Minimum Expected *-cast Time in DTNs

Andreea Picu and Thrasyvoulos Spyropoulos

ETH Zürich, Switzerland
lastname@tik.ee.ethz.ch

Abstract. Delay Tolerant Networks (DTNs) are wireless networks in which end-to-end connectivity is sporadic. Routing in DTNs uses past connectivity information to predict future node meeting opportunities. Recent research efforts consider the use of social network analysis (i.e., node communities, centralities etc.) for this forecast. However, most of these works focus on unicast. We believe that group communication is the natural basis of most applications envisioned for DTNs. To this end, we study constrained *-cast (broad-, multi- and anycast) in DTNs. The constraint is on the number of copies of a message and the goal is to find the best relay nodes for those copies, that will provide a small delivery delay and a good coverage. After defining a solid probabilistic model for DTNs collecting social information, we prove a near-optimal policy for our constrained *-cast problems that minimizes the expected delivery delay. We verify it through simulation on both real and synthetic mobility traces.

Keywords: DTN, social network, broadcast, multicast, anycast.

1 Introduction

Delay tolerant networks (DTNs) are sporadically-connected networks, experiencing frequent network partitioning. This is usually the result of one or more factors: high mobility, low node density, stringent power management, attacks etc. The DTN concept was first considered for challenged or exotic network conditions such as satellite networks with periodic connectivity, underwater acoustic networks with moderate delays and frequent disruptions, sensor networks for wildlife tracking, and Internet provision to developing regions. Nonetheless, the constant growth in number of mobile, networking-enabled devices has created the possibility for anyone to set up a network anywhere, anytime and with anybody, even in urban environments, something that has not remained unnoticed by researchers and industry. Numerous applications, e.g., peer-to-peer content exchange, localized content and service discovery, social networking, that are asynchronous and thus tolerant to delays, can be supported using a DTN architecture.

Nevertheless, communication in such networks is quite challenging due to the lack (or rapid change) of end-to-end paths. To this end, numerous novel DTN routing protocols for unicast traffic [23, 9, 19, 17] have been proposed, that attempt to infer the contact patterns between nodes and predict paths with high probability of delivery. Among these, a recent research thread has proposed to take advantage of *social* properties of such networks [8, 16]. Mobile devices are carried by humans whose mobility is governed by explicit (e.g., friends, colleagues) or implicit social factors (e.g., commuters on the same bus, shared

E. Altman et al. (Eds.): Bionetics 2009, LNICST 39, pp. 103–116, 2010.

cafeteria etc.). As a result, these protocols try to expressly capture social links between nodes (e.g., communities) and use complex network analysis tools [3, 22] to identify nodes that can be used to improve the routing efficiency.

Although unicast routing in DTNs has received a considerable amount of attention (see survey in [25]), other communication patterns, such as multicast or anycast, as well as appropriate metrics and algorithms to optimize these, have been largely neglected [13, 26]. However, as has been noted in [5], most applications envisioned for such "pocket switched networks" [15] are expected to not be point-to-point. For example, pushing data towards a subset of interested subscribers to a channel or service can be described better as multicast or broadcast, while service or data discovery that does not know (or care about) the identity of the node that provides the service in advance can be better described as anycast.

In this paper, we take a first step towards optimizing **-cast* (i.e., broadcast, multicast, and anycast) in DTN environments of human carried wireless devices. The approach we take will also be based on social network analysis, as we believe that the preliminary advantages that have been demonstrated for unicast routing [8, 16] are well-grounded, independently of the communication mode studied. Our work has two main contributions. First, we develop a solid probabilistic model of predictive value for any DTN gathering social information, e.g., node degrees. Second, we prove a near-optimal policy for a class of *-cast problems in a resource constrained setting that minimizes the expected delivery delay.

The rest of the paper is structured as follows. Section 2 presents the network model used throughout the study and we formally define the problems we analyze. Section 3 comprises the probabilistic model and the proof of delay optimization. A preliminary evaluation is shown in Section 4. We conclude in Section 5.

2 DTN Network Model

In this section, we present the network model used in our study and we formally define the problems we subsequently address.

2.1 Network Model

Let $\mathcal{N}$ be the set of all nodes in the network, $|\mathcal{N}| = N$. Each of the N nodes is identified by a unique ID and its mobility is assumed to be governed by (implicit or explicit) social relations. Specifically, (i) we can identify *node communities*, i.e., sets of nodes that tend to meet each other preferentially, and ii) nodes have different *sociability* or number of nodes they meet in a given time interval, ranging from *solitary* to *gregarious*. This seems a reasonable assumption, since mobile devices (network nodes) are carried by humans, who engage in socially meaningful relationships. Several previous DTN experiments, [10, 18], have confirmed this type of interaction patterns and information flow. In this paper, we will focus on the latter characteristic only, and defer studying community structure for future work. Specifically, we assume a *social graph* is created using past *contacts* between nodes and we will optimize around the degree distribution of the graph.

Contacts. A *contact* between two nodes happens when those nodes have setup a bi-directional wireless link between them. We assume that:

i) contacts last for a negligible time compared to that between two successive contacts, but long enough to allow all the required data exchanges to happen,
ii) contacts occur in sequence, i.e., there are no simultaneous contacts[1].

Each contact has a defining feature: its *type*. The type of a contact is uniquely defined by the IDs of the two nodes taking part in it. Hence, there are $\binom{N}{2}$ possible types of contact. We will assume that:

iii) the types of successive contacts are mutually independent random variables.

The distribution of each variable is fully defined by a $\binom{N}{2}$-sized vector of probabilities summing to 1. The probabilities depend on the mobility model and can be estimated in function of the information assumed available. In our case, mobility behavior is captured in a social graph, described next, and contact probabilities depend on the node degrees in this graph (as shown in Section 3.1).

