

Agenda-based routing in DTNs[†]

Mathias Boc¹ Anne Fladenmuller¹ Marcelo Dias de Amorim¹ Sergio Palazzo²

¹ LIP6/CNRS – Université Pierre et Marie Curie – 104 Avenue du Président Kennedy, 75016, Paris, France

² DIIT – Università di Catania – Viale Andrea Doria 6, 95125, Catania, Italy

We propose Price, an agenda-based greedy forwarding scheme for disruption-tolerant networks (DTN). Price combines greedy geographic forwarding with greedy contact-based forwarding by relying only on agendas of locations and agendas of contacts, which benefit from the inherent periodicity of mobility. While nodes construct their “social” environment of acquaintances through agendas of contacts, the pertinence of using greedy geographic forwarding is assessed by empirical evidences on the small world phenomenon in spatial mobility. The performance of the combination of these two forwarding modes in the Dartmouth College mobility data set comforts our analysis as we outperform PROPHET and Bubble in terms of delivery ratio and cost. Furthermore, Price shows equivalent performance of the centralized version of Bubble.

Keywords: Disruption-tolerant networks, routing, mobile environments, mobility prediction.

1 Introduction

Routing in disruption-tolerant networks is a challenging problem. In such networks, where mobile nodes show intermittent connectivity, an end-to-end path between a source and a destination may never exist. This implies that intermediate nodes store data messages (or bundles) until they have an appropriate opportunity to forward them towards the destination. Optimum delivery ratios can be achieved by flooding bundles, but at a two high overhead. As a consequence, more efficient routing/forwarding solutions have to be proposed.

Recently, a new class of forwarding algorithms relying on social and human interactions (e.g., mobility similarity, centrality, or communities) has emerged [LFC06, DH07, HCY08]. Such algorithms are interesting as they consider nodes through both their social sphere and more fundamental properties from graph theory. However, forwarding progression becomes tough when one considers wider environments where groups of nodes (communities) do not show any similarities. Unfortunately, most works in the literature rely on such an assumption. For instance, Leguay et al. propose a greedy forwarding scheme that relies on similarity of visited locations [LFC06]. They assume that a node can always find a relay that has more similarities with the destination. Daly et al. propose to use both ego-centrality and similarity, where bundles are transmitted to nodes having higher centrality to increase the probability of finding the destination [DH07]. However, the progression towards higher centrality values does not guarantee delivery if a destination that shows low centrality. Hui et al. have handled this problem by proposing a forwarding system based on communities and on both their intra-community (local) and global centralities [HCY08]. Nevertheless, the authors evaluated their approach using pre-computed centrality values, which does not consider how mobility behaviors and specifically activity patterns influence the attribution of a node to one community with a decentralized algorithm. Furthermore, such a solution might misjudge the popularity of a node in one environment if it has high contact events in another environment.

[†]This work has been partially supported by the ANR project Crowd under contract ANR-08-VERS-006 and by the European Commission in the framework of the FP7 Network of Excellence NEWCOM++ under contract 216715.

2 Price: Agenda-based routing in a DTN

In the literature, no solutions have been proposed to explore the situation where nodes have common locations with the destination. To fill this gap, we propose Price, a combination of geographic and social-based forwarding for disruption-tolerant networks. Price relies on the two available information that help characterize a node: (i) its geographic positions and (ii) its contacts. Through contacts, it is possible to define the relative distances among nodes according to their “acquaintances”. Through geographic information, we can benefit from the fact that users are most of the time restricted to a confined area (e.g., a city, a campus, at home, in office).

The idea behind Price is to rely on geographic forwarding when a source/relay node and its neighborhood are not in the ecosystem of the destination (i.e., when they cannot operate in contact-based mode). In this way, we help the bundle to circulate in the network rather than waiting for a node with a (greater) similarity with the destination node. In the following, we explain the two operation modes of Price.

Algorithm 1: Contact-based forwarding procedure

```

begin
   $DTW_i \leftarrow deliveryWindow(i, dest)$ 
  if ( $DTW_i < EDTW$ ) then  $EDTW \leftarrow DTW_i$ 
  forall neighborNode  $j$  do
     $DTW_j \leftarrow bestDeliveryWindow(j, dest)$ 
  if  $DTW_j < EDTW$  then
     $EDTW \leftarrow DTW_j$ 
     $sendCopy(bundle, j)$ 
  if ( $DTW_i == NULL$ ) then  $i.drop(bundle)$ 
end

```

Contact-based forwarding mode. A bundle transmission is performed whenever the relay node or at least one of its neighbors has a reasonable probability to get in contact with the destination. We consider that node i is likely to be in contact with node j if they have been previously in contact. The prediction algorithm assumes periodic mobility patterns and is supported by the use of *agendas of contacts* [Boc09]. All predicted contacts (which appear in the agenda of contacts of a node i) are nodes belonging to the “ecosystem” and are thus i ’s acquaintances. Based on the agenda of contacts, a delivery time window (DTW) allows defining when two nodes will likely be contact again. Within the ecosystem of the destination node, bundle replicas are generated with the purpose of reducing the expected delivery time window (EDTW) of the bundle, as shown in Algorithm 1. When the destination is no further an acquaintance of the relay node (i.e., the destination node is no more in the agenda of contact of the relay node), the bundle is dropped. The latter prevents uncontrolled diffusion of replicas in the network.

