Research Article

A Sociability-Based Routing Scheme for Delay-Tolerant Networks

Flavio Fabbri and Roberto Verdone

*Dipartimento di Elettronica, Informatica e Sistemistica (DEIS), WiLAB, University of Bologna, Viale Risorgimento 2,
40136 Bologna, Italy*

Correspondence should be addressed to Flavio Fabbri, flavio.fabbri@unibo.it

The problem of choosing the best forwarders in Delay-Tolerant Networks (DTNs) is crucial for minimizing the delay in packet
delivery and for keeping the amount of generated traffic under control. In this paper, we introduce *sociable routing*, a novel routing
strategy that selects a subset of optimal forwarders among all the nodes and relies on them for an efficient delivery. The key idea
is that of assigning to each network node a time-varying scalar parameter which captures its social behavior in terms of frequency
and types of encounters. This sociability concept is widely discussed and mathematically formalized. Simulation results of a DTN
of vehicles in urban environment, driven by real mobility traces, and employing sociable routing, is presented. Encouraging results
show that sociable routing, compared to other known protocols, achieves a good compromise in terms of delay performance and
amount of generated traffic.

1. Introduction

This paper introduces *sociable routing*, a novel routing
scheme for Delay-Tolerant Networks (DTNs) [1] and pro-
poses its evaluation and assessment with respect to other
existing protocols. The key idea of sociable routing is to
solve the routing problem in DTNs [2] by assigning to
each network node a time-varying scalar parameter, called
sociability indicator, depending on its social behavior, that
has to do with the frequency and type of node's encounters.
Then, each node forwards its data packets only to the most
sociable nodes. Thus, the chances of reaching the intended
endpoint are maximized and the amount of transmissions
kept under control.

After giving a detailed formalization of the sociability
concept, we simulate packet transmissions in a DTN in an
urban context. In particular, we propose a case study where
nodes are vehicles moving according to real traffic traces
[3]. Encouraging results show that sociable routing achieves
a good compromise in terms of delay performance and
amount of generated traffic. Along with result discussion,
we also mention some issues that are still open and discuss
possible improvements.

The main contribution of this paper is the formalization
of a sociability concept and a guideline to its exploitation
for efficient forwarding in DTNs. Additionally, a framework
for its evaluation and comparison with other schemes is also
presented.

In a typical DTN, nodes are mobile and of the same type,
they have wireless communication as well as buffering capa-
bilities. However, they can communicate and exchange data
only if they are within a certain distance, commonly called
transmission range. In a standard scenario, the transmission
range is small compared to network size. For this reason,
the network is most of the time partitioned and source-to-
destination paths do not exist. Nonetheless, the appearance
of new links when old ones break due to mobility, together
with a store-and-forward paradigm, can still make packet
delivery possible.

In the DTN jargon, data packets are referred to as *bundles*
[1], since it is often assumed that an overlay layer, called
bundle layer, is present above the existing protocol stack for
supporting interoperability.

The problem of routing in DTNs has recently deserved
a growing attention [2, 4–7]. When no information on
nodes schedule is available, epidemic routing, which basically

implies flooding, seems to be the only possible approach [8]. However, this reveals practically unfeasible because of the generated traffic which grows exponentially. The opposite approach consists of a perfect scheduling of transmissions, which requires deterministic knowledge of nodes behavior, as in the interplanetary paradigm [9]. In most of the cases, either partial information on nodes contacts and mobility patterns is available, or they possess some intelligence allowing them to learn such information and adapt their routing criteria consequently. When nodes are position-aware and can learn and share their mobility patterns, a solution is found in MobySpace routing [10], which uses the fact that nodes with similar patterns are likely to meet up. Other approaches consider the problem from a social perspective as, for example, in [5], where the authors introduce SimBet Routing, a strategy that exploits the notion of *centrality*. In [11], a general framework for context-aware adaptive routing in DTNs, called CAR, is proposed. CAR makes use of Kalman filter-based prediction techniques and utility theory in order to select the best carrier for a message. In [12], Hui et al. introduce *BUBBLE rap*, a forwarding algorithm based on social information and suitable for Pocket Switched Networks (PSNs). The authors ground their work on the concept of *community*, assuming each individual is doubly ranked based on its popularity in both the whole community and its local community. The ranking is based on the notion of *centrality*. Such an approach surely catches and exploits cooperation binds in people networks but is not easily applicable, for example, to vehicular networks, where communities are not so clearly definable.

Alternative scenarios envision the presence of additional nodes whose mobility can be controlled in order to maximize the amount of deliveries. In [13], data ferries are extra nodes whose paths are optimized based on a delay constraint. In [14], cars act as data mules and employ a carry-and-forward paradigm to transfer data packets to a portal. Finally, in [15], opportunistic data delivery is studied when both traditional routing and data mules techniques are jointly used.

Sociable routing can also be thought of as a protocol inspired by the concept of network opportunism [16], where resources offered by different nodes/networks are jointly exploited according to the needs of a specific application task: in such a vision, the sociability degree of a node is an information offered to the community.

The rest of the paper is structured as follows. In Section 2, the sociability concept and the core of sociable routing are introduced and discussed in detail. The model for the computation of sociability indicators is reported in Section 3. Then, in Section 4, the simulator is described, the performance metrics and results are shown and discussed. Finally, Section 5 reports concluding remarks and ideas for future work.

2. Sociability Concept

Our basic idea is that nodes having a high degree of sociability (i.e., frequently encounter many different nodes) are good candidate forwarders. Applying this simple rule to a delay-tolerant network is quite straightforward. As first step,

one needs to observe nodes behavior and learn their habits. Then, a synthetic scalar parameter will be assigned to each node depending on its social behavior. Finally, routing from a source to a destination node is performed by forwarding bundles to a restricted set of relays which show a high degree of sociability and, thus, are very likely to get in touch with all possible endpoints.

One relevant assumption that we need is the periodicity of behaviors, meaning that it is possible to make predictions on the social conduct of a node based on what has been observed before. Roughly speaking, we expect those nodes that showed very high sociability over a time period of a certain length to behave accordingly in the future for a period of at least the same length. This is a reasonable hypothesis in population networks [17], and we believe it still is in all scenarios where the mobility of nodes is governed by human behavior, as in vehicular networks, pedestrian networks, and so forth.

In this section, we illustrate the notion of sociability in more detail. In particular we give a mathematical characterization of it, showing how such information can be exploited by the nodes for enhancing routing performance and how it can be obtained.

2.1. Modeling. The way in which the social features are modeled should be very simple, on the one hand, in order for the nodes to produce and exchange such information in an inexpensive manner. On the other hand, the challenge stands in capturing as much as possible of the exploitable information in a single parameter, that we will call *sociability indicator*.

One way sociability could be quantified is by looking at the intercontact information of each node [18, 19]. In particular, the intercontact time analysis reveals how frequently a node meets with one another. As an example, an indication on the average intercontact time of a node with any other could give a rough idea of its social behavior. However, in the latter case, one can appear very sociable by having frequent meetings with a very restricted set of neighbors. Unfortunately, this does not make it a good candidate forwarder.

Moreover, an important aspect to be captured in analogy with human relationships, is that one person who only meets a single friend, the latter being very sociable, can itself be considered sociable. Turning to an information network perspective, a node being isolated most of the time with very sporadic links to a single neighbor, may appear very unsociable. Nonetheless, if the neighbor is very sociable and can reach many destinations, then the former node may also have chances to send its bundles to many destinations through a 2-hop path. As a consequence, the presence of sociable neighbors is an important addendum that should be incorporated into the sociability indicator of one node.

Intuitively, it is a natural assumption that mobility patterns of nodes are related to their social behavior. In fact, if a node visits a great number of different locations in a short time, it is likely to meet many others. Although this is true to some extent, there are plenty of scenarios where the concentration of users is not constant in space

(e.g., the union of a city center with its suburbs). Hence, the mere covering large distances does not necessarily result in high forwarding opportunities. For this reason, in order to maintain the overall idea detached to any specific environment, we chose not to include any direct information regarding mobility patterns in the sociability indicator. An important advantage of this approach is that no information on nodes position is ever required (see Section 2.2).