Social graph. A *social graph* represents our network: nodes (mobile devices) are vertices and contacts are edges. The graph seeks to capture the aforesaid social features of this network. Ways of creating the social graph of a DTN are implicitly used in various previous works [8, 16]. More recently, [14] explicitly addresses this as a standalone issue. Here, we build the social graph as follows.

Time is divided in W-sized windows. When a contact occurs, the respective edge is added to the graph. The node degree shows how many ***different*** devices that node contacted within one time window. Hence, the graph is simple: multiple contacts between the same two nodes result in only one edge. At the end of each window, nodes update their running average degrees over all time windows. This graph allows us to identify important actors in the underlying social network, by using complex network analysis metrics such as centrality. Following the practice of [20], we make the further assumption that the maximum degree in the graph is $o(\sqrt{N})$, reflecting the fact that in a large enough social network, a single person, even a very social one, cannot know a constant fraction of all users.

2.2 Constrained *-Cast Problem

We now describe the scenarios we will address in our probabilistic analysis. All scenarios use the above network model.

We focus on a subclass of constrained group communication problems in the DTN environment: broadcast, multicast and anycast. The constraint under consideration is on the number of copies of a message. The problem is formally defined below for the general case. Specific definitions for broadcast, multicast and respectively, anycast follow.

Definition 1 (General case). *At time* 0*, a source node* $s \in \mathcal{N}$*, creates a message* m *to be delivered to a fixed set* $\mathcal{D}$ *(*$|\mathcal{D}| = D$*) of distinct destination nodes in the network, with* $\mathcal{D} \in \wp(\mathcal{N})$ *and* $1 \leqslant D \leqslant N$*.* s *must then find a set of nodes* $\mathcal{L}$ *of size* L *which will each store a copy of* m*. The general goal is that every node in* $\mathcal{D}$ *receive* m*. Maximizing the number of nodes that receive it before a deadline (time to live (TTL)) or minimizing the delay until all nodes receive the message are different flavors of the optimization problem in hand*[2]*.*

---

[1] This is the case, for example, if we assume that the arrival process of contact events is Poisson. In general, this assumption just implies a relatively sparse network.

[2] $\wp(\mathcal{N})$ is the power set of $\mathcal{N}$, i.e., the set of all subsets of $\mathcal{N}$.

Definition 2 (Broadcast). *In broadcast,* $\mathcal{D} = \mathcal{N}$.

Definition 3 (Multicast). *In multicast,* $\mathcal{D}$ *is such that* $2 \leqslant D \leqslant N - 1$.

Definition 4 (Anycast). *In anycast, still* $\mathcal{D} \in \wp(\mathcal{N})$ *and* $1 \leqslant D \leqslant N$. *However, the goal of the routing changes. Instead of showing the message to the entire destination set,* m *need now only be seen by at least one node belonging to that set, whichever node that is.*

Therefore, with our routing requirements, the creator of a message must find a set $\mathcal{L}$ of *permanent message carriers* and give them each a copy of the message. By permanent message carriers we mean that the carriers cannot forward the message further, except to nodes in the destination set $\mathcal{D}$ who consume it (similar to the two-hop [4] and spray and wait schemes [23]). The size L of the set of permanent message carriers should be at most equal to the permitted number of copies of a message. Since the copies cannot be forwarded, the choice of permanent message carriers must be optimum to ensure the timely delivery of the messages.

Our analysis focuses on the delivery delay measured in number of contacts. Let T_d be the random variable counting the delivery delay of a message under the above routing constraints. Then, T_d can be expressed as a sum of two other random variables: i) *the finding time*, the time for choosing the carriers and handing them the copies, T_f and ii) *the showing time*, the time until all destination nodes have seen the message, T_s. Assume[3] $T_d = T_f + T_s$. Based on this decomposition of the delivery delay, we can identify three levels of the problem:

1. **Offline optimization with global knowledge**. Here, we assume the availability of a *graph oracle* that provides information about the degree of all nodes in the social graph. Our goal is to select the L permanent message carriers, according to their (globally known) degrees, that will maximize the delivery ratio and minimize the delivery delay. Furthermore, we assume that after the set $\mathcal{L}$ is chosen, copies of the message are transmitted instantaneously to the nodes in this set. Hence, we ignore T_f and focus on T_s, the showing time.
2. **Online optimization with global knowledge**. We still have the same oracle. Using the scheme devised above, we identified the L permanent message carriers and our new goal is to find these carriers as soon as possible and give them the message. This part will assess only the T_f element.
3. **Online optimization**. We no longer have the oracle. Using insight from the two schemes above, we want to optimize the delivery delay T_d, in its entirety.

In this paper, we address the first level, i.e., we prove an optimum for $\mathbb{E}[T_s]$ and sharp concentration of T_s around its mean. The analysis of the next two levels is part of our future work.

3 Probabilistic Analysis

In this section we provide a theoretical analysis of the evolution of T_s, the showing time component of the delivery delay, in function of the choice of permanent message carriers, $\mathcal{L}$. To do so, let us first provide some basic tools.

[3] In reality, T_f and T_s will overlap. We discuss this later on.

3.1 Probability Measure

Consider the *contact*, as defined in Section 2.1 and let us place ourselves at time t. Define the random variable C_t as the type of the first contact occurring after time t in the network. This type is unique, since there are no simultaneous contacts and we will denote it (i, j), where i and j are the two nodes coming in contact.

