeneity. Similarly to the random replication case, for increasing heterogeneity (in mobility and/or traffic) the gain of the extra copies clearly decreases (i.e. $P_{(src.)}$ increases) under utility-based schemes. In Fig. 10(b) we compare the probability $P_{(src.)}$ of scenarios with and without traffic heterogeneity in real networks. As before, the results are in agreement with our theory, i.e. the performance of Direct Transmission improves even compared to protocols that use more sophisticated techniques for relay selection. In these scenarios, similar results hold also for the ratio R which are presented in Fig. 11.

We have also been assuming unicast messages between a $\{s, d\}$ pair. However, our results apply also to multicast [13] or anycast (e.g. content sharing or service composition applications) [2] messages from s, with d being one of the destinations, since similar mechanism are often used for their dissemination. To demonstrate this, in Table 9 we present simulation results for two multicast scenarios, with homogeneous (*HOM*) and heterogeneous (*HET*) traffic ($\tau(x) = c \cdot x^4$), under varying mobility heterogeneity. A source sends messages to 5 destinations (each selected with a probability $\propto \tau_{ij}$) either by Direct Transmission or by Relay-Assisted routing with $L = 5$ copies. As delivery delay, we consider the delay till all the destinations get the message. It is evident that R and $P_{(src.)}$ (i) increase signif-

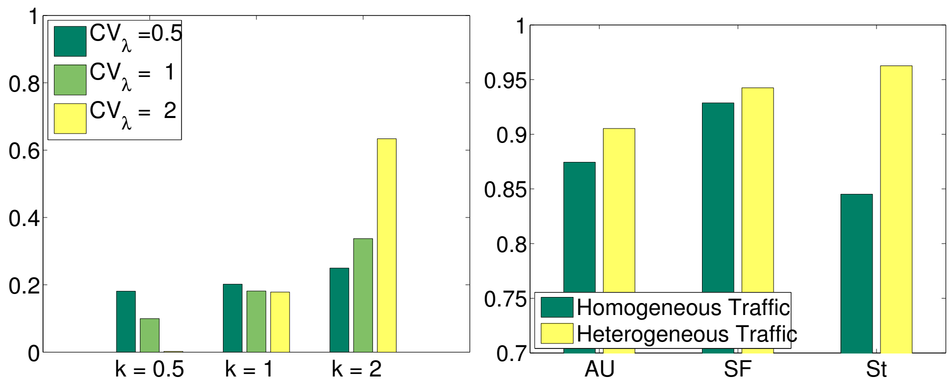

(a) Synthetic scenarios

(b) Real networks

Figure 11: Ratio R of utility-based routing in (a) synthetic scenarios with varying mobility (CV_λ) and traffic heterogeneity (k), and (b) real networks with homogeneous and heterogeneous traffic.

Table 9: Multicast Communication

	CV_λ	0.1	0.5	1	1.5	2
HOM	R	0.18	0.12	0.01	0	0
	$P_{(src.)}$	0.01	0.01	0.01	0.01	0.01
HET	R	0.18	0.26	0.39	0.52	0.61
	$P_{(src.)}$	0.01	0.03	0.12	0.26	0.41

icantly with mobility heterogeneity when traffic is heterogeneous, and (ii) become much larger compared to the homogeneous case (where R decreases and $P_{(src.)}$ is constant), which is in agreement with our results.

6 Related Work

Useful implications for opportunistic networking have arisen from the investigation of *mobility/social ties* and *social ties/communication traffic* correlations, which have been studied extensively and under different disciplines, like anthropology [24], sociology [33], social media [21] or pervasive social networks [9]. For example, [21] shown that the amount of exchanged communication traffic between users of OSNs depends on their social relationships. On the other hand, the *communication traffic/mobility* correlation has not been given similar attention. There exist only a few works [11, 12] studying it in a framework relevant to opportunistic networking. In [11], Hossmann *et al.* collected and analysed two datasets from online social networks (Facebook and Gowalla / Twitter), and investigated the relations among three dimensions: *mobility, social ties, communication traffic*. With respect to our study, they found that there is strong dependence between mobility and traffic, and, specifically, node pairs that contact during the experiments' duration, communicate with higher probability than the other pairs. Correspondingly, authors in [12] analysed a massive dataset of Call Detail Records (CDRs) of 6 million users and shown a positive correlation between the mobility and communication traffic patterns. Not only they shown that the higher the contact rate (λ_{ij}) of a node pair is, the higher the probability that the nodes communicate intensively, but also found that information inferred by the mobility patterns can work as a good predictor for future communication events.