In [10], the authors state that two people having similar mobility patterns (in terms of frequency of visits to specific locations) are more likely to meet each other, thus to be able to communicate. Then, they recognize that the main limitation of the previous statement is that even though two people visit the same locations, they do not necessarily do it synchronously. Thus, two such nodes may never be in the range of each other. This is not a rare event, especially at urban scale. Consider, for example, a public transportation fleet (e.g., buses). Two buses running on the same route have the exact same mobility patterns. However, if one follows the other few kilometers behind, they never reach each other. More generally, there are places like, for example, a big mall, that many people periodically visit at different time. This results in some similarity of their patterns which does not necessarily reflect meeting opportunities.

It is worth mentioning that sociability indicators are not based on the notion of *communities* (as in [12]), that is, groups of individuals that "stay in touch" for prolonged time due to shared habits, behaviors, believes, and so forth. The reason of this is that our approach is more vehicular traffic oriented and thus try to cope with a highly dynamical environment. This results in the fact that (i) we do not aim at identifying such groups of people; (ii) we do not keep track of contact duration, since it is not equally relevant in all types of networks (e.g., in vehicular networks contacts are all rather short in the majority of cases).

Finally, we emphasize that the sociability indicator only highlights what are the best forwarders in a given time period, in the sense of those having the highest degree of sociability. As a consequence, this information is not related to a specific destination to be reached but it is instead absolute. This descends from avoiding a sociability characterization based on mobility patterns and is consistent with the intent of minimizing the exchange of data. This also implies that no prior knowledge of the destination (e.g., its position, sociability indicator, etc.) is requested at the source.

A hybrid concept considering a mixture of sociability and mobility information could be evaluated in future studies. In Sections 3.1 and 3.2 we report a formal definition of the sociability indicator for the cases where only directed contacts enhance sociability, and also multihop contacts are considered, respectively.

2.2. Acquisition. Since we do not use information on positions, nodes are not requested to adopt any positioning technique, nor do they have to learn their mobility patterns as in [10]. The two main issues arising with the use of sociable routing are (i) how a node learns its own social behavior and (ii) how it communicates its social behavior to other nodes.

Note that the two issue are strictly connected, as a node needs to know the social behavior of its neighbors in order to derive its own. For this reason, a distributed strategy where nodes, upon encounters, update their own sociability parameter through the exchange of a minimum amount of data, could be the optimum. For example, the sociability updates could be appended to data bundles in order not to overwhelm the network with signaling information. Although this is not addressed here, since our aim is primarily that of presenting and validating the general idea at the base of sociable routing, we give a rough indication of the cost of acquiring *sociability indicators*.

Consider a network of N nodes at an initial state where no one knows its *sociability indicator*. This number is computed on the basis of the frequency and amount of encounters of a node. Thus, we can assume the ith node receives identity information from every encountered node. We will then estimate the scaling law for such transmissions. Let us denote as n_i the number of encounters of the ith node, $i = 1,\ldots,N$. On average, an arbitrary node has $\mathbb{E}\{n_i\} = \bar{n}$ encounters (i.e., transmissions/receptions) over a certain time period. From network perspective, the average number of exchanges is $K_n \propto N \cdot \bar{n}$. Now, $\bar{n}$ is a function of several parameters. In particular, $\bar{n} \propto T \cdot \bar{v} \cdot \rho$, where T is the observation period, $\bar{v}$ the average speed of nodes and ρ the density of nodes, seems a reasonable assumption. Moreover, ρ is in turn proportional to N. This yields the conclusion that $K_n = O(N^2)$.

In a successive step, when each node has computed a first estimation of *sociability indicator*, the exchange continues in such a way that the ith node receives from the neighbors not just their identities but also their *sociability indicators*, which are used for refining the estimation of its own. However, this has no effect on the above mentioned scaling law.

Hence, in the following we assume that nodes have knowledge of their social behavior referred to a specific time window. In particular, the analysis carried out in Section 3 is based on a centric perspective for the sake of mathematical treatment, without loss of generality, due to the feasibility of a distributed strategy at reasonable cost, as roughly discussed above.

2.3. Usage. As previously mentioned, the basic idea is to select a set of sociable nodes that can potentially reach any endpoint. This set should be kept small enough to avoid useless transmissions. To this end, the following strategy can be adopted. A node takes its routing decision at a given time t by (i) evaluating the sociability indicators of the current neighbors; (ii) comparing them to its own and (iii) choosing as forwarders a maximum of N_f nodes that have greater sociability than its. This simple scheme allows to limit the number of bundle transmissions at each encounter by setting a maximum, N_f. Moreover, a node does not transmit any bundle if it does not meet any more sociable node. As a further implication, when a bundle is generated by a node with low sociability degree, a large number of transmissions are permitted, since the source will certainly meet more sociable nodes. In fact the network copes with lack of encounters by generating multiple replicas of the

```
(1)   R_k(t) := ∅
(2)   if b ∈ W_k(t) then
(3)       R_k(t) = {b}
(4)   else
(5)       i = 1
(6)       while W_k(t) ≠ ∅ ∩ i ≤ N_f do
(7)           h := arg max_{j∈W_k(t)} s_j
(8)           W_k(t) ← W_k(t) \ {h}
(9)           if s_j > s_k ∩ not 1_j(t) then
(10)              R_k(t) ← R_k(t) ∩ {h}
(11)          end if
(12)          i ← i + 1
(13)      end while
(14)  end if
```

ALGORITHM 1: Routing decision algorithm.

original bundle. This happens because the algorithm pushes unsociable nodes, although they meet others sporadically, to transmit to almost everyone they meet. On the contrary, if the bundle is generated by the most sociable node, there will not be any transmission until the source is itself in the range of the destination, since it is also the best possible forwarder. This seeming imbalance is explainable as follows. Because an unsociable source is likely to remain isolated for a long time, it makes sense for the network to put a greater effort to route its message along by generating replicas. In the opposite case, when a source is highly sociable, only few transmissions are required because mobility will do the rest.

In a formal tone, by using a notation similar to that of [10], let U be the set of all nodes and $N = |U|$ their number. The sociability indicator of a node $k \in U$ at time t is $s_k(t) \in [0, 1]$. We also define a Boolean indicator, $\mathbf{1}_k(t)$, which is *true* if node k already possesses the bundle, and *false* otherwise. Assume also that at time t node k has a number of active direct links to some neighbors. Let us denote as $W_k(t) \subseteq U$ the neighborhood of k. The routing decision of k consists of either keeping the bundle or selecting up to N_f next forwarders belonging to $W_k(t)$, provided they do not already possess the bundle. With respect to a destination node, b, this can be performed by using a decision algorithm to be applied to the set $W_k(t)$ and b, and yields the set, $R_k(t) \subseteq W_k(t) \subseteq U$, $|R_k(t)| \leq N_f$, of next forwarders. The pseudocode is given in Algorithm 1.

3. Evaluation of Sociability Indicators

In order to evaluate the routing strategy based on the sociability concept, we first propose a simple model where the sociability indicator of each node is computed by looking at its direct encounters, meaning that it only considers single-hop neighbors. Then, we extend the latter definition to the case where the sociability degree of one node depends not only on its direct encounters but also on the encounters of its neighbors in an iterative fashion. Finally, we introduce a set of real mobility traces that we used to test our definitions.

3.1. First Hop-Based Sociability. As a first assumption, we consider the duration of any encounter to be constant, for simplicity, and equal to 1 second. Although duration is a relevant fact in that it is related to the amount of data than can be exchanged, the aim here is just to focus on the number and frequency of encounters, whereas a more advance concept of sociability incorporating data rates is left to future studies. A definition of sociability of node k limited to its direct encounters can be given as follows. Let T be a time window of finite length and $\mathbf{1}_c(k, j, t)$ be the *meeting indicator function* defined as

$$\mathbf{1}_c(k, j, t) = \begin{cases} 1, & \text{if } k \text{ is incontact with } j \text{ at time } t, \\ 0, & \text{otherwise.} \end{cases} \quad (1)$$

Then, the sociability indicator of node k at time t is

$$s_k^{(T)}(t) = \frac{1}{N \cdot T} \sum_{j \in U} \int_{t-T}^{t} \mathbf{1}_c(k, j, \tau) d\tau. \quad (2)$$

Such a definition quantifies the social behavior of a node by counting its encounters with all the other nodes in the network over a period T. In order to assess whether this can be considered a valid estimate of the future behavior, the implications of the choice of T will be discussed.