As stated before, estimates of the probability distribution of C_t can be obtained, using the information available. In our case, that information is twofold. First, we know that the structure of our network is strongly influenced by social relations and second, we dispose of a degree distribution oracle. Several studies, [22, 3], found that most social networks have skewed degree distributions and nodes are linked proportionally to their degrees. For us, this suggests that the probability of C_t being (i, j) should be proportional to d_i and d_j, the degrees of i and j, respectively. We will now provide a formal argument that it is indeed true.

Let us now define our probability space $(\Omega, \mathcal{F}, \mathbb{P})$:

Ω is the sample space, i.e., the set of all outcomes of our elementary experiment, the next contact. Then, Ω consists of all possible pairs of nodes and $|\Omega| = \binom{N}{2}$.

$\mathcal{F}$ is a σ-algebra of events, i.e. $\mathcal{F} \subseteq \wp(\Omega)$ and hence, $|\mathcal{F}| \leqslant 2^{|\Omega|} = 2^{\binom{N}{2}}$.

$\mathbb{P}$ is a measure on $(\Omega, \mathcal{F})$ called the probability measure. $\mathbb{P}(\Omega) = 1$.

In Section 2.1, we showed how to calculate node degrees based on a time window mechanism and we made the assumption that this is the only information available to nodes. We now place ourselves on the timeline, at the end of a window W. The current degree distribution reflects the contacts in all past windows and, if the network does not change erratically, we assume that it also provides a good estimate for the contacts in window $W + 1$. We will use this to calculate the probability distribution of C_t.

A degree distribution matches multiple graph instances. The probability of an edge existing in the network graph can be calculated as the number of graph instances containing that edge over the total number of graph instances matching our distribution. However, this is a difficult combinatorics problem. An easier way to calculate this probability is to use a graph construction model that generates fixed degree sequence graphs: the configuration model [2]. For every time window of size W, the configuration model guarantees the given degree distribution. In this model, each node is assigned a number of *stubs* or half edges, equal to its degree. The stubs are then paired with each other. The probability of an edge between nodes i and j is calculated as follows. For nodes i and j, and for $1 \leqslant s \leqslant d_i$ and $1 \leqslant t \leqslant d_j$, we define $I_{st,ij}$ to be the indicator of the event: "stub s is paired to the stub t", where the stubs are numbered arbitrarily. If $I_{st,ij} = 1$ for some st, then there is an edge between vertices i and j. It follows that the probability of an edge linking i and j is

$$p_{ij} = \sum_{\substack{1 \leqslant s \leqslant d_i \\ 1 \leqslant t \leqslant d_j}} Pr[I_{st,ij} = 1] = d_i d_j Pr[I_{11,ij} = 1], \tag{1}$$

since the probability of producing an edge between i and j by pairing the stubs s and t does not depend on s and t. Now, $Pr[I_{11,ij} = 1]$ is the probability that stubs 1 of i and 1 of j are paired to each other, which is equal to $\left(\sum_{i=1}^{N} d_i - 1\right)^{-1}$. Therefore, the probability of an edge between i and j is

$$p_{ij} = \frac{d_i d_j}{\sum_{1 \leqslant i \leqslant N} d_i - 1} \tag{2}$$

and the probability of no edge between i and j is $1 - p_{ij}$.

Hence, the probability for an edge linking i and j to appear in the graph of the next time window is p_{ij} in equation 2. However, this is **not** the probability of the next contact being between i and j. Based on p_{ij}, and assuming uniform sampling of the edges generated by the configuration model (i.e., each node pair has the same contact frequency), this probability is shown in equation 3.

$$\mathbb{P}[(i,j)] = \frac{d_i d_j}{\sum_{i=1}^{N} d_i - 1} \cdot \frac{1}{\frac{1}{2}\sum_{i=1}^{N} d_i} = \frac{d_i d_j}{\frac{1}{2}\left(\sum_{i=1}^{N} d_i\right)^2 - \frac{1}{2}\sum_{i=1}^{N} d_i} = \frac{d_i d_j}{\sum_{1 \leqslant i < j \leqslant N} d_i d_j + \sum_{i=1}^{N} \binom{d_i}{2}} \tag{3}$$

The term $\sum_{i=1}^{N} \binom{d_i}{2}$ in the denominator of equation 3 reflects the probability of self-loops ("choose 2 stubs out of d_i" are the self-loops of node i) in the graph generated by the configuration model, which translates into contacts where only one node is involved. To avoid this, we consider the *erased configuration model* [24]. Starting from the multigraph obtained through the configuration model, we merge all multiple edges into a single edge and erase all self-loops. It was shown in [24] that, provided that the maximum degree of the graph is $o(\sqrt{N})$, the *configuration model* and the *erased configuration model* are asymptotically equivalent, in probability. Since we justifiably made this assumption in Section 2.1, we can safely approximate the probability of the next contact being between i and j as

$$\mathbb{P}[(i,j)] = \frac{d_i d_j}{\sum_{1 \leqslant x < y \leqslant N} d_x d_y} \tag{4}$$

This concludes our estimation of the probability distribution of C_t, as we can now calculate the probability of any type of contact using the degree oracle.

3.2 Group Communication as a CCP

To analyse the problem in Definition 1 and its particular cases, we will use a notorious probabilistic method: the Coupon Collector Problem (CCP). This method was formally introduced in [11] and its setting is the following.

Lemma 1 (Coupon Collector). *Consider an unlimited supply of coupons of n distinct types. At each trial, we collect a coupon uniformly at random and independently of previous trials. Then, it takes on average nH_n trials, to collect at least one of each of the n coupon types (H_n is the Harmonic Number). In addition, the number of trials until the full collection is obtained is sharply concentrated around its expectation.*

In our problem, each contact is a trial, $\mathcal{L}$ - the set of carriers, is the collector and the nodes in $\mathcal{D}$ are the desired coupons. A node $d \in \mathcal{D}$ is collected when the first contact between d and any of the nodes in $\mathcal{L}$ has occurred. Unlike the coupon collector, each node has a different, known probability of being collected in our case. In [12], Flajolet et al. give formulas for the expected number of trials until both partial collections and a full collection, in the heterogeneous case.