However, despite the fact that [11, 12] show clearly that communication traffic is heterogeneous (and correlated to mobility), to our best knowledge, its *effects* on information dissemination mechanisms have not been studied previously.

7 Conclusion

Motivated by (i) recent findings indicating heterogeneous traffic patterns in mobile social networks and (ii) the lack of related studies, in this paper, we modelled traffic heterogeneity and studied how it affects the performance in opportunistic networking. We found that the effects can be significant, changing our understanding of common design principles in DTNs, such as the added value of relays. This indicates a necessity for revisiting the evaluation of protocols in scenarios that entail diversity in the traffic exchanged between different nodes. Moreover, our results seem to have some interesting implications about the usefulness of opportunistic networking for end-to-end applications.

Based on the initial understanding offered by our analysis, in future work we intend to investigate further traffic characteristics and study the implications of traffic heterogeneity for content-centric applications.

Acknowledgment

We would like to thank the authors of [11] for providing us the dataset they have collected, which was essential for the validation of our model and results in real networks' scenarios.

References

[1] L. Pelusi, A. Passarella, and M. Conti, "Opportunistic networking: data forwarding in disconnected mobile ad hoc networks," *Comm. Mag., IEEE*, vol. 44, pp. 134–141, Nov. 2006.

[2] M. Pitkänen, T. Kärkkäinen, and et al., "SCAMPI: service platform for social aware mobile and pervasive computing," *ACM Comput. Commun. Rev.*, vol. 42, pp. 503–508, Sept. 2012.

[3] B. Han, P. Hui, V. Kumar, M. Marathe, J. Shao, and A. Srinivasan, "Mobile data offloading through opportunistic communications and social participation," *IEEE Trans. on Mob. Comp.,*, vol. 11, no. 5, 2012.

[4] 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, no. 1, 2008.

[5] 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, no. 1, 2008.

[6] P. Hui, J. Crowcroft, and E. Yoneki, "Bubble rap: Social-based forwarding in delay-tolerant networks," *IEEE Transactions on Mobile Computing*, vol. 10, pp. 1576–1589, 2011.

[7] T. Spyropoulos, R. N. Rais, T. Turletti, K. Obraczka, and A. Vasilakos, "Routing for disruption tolerant networks: taxonomy and design," *Wirel. Netw.*, vol. 16, pp. 2349–2370, Nov. 2010.

[8] 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.

[9] 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, pp. 2201–2217, 2012.

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

[11] T. Hossmann, G. Nomikos, T. Spyropoulos, and F. Legendre, "Collection and analysis of multi-dimensional network data for opportunistic networking research," *Comput. Communications*, vol. 35, no. 13, 2012.

[12] D. Wang, D. Pedreschi, C. Song, F. Giannotti, and A.-L. Barabasi, "Human mobility, social ties, and link prediction," in *Proc. ACM KDD*, 2011.

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

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

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

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

[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] 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.

[19] 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.

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

[22] J. Ott, E. Hyytiä, P. Lassila, J. Kangasharju, and S. Santra, “Floating content for probabilistic information sharing,” *Pervasive Mob. Comput.*, vol. 7, pp. 671–689, Dec. 2011.

[23] A.-K. Pietiläinen, E. Oliver, J. LeBrun, G. Varghese, and C. Diot, “Mobiclique: middleware for mobile social networking,” in *ACM Proc. WOSN*, 2009.

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

[25] A. Gut, *An intermediate course in probability*. Springer Verlag, 2nd ed., 2009.

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

[27] M. Grossglauser and D. Tse, “Mobility increases the capacity of ad-hoc wireless networks,” in *Proc. IEEE INFOCOM*, 2001.

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

[29] S. M. Ross, *Introduction to Probability Models*. Academic Press, Elsevier, 9 ed., 2007.

[30] A. McDiarmid, J. Irvine, S. Bell, and J. Banford, “CRAWDAD data set strath/nodobo (v. 2011-03-23).” Downloaded from http://crawdad.cs.dartmouth.edu/strath/nodobo, Mar. 2011.

[31] P. Hui and J. Crowcroft, "How small labels create big improvements," in *IEEE PerCom Workshops*, 2007.

[32] S. Nelson, M. Bakht, and R. Kravets, "Encounter-based routing in dtns," in *Proc. IEEE INFOCOM*, 2009.

[33] N. Eagle and A. Pentland, "Inferring social network structure using mobile phone data," in *Proceedings of the National Academy of Sciences (PNAS)*, vol. Vol 106 (36), pp. 15274–15278, 2009.