As a first observation, T should be large enough to collect a sufficient statistic of encounters and let the indicator be significant. However, this time strongly depends on the characteristics of the network (e.g., topology, sparsity, etc.) as well as on those of mobility (e.g., velocity, correlatedness of movements, etc). For example, with reference to a vehicular network at urban scale, one user is likely to accomplish some daily tasks such as going to work in the morning, going out for lunch and go home again in the evening. In this case, a daily periodicity is clearly noticeable [18] and it is reasonable to assume that the information on social behavior of a user collected for a period $T = 1$ day is exploitable for the following day.

On the other hand, if T is so large as to allow users to change habits, the outcome parameters will no longer have a meaning. This could be the case, for instance, of a network of pedestrians carrying a mobile device in a campus [20, 21]. A student user that is observed for several semesters, is likely to modify its paths and encounters history when a new semester begins and it takes new courses.

In conclusion, T should somehow reflect the periodicity of human behavior and capture its coherence. However, since human interactions feature self-similarities at different scales [22], what periodicity scale it is more convenient to seek is a context dependent issue.

From (2), it is easy to see that $0 \leq s_k^{(T)}(t) \leq 1$. In particular, $s_k^{(T)}(t) = 1$ when the node k meets every other node at each time instant. Recent studies [23] showed that human contacts are governed by power-law behavior. In rough words, this means that a node that reaches all the others in a given period of time, will probably encounter few of them very frequently and have very rare opportunities of exchanging data with the rest. For this reason, we emphasize

the importance of evaluating not only the percentage of other nodes one gets in contact with, but also how many times. This is indeed the role of the integral in (2).

3.2. Kth Hop-Based Sociability.

As noted in Section 2.1, the sociability degree of one user should intuitively benefit from having highly sociable neighbors. With this in mind, we now aim at extending the previous definition. Let $s_k^{(n,T)}(t)$ be the sociability indicator of node k at time t, computed over a time range T, accounting for an n-hop dependence. For simplicity, we omit the dependence upon T, that is, $s_k^{(n,T)}(t) \equiv s_k^{(n)}(t)$. Then, we have as in (2)

$$s_k^{(1)}(t) = \frac{1}{N \cdot T} \sum_{j \in U} \int_{t-T}^{t} \mathbf{1}_c(k,j,\tau) d\tau = \frac{1}{N \cdot T} \sum_{j \in U} p_{k,j}^{(1)},$$

(3)

where $p_{k,j}^{(1)} = \int_{t-T}^{t} \mathbf{1}_c(k,j,\tau) \, d\tau$ and the dependence on t has been suppressed for conciseness. An immediate extension for incorporating into one node's sociability indicator the sociability of first hop neighbors, is obtained as

$$s_k^{(2)}(t) = \frac{1}{N \cdot T} \sum_{j \in U} \max\left(p_{k,j}^{(1)}, \ p_{k,j}^{(2)}\right),$$

(4)

where

$$p_{k,j}^{(2)} = \sum_{h \in U} \min\left(p_{k,h}^{(1)}, \ p_{h,j}^{(1)}\right) \cdot w_{k,h,j},$$

(5)

with $w_{k,h,j}$ being a weight parameter to be conveniently defined. Starting from the redefinition (3), $p_{k,j}^{(1)}$ represents the number of direct contacts between nodes k and j over T. In order to include indirect contacts as well, we need to define $p_{k,j}^{(2)}$, which counts the number of contacts between k and j through a third relay node. In (4) we compute the 2-hop sociability indicator by considering either direct or 2-hop connections, depending on which modality of the two gives greater contact opportunities. To explain (5), refer to the scenario of Figure 1. A node N_1 may connect to a node N_2 by exploiting a 2 hop link involving node N_3. In particular, N_1 may send its bundle to N_3 as soon as the link A is active. N_3 keeps it in a buffer and sends it to N_2 when the link B becomes active. Due to the dynamic nature of the network, the links A and B are intermittent and thus may exists or not at a given time instant depending on the mobility patterns of the nodes. Assume that, in the interval $[t-T, t]$, the two links appear 4 times each in the order shown in the bottom part of Figure 1.

Observe that at the beginning, link A appears right before link B. This makes it possible for node N_1 to send bundles to node N_2 through N_3, and should indeed be regarded as a contact opportunity. Conversely, when B appears before A (as it happens later on), no transmisson is possible from N_1 to N_2. By simple observation, it is straightforward to realize that one contact opportunity arises whenever there is an ordered sequence A, B on the timeline (for a thorough analysis of intermittent links problems in DTNs, refer to [24],
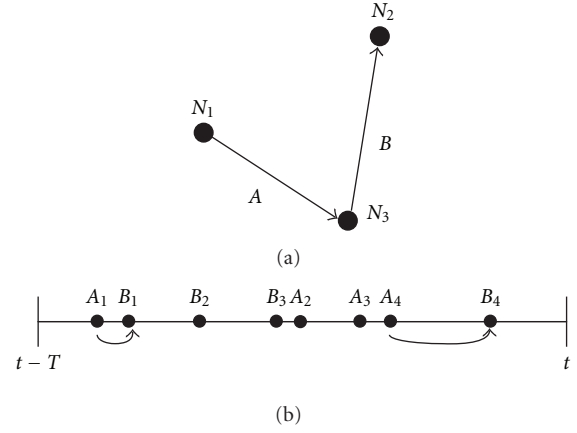

FIGURE 1: (a) simple 3 nodes network with intermittent links. (b) temporal occurrence of the links.

where the issue is addressed from the theoretical perspective of time-varying graphs.) In our example this happens twice, although links A and B appear 4 times each. Note also that link A appears 3 times before the last apparition of B. Even though N_3 can buffer all the bundles received by N_1 in the 3 transmissions, it then has only one opportunity to send them to N_2 and thus the temporal sequence of links A_2, A_3, A_4, B_4 gives rise to a single contact opportunity from N_1 to N_2. As a natural consequence, we can state that, given a sequence of apparitions of links A and B, where they appear n_A and n_B times, respectively, the number of contact opportunities from N_1 to N_2 can never exceed $\min(n_A, n_B)$. This explains the presence of the min function in (5).

Although we know that the number of contact opportunities of k with j through h is in the range $[0, \min(p_{k,h}^{(1)}, p_{h,j}^{(1)})]$, we cannot give an exact estimate of such number, because it depends on the sequence of apparition of the two links, which we do not keep track of in our model. However, it is possible to obtain an approximated average expression for it by means of simple statistical considerations.