Another difference between our setting and the CCP is the *null coupon.* The CCP has n types of coupons and $\sum_{k=1}^{n} p_n = 1$. We have D distinct nodes to be collected, plus a null event representing contacts in which none or both of the two nodes involved, carries the message m. In other words, there exist $n = D+1$ types of coupons (D destination nodes and no message exchange) and successful message dissemination means collecting all but the null coupon, that is D coupons. Let p_0 be the probability of the null event: $p_0 + \sum_{k=1}^{D} p_k = 1$.

Lemma 2. *For each of the definitions in Section 2.2, the probabilities p_i, $0 \leqslant i \leqslant D$ are fully defined using the probability distribution from Section 3.1:*

$$p_k = \frac{d_{n_k} \sum_{j=1}^{L} d_{c_j}}{\sum_{1\leqslant x<y\leqslant N} d_x d_y} \quad \textit{and} \quad p_0 = \frac{\sum_{1\leqslant i<j\leqslant L} d_{c_i} d_{c_j} + \sum_{1\leqslant i<j\leqslant N-L} d_{n_i} d_{n_j} + \sum_{\substack{1\leqslant i\leqslant L \\ 1\leqslant j\leqslant N-L-D}} d_{c_i} d_{n_{D+j}}}{\sum_{1\leqslant x<y\leqslant N} d_x d_y},$$

where c_1 to c_L are the L collector nodes, n_1 to n_D are the D[4] coupon nodes and n_{D+1} to n_{N-L} are the rest of the nodes.

Proof. The probability for the collector $\mathcal{L}$ to collect a fixed coupon type n_k, $1 \leqslant k \leqslant D$ is the probability that n_k meets any of the collector nodes, i.e.

$$p_k = \mathbb{P}\left[\bigcup_{1\leqslant j\leqslant L} (n_k, c_j)\right] \overset{\text{mut. ex.}}{=} \sum_{1\leqslant j\leqslant L} \mathbb{P}\left[(n_k, c_j)\right] \overset{\text{eq. 4}}{=} \frac{d_{n_k} \sum_{1\leqslant j\leqslant L} d_{c_j}}{\sum_{1\leqslant x<y\leqslant N} d_x d_y}. \tag{5}$$

Equation 5 defines all p_k $1 \leqslant k \leqslant D$. The null coupon probability is defined as follows (1st term: collectors meeting each other, 2nd term: non-collectors meeting each other, 3rd term: collectors meeting non-destinations):

$$\begin{aligned} p_0 &= \mathbb{P}\left[\bigcup_{1\leqslant i<j\leqslant L} (c_i, c_j) \cup \bigcup_{1\leqslant i<j\leqslant N-L} (n_i, n_j) \cup \bigcup_{\substack{1\leqslant i\leqslant L \\ 1\leqslant j\leqslant N-L-D}} (c_i, n_{D+j})\right] \qquad (6) \\ &\overset{\text{mut. ex.}}{=} \sum_{1\leqslant i<j\leqslant L} \mathbb{P}\left[(c_i, c_j)\right] + \sum_{1\leqslant i<j\leqslant N-L} \mathbb{P}\left[(n_i, n_j)\right] + \sum_{\substack{1\leqslant i\leqslant L \\ 1\leqslant j\leqslant N-L-D}} \mathbb{P}\left[(c_i, n_{D+j})\right] \\ &\overset{\text{eq. 4}}{=} \frac{\sum_{1\leqslant i<j\leqslant L} d_{c_i} d_{c_j} + \sum_{1\leqslant i<j\leqslant N-L} d_{n_i} d_{n_j} + \sum_{\substack{1\leqslant i\leqslant L \\ 1\leqslant j\leqslant N-L-D}} d_{c_i} d_{n_{D+j}}}{\sum_{1\leqslant x<y\leqslant N} d_x d_y}. \end{aligned}$$

Thus, we have a fully defined coupon collector probability vector. The vector is a function of the chosen set of *permanent message carriers*, i.e., the chosen collector $\mathcal{L}$ and it does not include the null coupon: $p(\mathcal{L}) = (p_1(\mathcal{L}), \ldots, p_D(\mathcal{L}))$. ∎

[4] We consider the case where $\mathcal{L} \cap \mathcal{D} = \varnothing$. Otherwise, some of the carriers (or all, if $D \leqslant L$) could be trivially assigned among the destinations.

We will now define an order relation for probability vectors of this type.

Lemma 3. *The binary relation $\leqslant$ defined on $X = \{p(\mathcal{L}) \mid \mathcal{L} \in \wp(\mathcal{N})$ and $|\mathcal{L}| = L\}$*

$$p(\mathcal{L}_a) \leqslant p(\mathcal{L}_b) \quad \Leftrightarrow \quad \min(p(\mathcal{L}_a)) \leqslant \min(p(\mathcal{L}_b)). \tag{7}$$

is a total order relation.

Proof. The minimum of a vector is a real number, hence our relation is equivalent to the total order relation $\leqslant$ on $\mathbb{R}$. Thus it is a total order relation. ■

Using the total order relation, we frame a lemma on the monotonicity of $p(\mathcal{L})$.

Lemma 4. *The coupon collector probability vector, $p(\mathcal{L})$, is monotonous:*
*Broadcast: $p(\mathcal{L})$ increases with $\sum_{c\in\mathcal{L}} d_c$, **with skewed degree distributions**.*
*Multi-, Anycast: $p(\mathcal{L})$ **always** increases with $\sum_{c\in\mathcal{L}} d_c$.*

The proof and broadcast conditions for monotonicity are in the Appendix A.

In the next section, we analyze the variable T_s. This represents the number of contacts until the D distinct coupons have been collected for broadcast and multicast and the number of contacts until one coupon has been collected for anycast.

3.3 Expected *-Cast Time

Using the probability distributions, we express the expected number of contacts until successful message dissemination. As before, the vector of collection probabilities is: $p(\mathcal{L}) = (p_1(\mathcal{L}), \ldots, p_D(\mathcal{L}))$, where $p_0 + \sum_{i=1}^{D} p_i = 1$.

Broadcast (Def. 2) and Multicast (Def. 3). For broadcast and multicast, the expected number of contacts until successful message dissemination is

$$\mathbb{E}[T_s] = \sum_{k=1}^{D} (-1)^{k+1} \sum_{1\leqslant x_1<\ldots<x_k\leqslant D} \frac{1}{p_{x_1} + \cdots + p_{x_k}} \overset{[7]}{=} \int_0^1 \left(1 - \prod_{i=1}^{D} (1 - t^{p_i})\right) \frac{dt}{t}. \tag{8}$$

by the inclusion-exclusion principle. Note that this formula holds regardless of whether $\sum_{k=1}^{D} p_k = 1$ holds (here, it does not). This function has been studied extensively in the mathematics literature and not only [21, 6, 1, 12]. In this study, we want to prove that it is positive and decreasing in p. As shown in Lemma 4, maximizing p amounts to maximizing the degrees of the carrier nodes. Therefore, we will also prove that choosing nodes of maximum degree as carriers is the optimum solution in terms of delivery delay.

Theorem 1. *With skewed degree distributions, for each $N \in \mathbb{N}$ and for each $L \in \mathbb{N}$, $L < N$, the function $f_D(p) = \mathbb{E}[T_s]$ is positive and decreasing in p, using the order relation defined in Lemma 3.*

Proof. By a simple change of variable ($t = e^{-u}$) in the integral in equation 8, $f_D(p)$ can also be written as

$$f_D(p) = \mathbb{E}[T_s] = \int_0^{\infty} \left(1 - \prod_{i=1}^{D} \left(1 - e^{-tp_i}\right)\right) dt. \tag{9}$$

In [7], Borwein et al. showed that the key to proving an entire series of interesting properties of this function (including monotonicity) lies in a convenient way of writing it, derived from equation 9. Let X_i, $(1 \leqslant i \leqslant D)$ be independent positive exponentially distributed random variables with parameter $\lambda = 1$. Then,

$$f_D(p) = \mathbb{E}[T_s] = \int_0^{\infty} \left(1 - \prod_{i=1}^{D} \left(1 - e^{-tp_i}\right)\right) dt = \mathbb{E}\left(max\left(\frac{X_1}{p_1}, \ldots, \frac{X_D}{p_D}\right)\right). \tag{10}$$

Then, from the representation in equation 10, it is straightforward that $f_D(p)$ is positive and decreasing with the increase of all p_1 to p_D.

In the case of *multicast* and *anycast*, all p_1 to p_D increase with the increase of $\sum_{c \in \mathcal{L}} d_c$ and the proof is finished.

For the *broadcast* case, all p_1 to p_D except one, will increase with the increase of $\sum_{c \in \mathcal{L}} d_c$ (see proof of Lemma 3). With a skewed degree distribution, the network has numerous nodes of minimum degree. This means, a considerable number of the probabilities p_1 to p_D will be equal and of minimum value. These minimum probabilities will have an overwhelming mass in $f_D(p)$. Together with all the other increasing probabilities, they will almost surely, largely outbalance the unique decreasing value of the vector $p(\mathcal{L})$.

Hence, $f_D(p)$ will always decrease with the increase of vector p, where the order of vectors $p(\mathcal{L})$ is defined as in equation 7 of Lemma 3. ■

Anycast (Def. 4). As far as anycast from Definition 4 is concerned, the problem is much simpler. Indeed, for the anycast case, one only need to deliver the message to one node in the destination set $\mathcal{D}$. Hence, the showing time, T_s is a geometric random variable with success probability

$$p_{\text{any}} = \mathbb{P}\left[\bigcup_{\substack{1 \leqslant i \leqslant D \\ 1 \leqslant j \leqslant L}} (n_i, c_j)\right] \overset{\text{mut. ex.}}{=} \sum_{\substack{1 \leqslant i \leqslant D \\ 1 \leqslant j \leqslant L}} \mathbb{P}[(n_i, c_j)] \overset{\text{eq. 4}}{=} \frac{\sum_{i=1}^{D} d_{n_i} \sum_{j=1}^{L} d_{c_j}}{\sum_{1 \leqslant x < y \leqslant N} d_x d_y}. \tag{11}$$

Therefore, the expected value of T_s is

$$\mathbb{E}[T_s] = \frac{1}{p_{\text{any}}} = \frac{\sum_{1 \leqslant x < y \leqslant N} d_x d_y}{\sum_{i=1}^{D} d_{n_i} \sum_{j=1}^{L} d_{c_j}}, \tag{12}$$

and it is obviously also decreasing with the increase of vector p, where the order of vectors $p(\mathcal{L})$ is defined as in equation 7. Finally, Theorem 1 and equation 12 allow us to conclude that choosing the nodes of maximum degree as permanent message carriers does indeed minimize the showing time portion of the delivery time in all message dissemination schemes considered.