Consider the network $k \rightarrow h \rightarrow j$. Assume the links $k \rightarrow h$ and $h \rightarrow j$, which are activated $p_{k,h}^{(1)}$ and $p_{h,j}^{(1)}$ times, respectively, appear uniformly at random on $[t-T, t]$. Define the random variable $t'_{k,h}$ as

$$t'_{k,h} := \{\text{time of 1st appearance of link } k \longrightarrow h\}$$
$$= \min\left(t_{k,h}^{(1)}, \ldots, t_{k,h}^{(a)}\right),$$

(6)

where $t_{k,h}^{(m)} \sim \mathcal{U}[t-T, t]$, for all m, and $a = p_{k,h}^{(1)}$. By noting the equivalence of the events

$$\{t'_{k,h} \leq t\} = \{t_{k,h}^{(m)} > t, \forall m\}^c,$$

(7)

with ^{c} denoting the complementary event, we have the CDF

$$F_{t'_{k,h}}(\tau) = 1 - \left[1 - F_{t_{k,h}}(\tau)\right]^a$$

$$= \begin{cases} 0, & \tau \leq t - T, \\ 1 - \left(1 - \dfrac{\tau}{T}\right)^a, & t - T \leq \tau \leq t, \\ 1, & \tau > T, \end{cases} \tag{8}$$

and the expectation

$$\mathbb{E}\left\{t'_{k,h}\right\} = \frac{T}{a+1} = \frac{T}{p_{k,h}^{(1)} + 1}. \tag{9}$$

A sufficient (but not necessary) condition for outage of the 2-hop link connecting k to j through h over a period T, is when all the instances of the $h \rightarrow j$ link appear before the expected first appearance of the $k \rightarrow h$ link. This outage probability, P_{out}, is obtained as

$$P_{\text{out}} = \left[F_{t_{h,j}}\left(\mathbb{E}\left\{t'_{k,h}\right\}\right)\right]^b = \left(\frac{\mathbb{E}\left\{t'_{k,h}\right\}}{T}\right)^b, \tag{10}$$

where $b = p_{h,j}^{(1)}$. Finally, substituting (9) into (10) yields

$$P_{\text{out}} = \left(\frac{1}{1 + p_{k,h}^{(1)}}\right)^{p_{h,j}^{(1)}}. \tag{11}$$

Hence, $1 - P_{\text{out}}$ is an upper bound to the probability that the 2-hop link connecting k to j is available at least once over the period T. This suggests that the weight $w_{k,h,j}$ in (5) should acquire the same meaning. For this reason we let $w_{k,h,j} = 1 - P_{\text{out}}$ and (5) becomes

$$p_{k,j}^{(2)} = \sum_{h \in U} \min\left(p_{k,h}^{(1)}, p_{h,j}^{(1)}\right) \cdot \left[1 - \left(\frac{1}{1 + p_{k,h}^{(1)}}\right)\right]^{p_{h,j}^{(1)}}. \tag{12}$$

It is worth noting that the expression $\min(p_{k,h}^{(1)}, p_{h,j}^{(1)}) \cdot [1 - (1/(1 + p_{k,h}^{(1)}))]^{p_{h,j}^{(1)}}$ could be interpreted as an approximation to the number of contact opportunities between k and j through h. In order to test its tightness, we simulated a three nodes network where the two links $k \rightarrow h$ and $h \rightarrow j$ appear uniformly at random on $[t - T, t]$. Results are reported in Figure 2, where the expected number of contact opportunities is plotted as a function of $p_{k,h}^{(1)}$ for different values of $p_{k,h}^{(2)}$. It can be observed that the analytical expression may be regarded as an upper bound, which is tighter for smaller values of $p_{k,h}^{(1)}$ and $p_{h,j}^{(1)}$. As a consequence, this model may be employed as long as (i) the period T is taken such that a small number of encounters between k and h, h and j, is recorded and (ii) the encounters do not deviate too much from a uniform distribution.

While the first assumption can be arbitrarily noninfluential by adjusting T, the second assumption can only
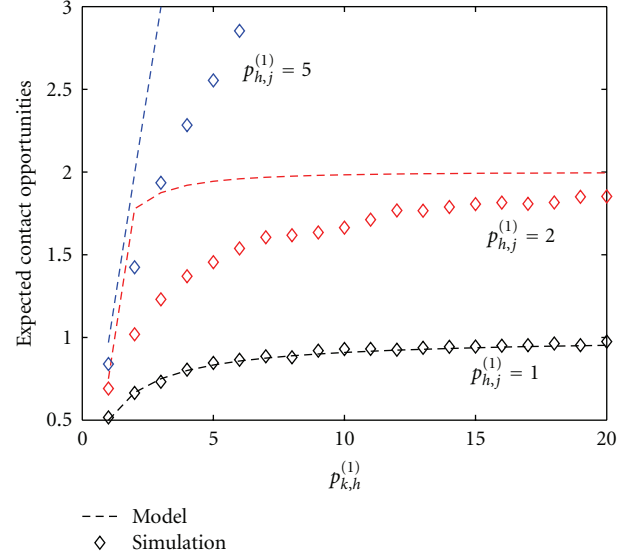

FIGURE 2: Comparison between model and simulation for the expected number of contact opportunities when $p_{k,h}^{(1)}$ and $p_{h,j}^{(1)}$ vary.

be verified by examining real traffic traces. However, more sophisticated models may be formulated when some a priori information on the traffic is available and contact statistics is inferred accordingly.

The extension of (4) to K hop is simply

$$s_k^{(K)}(t) = \frac{1}{N \cdot T} \sum_{j \in U} \max_K \left\{p_{k,j}^{(K)}\right\}, \tag{13}$$

where

$$p_{k,j}^{(K)} = \sum_{h \in U} \min\left(p_{k,h}^{(K-1)}, p_{h,j}^{(K-1)}\right) \cdot \left[1 - \left(\frac{1}{1 + p_{k,h}^{(K-1)}}\right)\right]^{p_{h,j}^{(K-1)}}. \tag{14}$$

3.3. Mobility Traces Used and Sociability Plots. Recent measurement campaigns have been conducted in the context of ambient mobile networks, with particular emphasis on vehicular networks at urban scale and pedestrian networks in a building scenario. Some of them (e.g., [25]), required the help of voluntary attendees of a conference who carried mobile devices during several days period for recording spontaneous contacts among users. At urban level, although the difficulty of finding volunteers between private users, analogous experiments could be performed on vehicles belonging to a specific entity, such us public transportation fleets. Such precious data, especially that of contacts among users, reveals very important for studying the social behavior of nodes and providing insight for potential delay-tolerant applications.

A variety of measurements have been made recently available on the Internet [3, 26] in the form of traffic traces or contact patterns. When a historical database of contacts is available, it is possible to study the social behavior. This, however, cannot be done in conjunction with mobility

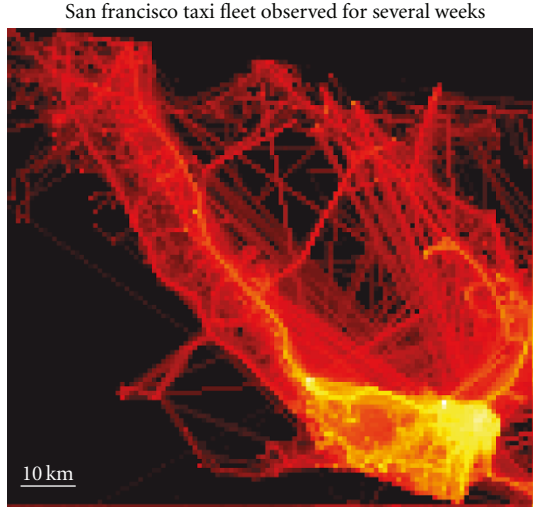San francisco taxi fleet observed for several weeks

10 km

FIGURE 3: Superposition of the mobility patterns of all San Francisco taxicabs: intensity of color is proportional to the time globally spent on each location.

consideration, since the information on mobility patterns is not directly present. In some studies (see, e.g., [10]), the log information of Wi-fi users who connect to a set of access points (APs) is examined. APs may be regarded as *locations* and consequently the mobility patterns of the users (consisting of a sequence of visits to the locations) can be indirectly inferred. By following this rationale, it is also natural to assume that two users that are connected to the same AP at a given time, they are in contact with each other.

We believe it is not possible to assess to which extent the latter assumptions hold. For this reason, we seek traces where the exact position of users is sampled, at least randomly in time. Then, from a complete mobility information, contacts history can be easily extracted.

In this paper we base our analysis on the traffic traces from the taxicabs of the city of San Francisco, CA [3], consisting of approximately 500 units. Such data report the GPS coordinates of each vehicle collected over 30 days in the San Francisco Bay Area. Each taxi is equipped with a GPS receiver and sends a location-update (timestamp, identifier, geo-coordinates) to a central server. The location-updates are quite fine-grained—the average time interval between two consecutive location updates is less than 10 seconds, allowing us to accurately interpolate node positions between location-updates. In the heatmap of Figure 3, a spatial plot is reported where the intensity of color is proportional to the time spent by the totality of taxicabs in each location.