4 Experimental Evaluation

We provide here a brief experimental verification of our main finding in Section 3.3: having maximum degree nodes as carriers minimizes the showing time.

4.1 Contact Generators

We use three contact generators: two real mobility traces and one synthetic. The synthetic contacts are based on a scale-free graph model. The real mobility traces originate from two data collection experiments conducted by universities.

Scale Free Contacts: The contacts whereof the aggregation results in a scale free (SF) social graph are generated as follows. First, each node is assigned a popularity according to a power law (exponent 3). Then, the two nodes participating in a contact are randomly and independently chosen according to their respective popularities. We use a scale free model with 500 nodes and 50 000 contacts.

ETH Contacts: The first trace comes from an experiment of ETH Zürich [18]. 20 students and staff working on the same floor of an ETH building carried 802.11-enabled devices for 5 days. Every 0.5 s, each device sent a beacon message, the reception of which was logged by all devices in 802.11 radio proximity. This trace contains more than 23 000 reported contacts and is unique in terms of time granularity and reliability. Although the ETH trace measurement period spans a relatively short time, there are on average more than 1 000 contacts per device. This is comparable in number of contacts to similar longer traces.

MIT Contacts: The second trace comes from the Reality Mining [10] project. 97 students and employees of MIT were equipped with mobile phones scanning every 5 minutes for Bluetooth devices in proximity during 9 months. This trace is unique in terms of number of devices and duration. Nevertheless, with a time granularity of 5 minutes, many short contacts were presumably not logged. For our simulations, we cut the trace at both ends and used 100 000 contacts reported between September 2004 and March 2005. Note that this time period contains holidays and semester breaks and thus still captures varying user behavior.

For both real mobility traces we ignored logged timing information and just ordered the reported contacts according to their start times (i.e., slotted contacts). We obtain the social graph from the contacts using the method discussed and the results (optimal parameters) obtained in [14].

4.2 Simulation Results

To confirm our analysis, we evaluated the performance of constrained broadcast and multicast as described in Definitions 2 and 3, for Uniform Randomly chosen carriers (UR) versus message carriers with Highest Degree (HD) in the network.

For each contact generator, the simulation has a warmup period, to allow the collection of information and a cool down period, to allow the messages created last to be in the simulation for one TTL, as well. The TTLs were found empirically, with the aim of a reasonable coverage. Messages are generated randomly with probability 0.05 at each contact, during the period between the warmup and the cool down times. For multicast, we consider groups of size 25% of all nodes, chosen uniformly at random.

The two metrics considered are the **coverage**: the percentage of nodes in $\mathcal{D}$ receiving the message within TTL steps, and the **number of delivered messages** over time. This last metric clearly captures the average delay of the two strategies, as well.

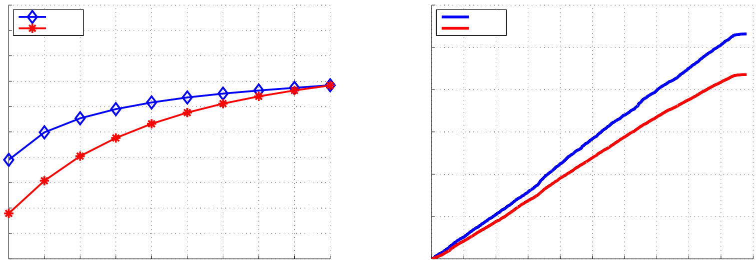

(a) # carriers vs. % coverage (b) Time vs. # delivered msg. (4 carriers)

Fig. 1. Broadcast Results (MIT trace)

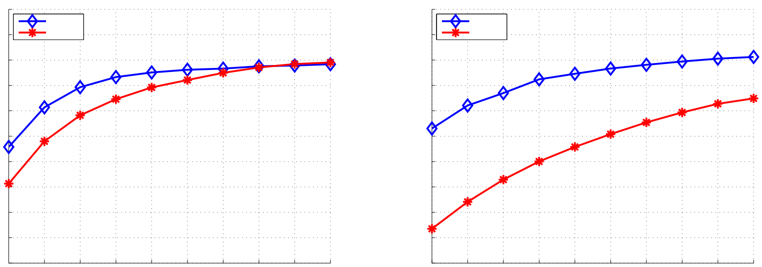

(a) ETH trace (b) Scale-free model

Fig. 2. Multicast Results: # carriers vs. % coverage

Figure 1 shows the two metrics for broadcast using the MIT trace. Plots for the ETH trace and the scale-free model are qualitatively similar and were left out due to space limitations. Figure 1(a) shows that the coverage is indeed better with the HD strategy. However, as the number of carriers approaches the total number of nodes in the network, the two curves converge, since the two sets of carriers (HD based and UR chosen) overlap more and more.

Figure 1(b) shows the number of messages delivered during the trace, for one point in Figure 1(a). It confirms that the HD scheme does indeed take less time than the UR scheme to deliver the same amount of messages. Figures 2(a) and 2(b) show the coverage for multicast on the ETH trace and respectively, the scale-free model. Again, HD has better coverage and delivery time in both cases.

As a final note, we observed that, using the two traces with a higher numbers of carriers, often the UR scheme slightly outperforms the HD scheme. We believe it is an effect of one social relation, our current analysis does not account for: node communities. Indeed, if all high degree nodes are part of the same community, the remaining communities will not have any member who is a carrier and will thus have a low chance of receiving the message. The UR scheme, however, has better odds of sampling carriers from smaller communities, hence the better performance with sufficient carriers. Accounting for node communities is the next step in our work.

5 Conclusion

Our goal in this paper has been to take a first look into a neglected area of DTN routing: *-cast. More precisely, we have undertaken a study of broad-, multi- and anycast in DTNs using social information. We identified a class of relevant optimization problems for *-cast, where the amount of resources (message copies) is constrained, and the optimization goal is to allocate resources to relays so as to minimize the delay until all nodes in the target *-cast group receive the message.

To solve these problems, we used social network analysis to render complex mobility patterns into a graph, transforming the problem into choosing an optimal set of vertices. Our first contribution is a probabilistic model that maps this graph to future contact probabilities between nodes. Moreover, in a setting where the degree distribution of this graph is skewed and known a priori, we prove, using a coupon collector analogy, that when the highest degree nodes are message relays the expected *-cast delivery time is optimum. We corroborate our findings using an evaluation of this policy on both synthetic and real mobility traces.

This is merely a first step into an area of many interesting, open issues. With the insight gained from this study, our next aim is to prove optimal policies for the problem where nodes are met online and degrees of nodes not yet met are unknown. Moreover, in practice, social graphs are expected to have high clustering coefficients and community structure. This makes the choice of highest degree nodes as carriers not necessarily optimal, due to the potential similarities among the neighborsets of high degree nodes. We are currently studying optimal policies accounting for both degrees and communities to optimally allocate resources.

References

1. Adler, I., Oren, S., et al.: The coupon-collector's problem revisited. J. Appl. Probab. 40(2), 513–518 (2003)
2. Aiello, W., Chung, F., et al.: A random graph model for massive graphs. In: STOC 2000, pp. 171–180. ACM, New York (2000)
3. Albert, R., Barabási, A.L.: Statistical mechanics of complex networks. Rev. Mod. Phys. 74(1), 47–97 (2002)
4. Altman, E., Başar, T., et al.: Optimal monotone forwarding policies in delay tolerant mobile ad-hoc networks. In: ValueTools 2008, pp. 1–13. ICST, Brussels (2008)
5. Boldrini, C., Conti, M., et al.: ContentPlace: social-aware data dissemination in opportunistic networks. In: MSWiM 2008, pp. 203–210. ACM Press, New York (2008)
6. Boneh, S., Papanicolaou, V.: General asymptotic estimates for the coupon collector problem. J. Comput. Appl. Math. 67(2), 277–289 (1996)
7. Borwein, J., Affleck, I., et al.: Convex? Convex! Convex II. SIAM: Problems and Solutions (2000)
8. Daly, E., Haahr, M.: Social network analysis for routing in disconnected delay-tolerant MANETs. In: MobiHoc 2007, pp. 32–40. ACM Press, New York (2007)
9. Dubois-Ferriere, H., Grossglauser, M., et al.: Age matters: efficient route discovery in mobile ad hoc networks using encounter ages. In: MobiHoc 2003, pp. 257–266. ACM Press, New York (2003)
10. Eagle, N., Pentland, A.: Reality mining: sensing complex social systems. Personal Ubiquitous Comput. 10(4), 255–268 (2006)

11. Feller, W.: An introduction to probability theory and its applications, 2nd edn., vol. 2. John Wiley & Sons, Chichester (1957)
12. Flajolet, P., Gardy, D., et al.: Birthday paradox, coupon collectors, caching algorithms and self-organizing search. Disc. Appl. Math. 39(3), 207–229 (1992)
13. Gao, W., Li, Q., et al.: Multicasting in delay tolerant networks: a social network perspective. In: MobiHoc 2009, pp. 299–308. ACM, New York (2009)
14. Hossmann, T., Legendre, F., et al.: From Contacts to Graphs: Pitfalls in Using Complex Network Analysis for DTN Routing. In: INFOCOM NetSciCom 2009, pp. 1–6 (2009)
15. Hui, P., Chaintreau, A., et al.: Pocket switched networks and human mobility in conference environments. In: SIGCOMM WDTN 2005, pp. 244–251. ACM, New York (2005)
16. Hui, P., Crowcroft, J., et al.: Bubble rap: social-based forwarding in delay tolerant networks. In: MobiHoc 2008, pp. 241–250. ACM Press, New York (2008)
17. Leguay, J., Friedman, T., et al.: Evaluating Mobility Pattern Space Routing for DTNs. In: INFOCOM 2006, pp. 1–10 (2006)
18. Lenders, V., Wagner, J., et al.: Measurements from an 802.11b Mobile Ad Hoc Network. In: WOWMOM 2006, pp. 519–524. IEEE Computer Society, Washington (2006)
19. Lindgren, A., Doria, A., et al.: Probabilistic routing in intermittently connected networks. SIGMOBILE Mob. Comput. Commun. Rev. 7(3), 19–20 (2003)
20. Mihail, M., Saberi, A., et al.: Random walks with lookahead on power law random graphs. Internet Math. 3(2), 147–152 (2006)
21. Nath, H.: Waiting Time in the Coupon-Collector's Problem. Austr. & New Zeal. J. of Stat. 15(2), 132–135 (1973)
22. Newman, M.E.J.: The structure and function of complex networks. SIAM Review 45, 167 (2003)
23. Spyropoulos, T., Psounis, K., et al.: Spray and wait: an efficient routing scheme for intermittently connected mobile networks. In: SIGCOMM WDTN 2005, pp. 252–259. ACM Press, New York (2005)
24. van der Hofstad, R.: Random Graphs and Complex Networks (2009)
25. Zhang, Z.: Routing in intermittently connected mobile ad hoc networks and delay tolerant networks: overview and challenges. IEEE Communications Surveys & Tutorials 8(1), 24–37 (2007)
26. Zhao, W., Ammar, M., et al.: Multicasting in delay tolerant networks: semantic models and routing algorithms. In: SIGCOMM WDTN 2005, pp. 268–275. ACM, New York (2005)

A Proof of Lemma 4

Proof. We want to prove that, as we increase $\sum_{c\in\mathcal{L}} d_c$, $p(\mathcal{L})$ will also increase, which in light of Lemma 3, means that $\min(p(\mathcal{L}))$ will increase.