With respect to this data, we report in Figure 4, as examples, the sociability indicators $s_k^{(1)}$, $s_k^{(2)}$, $s_k^{(3)}$, $s_k^{(4)}$ for the different vehicles (i.e., $1 \leq k \leq 500$) (Figures 4(a) and 4(c)) and the Complementary Cumulative Distribution Function (CCDF) of $s_k^{(1)}$, $s_k^{(2)}$, $s_k^{(3)}$, $s_k^{(4)}$ (Figures 4(b) and 4(d)). All the plots refer to a $T = 100$ seconds observation time. In particular, plot pairs Figures 4(a), 4(b), 4(c), and 4(d) are taken over two subsequent time windows, randomly sampled over the whole trace.

As one can see, the sociability indicators are on average smaller in plot Figure 4(c) compared to Figure 4(a): this means that in the second observation period, a smaller number of contacts has been recorded. For the 1 hop case, very few nodes have a significant indicator, while the others have almost zero indicators. This reflects in plots Figures 4(b) and 4(d), where one can observe that less than 5% of nodes have indicators greater than $3 \cdot e - 3$. When a multihop sociability definition is considered, the indicators on average increase. This means that most of the nodes having rare contacts, happen indeed to be in contact with highly sociable nodes.

However, although this effect is remarkable when moving from single hop to two hops sociability, it is not significant when the number of hops considered is greater than 3. This is coherent with the fact that multihop connections, although exponentially more numerous when K is greater, are less likely to be successful since links must appear in the correct temporal order. Finally, it bears highlighting that nodes which are completely isolated, do remain so no matter how many hops we allow. For this reason, it appears in plots Figures 4(b) and 4(d) that the probability of having a sociability indicator greater than zero, never approaches one.

4. Simulation Results

In the present section we introduce the simulator that allows us to test the forwarding scheme proposed and to compare it to other existing protocols. Then, before showing the numerical results, a brief overview of performance metrics and a short description of our benchmarking schemes, are given.

4.1. Methodology. We have designed an autonomous network simulator for testing the routing scheme. It takes as input a mobility trace like the one presented in Section 3.3 and generates mobile nodes accordingly. The time is discretized and resolution is 1 second. Each node has an infinite buffer for storing the exchanged bundles. In a realistic setup, a routing protocol should be evaluated by accounting for limited buffering capabilities. Nonetheless, although we do not address it here, we assess the validity of protocols by also counting the amount of extra bundles generated, as a rough measure of resources consumption at network level.

In addition, we make very simple assumptions at physical and MAC layers, namely, nodes are in contact when their distance is less than the transmission range, TR; channels are interference-free; and transmissions are instantaneous. Furthermore, although a node is not aware of its absolute geographical position, it has a complete knowledge of its logical connectivity, (i.e., what other nodes are within its transmission range), and it is always willing to cooperate with others.

A simulation run starts when two nodes are randomly selected as source and destination of a bundle, respectively, and terminates when the bundle is either successfully received by the recipient or discarded for exceeding a timeout threshold.
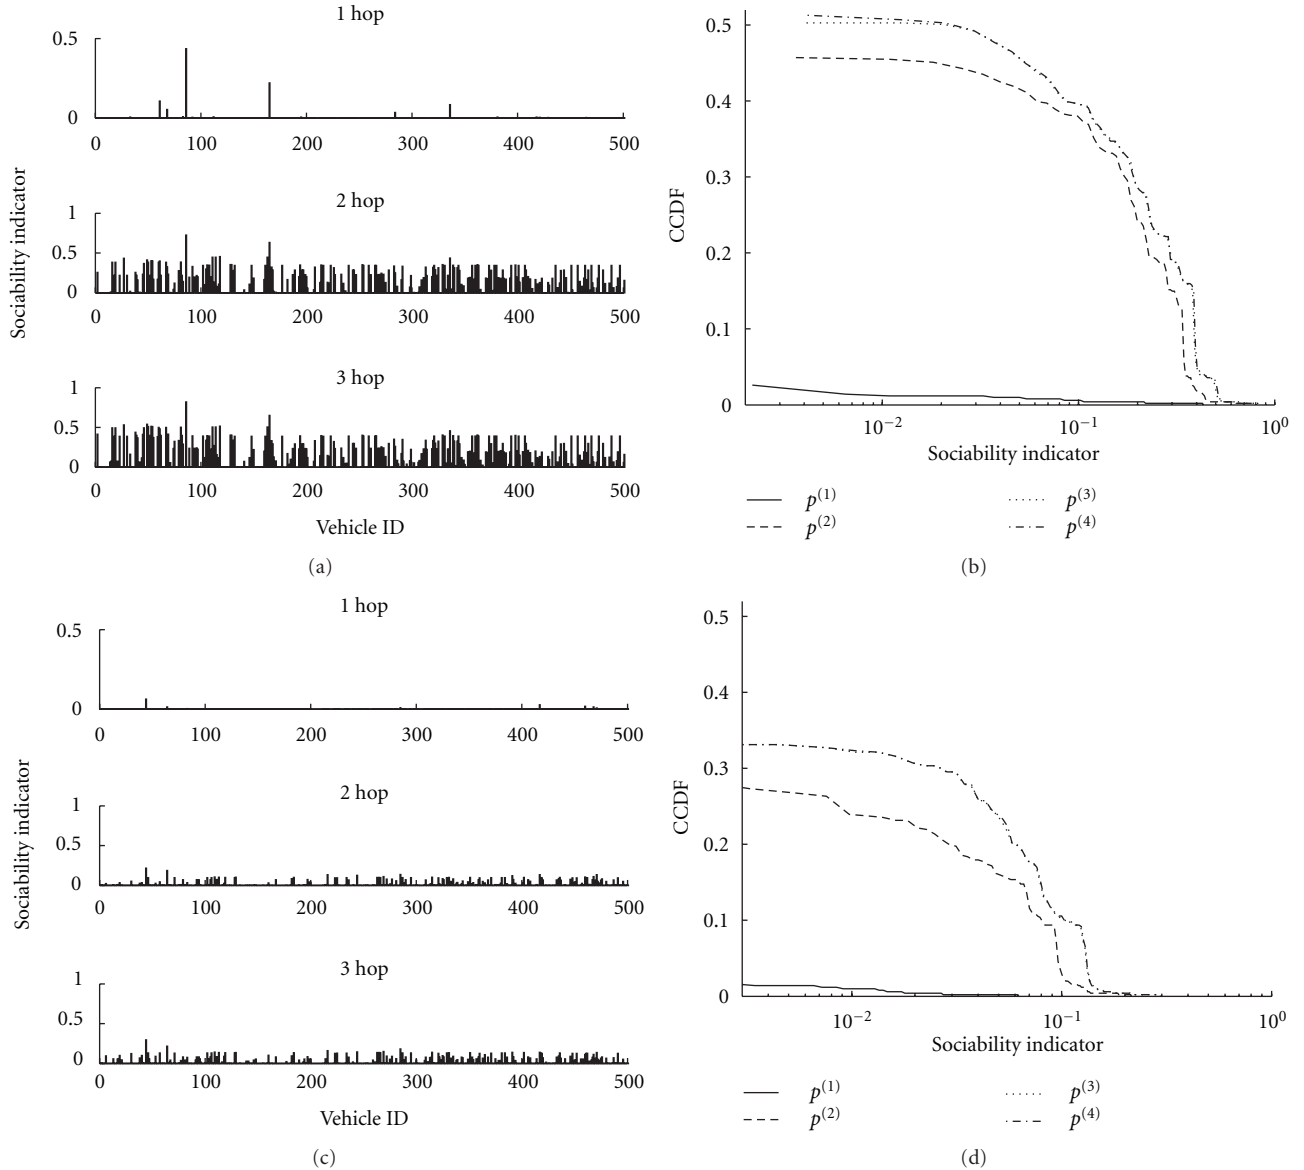
(a)

(b)

(c)

(d)

FIGURE 4: Bar plots ((a), (c)) and Complementary Cumulative Distribution Function (CCDF) plots ((b), (d)) of sociability indicators computed over two different time windows of duration $T = 100$ second each.

4.2. Input Mobility and Parameters. As input mobility, we consider the taxi cab traces introduced in Section 3.3. It must be noted that taxi cab's movements are not particularly predictable as can be those of a private citizen or even a public transportation vehicle (e.g., a bus). In fact, apart from the most frequent routes (e.g., airport to train station), each time a passenger is collected, a destination which potentially differs from the previous one has to be reached. For this reason, if we can appreciate any benefit from the sociable routing scheme in this scenario, we expect even better performance when using, for example, Seattle city bus traces [26] as input mobility.

We put two constraints in order to speed up the simulations. First, source and destination nodes are randomly picked among those that are located, at the generation

instant, in a 10×10 km square centered in downtown San Francisco. This indeed decreases the average delivering time by avoiding too far away source-destination pairs. Secondly, nodes that have not been moving for more than 1 hour cannot be source candidates. This avoids extra delays due to when a bundle is generated by a cab that is not in service, and thus has greater chances to remain isolated for long.

The number of nodes, all included, is then 535 and the traces are two weeks long. Every simulation is composed of 1000 runs (i.e., 1000 bundles are either successfully received or dropped due to excess delay) and is started at a random time on the first day of traced period. We set a timeout of 1 day and a transmission range TR = 500 meters. This value is in accordance, for example, with the standard IEEE 802.11p

[27], which is meant to be employed in vehicular networks. Finally, in case of multiple contemporaneous encounters, one node is allowed to forward the bundle to only $N_f = 1$ neighbor.

4.3. Performance Metrics and Benchmarks. For each received bundle, several measures are performed. First, the *delay*, that is, the elapsed time from generation to delivery, is recorded. Delay, which is usually imposed by an application, is a meaningful parameter for discriminating forwarding schemes. Similarly, a *cost* parameter, intended as how much of network resources a routing scheme consumes, will also be considered. In our case, we define as cost of a routing scheme the average number of network nodes that receive the bundle, apart from the intended destination. Although simplistic, this serves as an indication of how much extra traffic is generated in the network (recall that we neglect signaling traffic by assuming that nodes have perfect knowledge of the logical connectivity), how intensively the buffers are employed, and it is also related to the amount of overhead introduced at lower layers.

Generally, as it will be observed, delay is inversely proportional to cost, whereas good protocols are expected to achieve low delay at a low-cost.

We also consider the *path length*, defined as the number of hops from source to destination, as well as the *keeping time*, defined as the average time a node keeps the bundle before forwarding it to the next hop. The latter two are complementary, since the product path length × keeping time, approximately equals the delay. On equal delays, a long path length (equivalently, a short keeping time) may indicate a waste of resources and thus result in a high cost.

The performance of our routing scheme, sociable routing, is compared against that of other known protocols.

(i) *Epidemic routing.* This naive strategy [8] belongs to a category of routing protocols achieving very low delay at very high cost. It is indeed the optimum for what concern the delay performance. Practically speaking, every time a node is in contact with any other node it sends the bundle. It is easy to realize that the number of bundles present in the network grows exponentially in time. This diffusion enhances the probability that one of the bundles reaches the destination but, most of the time, its cost is unbearable for real networks. We use epidemic routing for a lower bound delay performance.

(ii) *MobySpace routing.* This scheme, introduced in [10], considers the mobility patterns of the nodes and assigns to each node a descriptor vector containing the frequencies of visits to each location. The basic idea is that nodes having similar patterns are likely to meet. Hence, a node forwards bundles to nodes whose patterns are more and more similar to that of the destination (which should be known at the source). No notion of sociability is employed but only topological considerations. MobySpace routing achieves low-cost but has a poor delay performance.

(iii) *Random routing.* This protocol is created ad hoc for comparison with sociable routing. Basically, it has the same functionalities as sociable routing (i.e., it employs Algorithm 1) but is fed with "fake" sociability indicators, meaning that they are not related to the actual social behavior of the nodes but they are just random numbers. By so doing, we expect Random routing to achieve a cost similar to that of sociable routing and a delay performance to be compared with the latter.

4.4. Results. When simulating sociable routing, the time interval between two refreshes of the sociability indicators must be set. This should be calibrated based on the nature of mobility traces. We assume no a priori information is available about the social behavior of the nodes. We then take $T = 1000$ second as initial guess. We also choose to evaluate only the first and second hop based sociability schemes, since we do not expect significant changes for a number of hops $K > 2$, as observed in Section 3.3.

In Figure 5, we report the cumulative distribution of delivered bundles over time, for the 1st and 2nd hop sociable routing, as well as for the benchmarking protocols. By observing a time window of approximately 1 day, it clearly appears how epidemic delivers a much larger amount of bundles compared to other solutions. However, as previously noted, this scheme is practically unfeasible.

Conversely, MobySpace is the one delivering the smallest amount of bundles. The reason seems to be the presence of large deviations from the mean delay, occurring when a node does not find a suitable relay and keeps the bundle for long. A deeper consideration is that the basic assumption of the protocol, according to which two nodes having similar patterns are likely to meet, is not easily applicable to the case of taxi, where all nodes tend to visit a small set of locations (e.g., airport, main square, etc.) with approximately the same frequencies. 1-hop sociable routing seems to be delivering the largest amount of bundles at a fairly constant rate. Random Routing, instead, which employs the same scheme as 1-hop Sociable but with "fake" sociability indicators, shows a more irregular trend. The reason is that when bundles are sent to not very sociable nodes, they are likely to be stuck, since they do not meet other nodes, and consequently cause extra delays. 2-hop Sociable has a slightly poorer performance than 1-hop Sociable, at least in terms of number of deliveries. Finally, all the protocols could deliver 100% of packets before timeout except MobySpace, which dropped 1.8% of bundles.