Suppose we replace one collector node, c_i with one non-collector node n_j. Denote the new collector set $\mathcal{L}'$, n_j is relabeled c_i' and c_i is relabeled n_j'. Assume $d_{c_i} < d_{n_j}$. It follows that

$$\sum_{c\in\mathcal{L}} d_c < \sum_{c\in\mathcal{L}'} d_c, \tag{13}$$

We must prove that $\min(p(\mathcal{L})) < \min(p(\mathcal{L}\prime))$.

Multi-, Anycast: We start with the easier case of multicast and anycast, with $\mathcal{L} \cap \mathcal{D} = \varnothing$. Otherwise, some of the carriers (or all, if $D \leqslant L$) could be trivially assigned among the destinations. Under this condition, $c_i \notin \mathcal{D}$. Therefore, n_j'

will not be added to the destination set and there will be no extra p'_j. Whether n_j is a destination or not is irrelevant, as it becomes a collector node and will receive a copy of the message for storage. By equations 13 and 5, **p_1 to p_D will all increase** and thus, so will their minimum $\min(p(\mathcal{L}'))$. That is $\min(p(\mathcal{L})) < \min(p(\mathcal{L}'))$ and the proof for multicast and anycast in Lemma 4 is finished.

Broadcast: In broadcast, $\mathcal{D} = \mathcal{N}$, therefore, all nodes are also destinations, consequently $c_i \in \mathcal{D}$. This makes the evolution of p_1 to p_D slightly more ambiguous than previously. Whereas above, the destination set is not affected by the replacement, here the destination set, $\mathcal{N} \setminus \mathcal{L}$, is inevitably altered.
In particular, node c_i will now be part of the destination set, whereas it was not, before the replacement. Removing c_i from the collector set means, c_i will no longer get a copy of the message for storage and it will have to be counted in the probabilities p_1 to p_D, i.e., there will be a new probability, p'_j. **The original probabilities will all increase** by equations 5 and 13. The fate of the new minimum, $p'_{\min} = \min(p(\mathcal{L}'))$, depends on $p'_j = \frac{d_{n'_j} \sum_{c\in\mathcal{L}'} d_c}{\sum\limits_{1\leqslant x<y\leqslant N} d_x d_y}$. Denote by $\min(d_{\mathcal{N}\setminus\mathcal{L}})$ and respectively $\min(d_{\mathcal{N}\setminus\mathcal{L}'})$, the minimum degree among the nodes in the destination set. Then, there are two possibilities:

a. $\boldsymbol{d_{n'_j} \geqslant \min(d_{\mathcal{N}\setminus\mathcal{L}})}$. The return of n'_j to the destination set does not change the minimum $\min(p(\mathcal{L}))$. It will have increased, but it will still correspond to the same node as before the replacement. Therefore, $p(\mathcal{L}) \leqslant p(\mathcal{L}')$.

b. $\boldsymbol{d_{n'_j} < \min(d_{\mathcal{N}\setminus\mathcal{L}})}$. The return of n'_j to the destination set could change the minimum. This can be determined by checking the sign of $p_{\min} - p'_{\min}$, the difference between the minimum before and respectively after the replacement. Denote $C = \sum_{c\in\mathcal{L}\setminus\{c_i\}} d_c = \sum_{c\in\mathcal{L}'\setminus\{c'_i\}} d_c$ and $X = \left(\sum_{1\leqslant x<y\leqslant N} d_x d_y\right)^{-1}$. According to equation 5, $p_{\min} = \min(d_{\mathcal{N}\setminus\mathcal{L}})\cdot\left(C + d_{n'_j}\right)\cdot X$ and $p'_{\min} = d_{n'_j}\cdot\left(C + d_{c'_i}\right)\cdot X$. Then,

$$p_{\min} - p'_{\min} = \left[C(\min(d_{\mathcal{N}\setminus\mathcal{L}}) - d_{n'_j}) + d_{n'_j}(\min(d_{\mathcal{N}\setminus\mathcal{L}}) - d_{c'_i})\right]\cdot X. \tag{14}$$

From the assumption of item b, the first term of the sum in equation 14 is positive. Moreover, as $\min(d_{\mathcal{N}\setminus\mathcal{L}})$ was the minimum before the replacement, clearly $\min(d_{\mathcal{N}\setminus\mathcal{L}}) \leqslant d_{n_j} = d_{c'_i}$ and thus, the second term of the sum in equation 14 is negative. This means that the monotonicity of $p(\mathcal{L})$ might not hold for a certain choice of parameters, when the positive term in equation 14 outweighs the negative term. However, as the discussion below will show, in practice, the monotonicity is always respected.

Discussion. To sum up the above proof, probability vectors $p(\mathcal{L})$ increase with $\sum_{c\in\mathcal{L}} d_c$, as per the order relation in Lemma 3, with one exception: item b of the broadcast case. Under an arbitrary degree distribution, this implies it might be optimal to have very low degree nodes as message carriers (i.e., they receive a message copy to store). This is merely an artifact of the fact that, currently, we analyze the showing time, T_s alone. It will be reconciled when we consider both the finding time, T_f and the showing time, T_s together. In practice, degree distributions are, more often than not, skewed. A skewed degree distribution means there are numerous minimum degree nodes. As is evident from equation 14, giving messages copies to small degree nodes worsens the probability vector. ∎