In Table 1 we introduce, besides the cost, other metrics among those discussed in Section 4.3. Average values, taken over 1000 simulation runs, together with the 95% confidence interval, obtained through the Student's t distribution, are reported. This table reveals the opposite trend of cost with respect to delay. In fact, the delay performance of epidemic, for example, is payed off by a large waste of resources (2.5 times more than 1-hop Sociable). By looking at path lengths, it can be seen how low-cost strategies lead to short paths from source to destination. As an extreme case, simulation of MobySpace reveals that most of successful deliveries
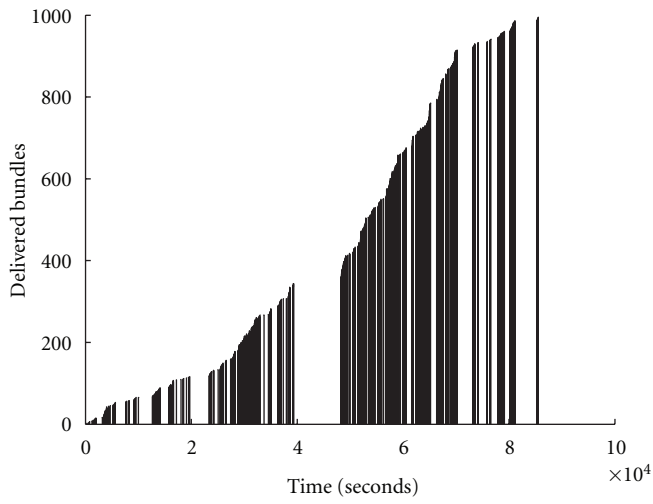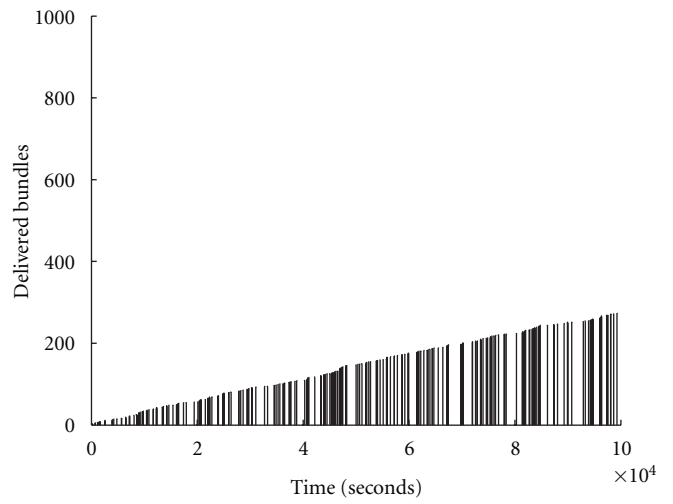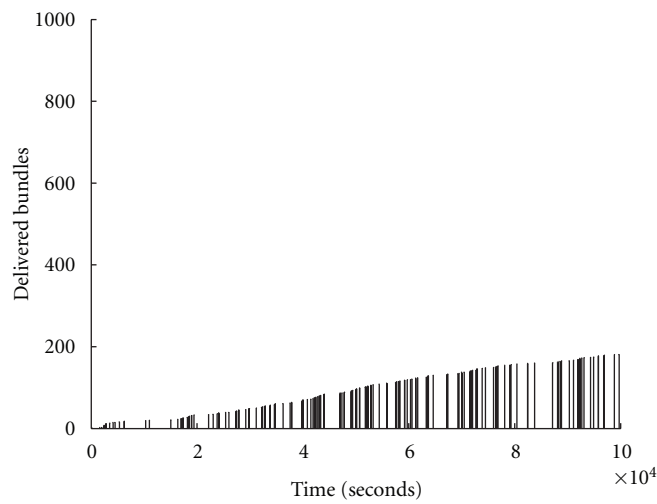
(a) Epidemic
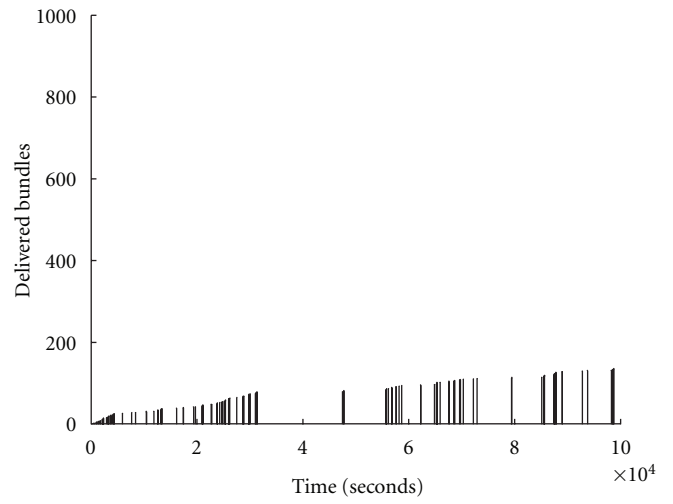
(b) MobySpace
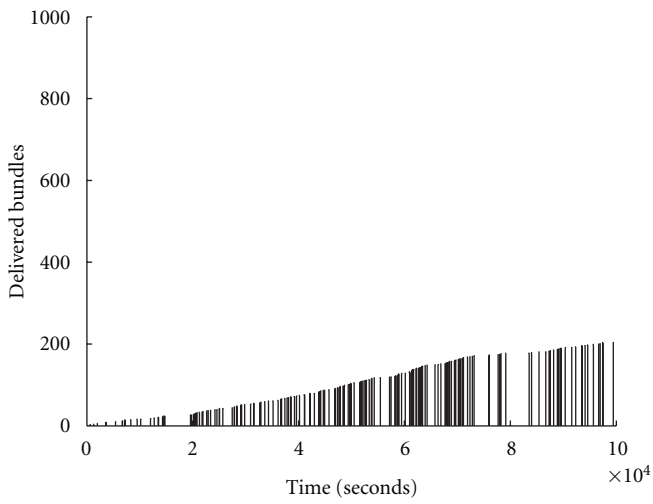
(c) Random

(d) 1hop Sociable

(e) 2hop Sociable

FIGURE 5: Cumulative bundles delivery over time for the routing scheme considered.

TABLE 1: Numerical results in terms of average delay, cost, and path length for the protocols considered.

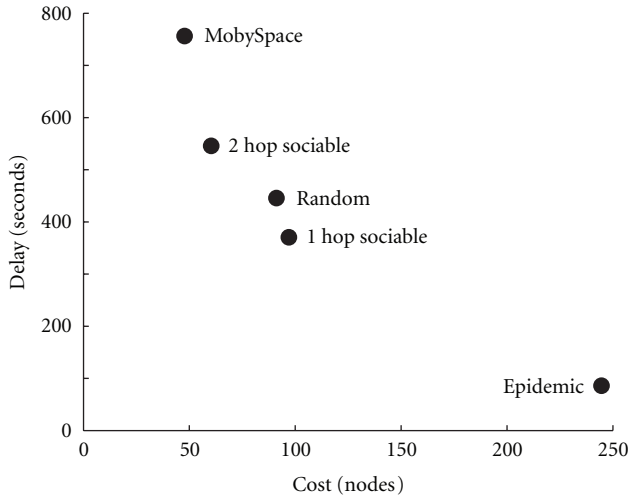|	Delay (seconds)	Cost (nodes)	Path length (hops)
Epidemic	85.98 ± 20.61	244.63 ± 6.17	10.65 ± 0.33
MobySpace	756.48 ± 240.78	47.73 ± 5.67	3.40 ± 0.21
Random	445.62 ± 29.13	91.14 ± 4.90	4.77 ± 0.13
1hop Sociable	370.23 ± 26.99	97.06 ± 5.25	4.82 ± 0.14
2hop Sociable	545.90 ± 50.44	60.28 ± 3.29	4.25 ± 0.12

FIGURE 6: Position of the examined routing schemes on the delay versus cost plane.

employed no more than $3 \div 4$ hops. Finally, by noting that the order of magnitude of delays is $100 \div 1000$ second, we can now state that the choice $T = 1000$ second for the refreshing time was a reasonable one. In fact, the situations where many refreshes occur while the same bundle is in the network are intuitively non-optimal.

As one can see, the performance of 2-hop Sociable is poorer than Random in terms of delay and better than 1-hop Sociable in terms of cost. This deserves further considerations. A greater insight is gained by looking at Figure 4 once again. Consider, for example, Figure 4(a): in the scenario we have selected, by only accounting for 1-hop sociability on a $T = 100$ seconds period, we observe a situation where very few nodes have *sociability indicators* remarkably greater than zero, while the rest is very close to zero or zero (Figure 4(b) suggests in fact a power law distribution). In rough words, there are few suitable forwarders and many unsociable nodes. This imply the following. A node (statistically a poorly sociable one) generates a bundle. Then, it transmits it to its slightly more sociable neighbors. This part of the process is slow, since encounters are sporadic for the origin node but, at the same time, transmissions are likely to occur upon every encounter. At some point the most sociable forwarder is reached. Now the process rapidly converges to delivery, thanks to the identification of the most suitable forwarder. We may equivalently say that there is a little bit of flooding

at the very beginning of the process until the key forwarder is met. Then what it takes is just waiting for it to meet the destination node. Note that the latter resembles the *Spray and Wait* approach proposed in [28], with the difference that the spray phase intelligently stops when the right forwarder is met.

Consider now the 2nd plot of Figure 4(a). 2-hop based *sociability indicators* still show few best forwarders. However many other nodes have non-negligible indicators. In particular, a large subset has sociability around 0.5. Now imagine a bundle is originated by one such node. According to the algorithm, the latter will automatically exclude from the list of potential next-hop relays all its peers (i.e., those having equal or smaller indicator—a great percentage). The key implication is that there is no initial flooding. Instead, the origin node tends to wait to meet one of the 2 or 3 strong forwarders before initiating a transmission. Thus, while in the 1-hop case all nodes are almost equally likely to potentially receive the bundle, this different definition of *sociability indicators* automatically excludes a subset of them. This does not necessarily result in worse performance. In fact, although we observe a greater delivery delay with respect to 1-hop sociability (see Table 1), we considerably reduce the cost, which makes it better scalable.

In conclusion, 1-hop and 2-hop sociability are two very different routing strategies with pros and cons, among which the network designer should choose the preferred one based on the requirements (i.e., whether to minimize cost or delay).

For a visually more effective positioning of the examined protocols on a delay versus cost plane, we refer to Figure 6. By focusing on the left part of the plot, where the "affordable" schemes are positioned, it appears that 1-hop Sociable represents a good compromise.

5. Conclusions and Future Work

The contribution of this paper has been the proposal of a novel routing scheme for DTNs. sociable routing chooses the set of best forwarders among those having high sociability indicators, the latter being time-varying scalar parameters. Sociability indicators relate to the social characteristics of network nodes, by capturing the frequency and type of their encounters. The routing strategy has been widely discussed and evaluated by simulation on a DTN of vehicles in urban environment. Results and comparison with other existing protocols showed that sociability based routing can achieve a good compromise in terms of delay performance and cost.

Several open issues throughout the paper have been mentioned and deserve attention in future works. First, the study of distributed algorithms for computing sociability metrics is currently being addressed. Finally, an additional simulation campaign using other mobility traces (possibly with different social characteristics) could reveal the robustness of Social Routing against diverse scenarios and prove the generality of the key ideas.

Acknowledgment

References

[1] K. Fall, "A delay-tolerant network architecture for challenged internets," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '03)*, pp. 27–34, ACM, Karlsruhe, Germany, August 2003.

[2] S. Jain, K. Fall, and R. Patra, "Routing in a delay tolerant network," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '04)*, pp. 145–157, ACM, August-September 2004.

[3] M. Piorkowski, N. Sarafijanovic-Djukic, and M. Grossglauser, "CRAWDAD data set epfl/mobility (v. 2009-02-24)," February 2009.

[4] J. Burgess, B. Gallagher, D. Jensen, and B. N. Levine, "Max-Prop: routing for vehicle-based disruption-tolerant networks," in *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM '06)*, pp. 1–11, April 2006.

[5] E. M. Daly and M. Haahr, "Social network analysis for routing in disconnected delay-tolerant MANETs," in *Proceedings of the 8th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '07)*, pp. 32–40, ACM, September 2007.

[6] A. Balasubramanian, B. Levine, and A. Venkataramani, "DTN routing as a resource allocation problem," in *Proceedings of the Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications (SIGCOMM '07)*, pp. 373–384, ACM, Kyoto, Japan, August 2007.

[7] E. P. C. Jones, L. Li, and P. A. S. Ward, "Practical routing in delay-tolerant networks," in *Proceedings of the ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN '05)*, pp. 237–243, ACM, August 2005.

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

[9] S. Burleigh, A. Hooke, L. Torgerson, K. Fall, V. Cerf, B. Durst, K. Scott, and H. Weiss, "Delay-tolerant networking: an approach to interplanetary internet," *IEEE Communications Magazine*, vol. 41, no. 6, pp. 128–136, 2003.

[10] J. Leguay, T. Friedman, and V. Conan, "Evaluating mobility pattern space routing for DTNs," in *Proceedings of the 25th IEEE International Conference on Computer Communications (INFOCOM '06)*, pp. 1–10, Barcelona, Spain, April 2006.

[11] M. Musolesi and C. Mascolo, "Car: context-aware adaptive routing for delay-tolerant mobile networks," *IEEE Transactions on Mobile Computing*, vol. 8, no. 2, pp. 246–260, 2009.

[12] P. Hui, J. Crowcroft, and E. Yoneki, "BUBBLE rap: social-based forwarding in delay tolerant networks," in *Proceedings of the 9th ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc '08)*, pp. 241–250, ACM, May 2008.

[13] W. Zhao, M. Ammar, and E. Zegura, "Controlling the mobility of multiple data transport ferries in a delay-tolerant network," in *Proceedings of the 24th Annual Joint Conference of the IEEE Computer and Communications Societies (INFOCOM '05)*, vol. 2, pp. 1407–1418, March 2005.

[14] B. Hull, V. Bychkovsky, Y. Zhang, K. Chen, M. Goraczko, A. Miu, E. Shih, H. Balakrishnan, and S. Madden, "CarTel: a distributed mobile sensor computing system," in *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems (SenSys '06)*, pp. 125–138, ACM, November 2006.

[15] M. Abuelela, S. Olariu, and I. Stojrnenović, "OPERA: opportunistic packet relaying in disconnected vehicular ad hoc networks," in *Proceedings of the 5th IEEE International Conference on Mobile Ad-Hoc and Sensor Systems (MASS '08)*, pp. 285–294, October 2008.

[16] F. Fabbri and R. Verdone, "Sociability based routing for environmental opportunistic networks," *Radio Communication Series: Recent Advances and Future Trends in Wireless Communication*, vol. 1, no. 1, pp. 98–103, 2010.

[17] R. Kumar, D. Liben-Nowell, and A. Tomkins, "Navigating low-dimensional and hierarchical population networks," in *Proceedings of the 14th Annual European Symposium on Algorithms (ESA '06)*, vol. 4168 of *Lecture Notes in Computer Science*, pp. 480–491, Springer, 2006.

[18] T. Karagiannis, J.-Y. Le Boudec, and M. Vojnovic, "Power law and exponential decay of inter contact times between mobile devices," in *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking (MobiCom '07)*, pp. 183–194, ACM, September 2007.

[19] V. Conan, J. Leguay, and T. Friedman, "Characterizing pairwise inter-contact patterns in delay tolerant networks," in *Proceedings of the 1st international conference on Autonomic Computing and Communication Systems (Autonomics '07)*, pp. 1–9, Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, 2007.

[20] T. Henderson, D. Kotz, and L. Abyzov, "The changing usage of a mature campus-wide wireless network," in *Proceedings of the 10th Annual International Conference on Mobile Computing and Networking (MobiCom '04)*, pp. 187–201, ACM, October 2004.

[21] M. McNett and G. M. Voelker, "Access and mobility of wireless PDA users," *ACM SIGMOBILE Mobile Computing and Communications Review*, vol. 9, no. 2, pp. 40–55, 2005.

[22] R. Guimerà, L. Danon, A. Díaz-Guilera, F. Giralt, and A. Arenas, "Self-similar community structure in a network of human interactions," *Physical Review E*, vol. 68, no. 6, Article ID 065103, 4 pages, 2003.

[23] A. Chaintreau, P. Hui, J. Crowcroft, C. Diot, R. Gass, and J. Scott, "Impact of human mobility on opportunistic forwarding algorithms," *IEEE Transactions on Mobile Computing*, vol. 6, no. 6, pp. 606–620, 2007.

[24] J. Tang, M. Musolesi, C. Mascolo, and V. Latora, "Temporal distance metrics for social network analysis," in *Proceedings of the 2nd ACM Workshop on Online Social Networks (WOSN '09)*, pp. 31–36, ACM, 2009.

[25] P. Hui, A. Chaintreau, J. Scott, R. Gass, J. Crowcroft, and C. Diot, "Pocket switched networks and human mobility in conference environments," in *Proceedings of the ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN '05)*, pp. 244–251, ACM, August 2005.

[26] J. G. Jetcheva, Y.-C. Hu, S. PalChaudhuri, A. K. Saha, and D. B. Johnson, "CRAWDAD data set rice/adhoccity (v. 2003-09-11)," September 2003.

[27] L. Stibor, Y. Zang, and H.-J. Reumerman, "Neighborhood evaluation of vehicular ad-hoc network using IEEE 802.11p," in *Proceedings of European Wireless Conference*, 2007.

[28] T. Spyropoulos, K. Psounis, and C. S. Raghavendra, "Spray and wait: an efficient routing scheme for intermittently connected mobile networks," in *Proceedings of the ACM SIGCOMM Workshop on Delay-Tolerant Networking (WDTN '05)*, pp. 252–259, ACM, August 2005.