Algorithm 2: Geographic forwarding procedure

```

begin
   $D_i \leftarrow agendaDistance(i, P_{dest})$ 
  if  $D_i > r$  then
    forall neighborNode  $j$  do
       $D_j \leftarrow bestAgendaDistance(j, P_{dest})$ 
    end
    if  $D_j < (D_i - r)$  and  $isSource(bundle, i)$  then  $sendCopy(bundle, j)$ 
    if  $D_j < (D_i - r)$  and  $isNotSource(bundle, i)$  then  $send(bundle, j)$ 
  end
end

```

Geographic forwarding mode. By analyzing its own mobility, any node is able to determine a region where it spends most of its active time or where it is in contact with most of its acquaintances. For now,

we consider that these two locations are the same; we refer to them as “home region”. We rely on the assumption that all nodes have a home region and the closer we get to this region the more likely we find its acquaintances. Potential corresponding nodes can fetch the home region of other nodes, provided by a location management system (out of the scope of this paper). With the home region’s coordinates as geographic destination, transmissions are operated greedily using spatial mobility predictions. Predictions are assured by *agendas of locations* presented and evaluated in [Boc09]. Hence, the progress achieved by a neighbor is independent from its current positions, but dependent on its predicted positions (cf., Algorithm 2). The performance of this forwarding mode relies on the capacity of the bundles to navigate in the environment based on geographic information. As soon as a relay gets in contact with a node that has the destination node as an acquaintance, the bundle forwarding strategy is switched to the contact-based mode.

3 Evaluation

Data set. We consider the Dartmouth campus data set to evaluate our approach. We select mobility information over one month and focus on a subset composed of the 200 most active nodes (in terms of duration of presence in the network). To obtain spatial geographic trajectories, we adopt the scheme proposed by Kim et al. to retrieve the coordinates of patterns of mobility by using a Kalman filter [KKK06]. In the resulting Euclidian space, we consider that two nodes are in contact if their distance is within 50 meters. For contact and spatial predictions, time is sliced in slots of 1 hour and the prediction period is set to 24 hours [Boc09].

Greedy geographic reachability. When a greedy forwarding algorithm is able to perform high delivery ratio while finding short path lengths towards a destination, the network is considered as searchable. We first evaluate the capacity of the greedy geographic forwarding scheme to reach, alone, nodes’ home locations. A location is defined by its geographic coordinates and a radius of 50 meters. We observe increasing delivery ratios when node density increases: with 105 nodes, we deliver 20% of bundles; with 200 nodes, we deliver 67%; and with 300 nodes, 81% of the bundles reach the destination. Fig. 1 shows the path lengths (number of hops) of the delivered bundles. A bundle that makes a progression of 50 meters at each transmission would follow a squared root function. One observes that the resulting path lengths are only lightly correlated to the density of nodes and the traveled distances. They follow closely a logarithmic function rather than a squared root one.




Fig. 1: Path lengths of delivered bundles to home locations.

Performance analysis. We evaluate the efficiency of Price to deliver bundles to destinations. We gauge the performance of Price with 4 well-known routing alternatives: flooding, wait, Prophet [LDS03], and Bubble [HCY08]. Flooding consists in broadcasting bundles replicas to all neighbor nodes at each contact. In wait, the source node does not do anything until it is in contact with the destination to deliver the bundle. In Prophet, the routing strategy relies on the history of encounters and probability transitivity to predict the probability of each node to deliver bundles to destinations. We use the same parameters proposed by the authors and set the aging function to a period of 24 hours. For Bubble (cf., Section 1), we consider the two algorithms proposed by the authors: the global/local centrality value of all nodes and the communities are computed offline (called “Bubble” hereafter) and the algorithm where the global centrality value of each node is computed with a sliding window (called “Bubble-C” hereafter).[‡]

The delivery delay is computed from the first contact of the source with any node. We generate data traffic of 1,000 bundles between randomly chosen pairs of nodes at the beginning of the observation period.

[‡] We consider cliques of sizes 3 and a contact threshold of 10,800 seconds (to create communities) and a sliding window of 6 hours to compute the global centrality values.




Fig. 2: Performance comparison between Price and other approaches (cumulative distribution functions).

Each emulation is run 20 times.

Fig. 2(a) presents the delivery success ratio as a function of the delay. Flooding provides the upper-bound on the delivery ratio. It delivers 99% of bundles during the whole period of observation. In terms of delay, it is important to note that it takes one week to deliver 90% of bundles. On the other hand, wait delivers only 9% of the bundles. When compared to Prophet, Price delivers 27% more bundles. 45% of bundles are delivered within 1 week, while Prophet delivers less than 13% during the same period. As for Bubble, we were surprised to observe that Price delivers only 9% less bundles when compared with the centralized version of Bubble! When compared on the same basis, i.e., with Bubble-C, Price outperforms the latter by delivering 30% more bundles.

In terms of cost, Price and Bubble-C are the most efficient ones. To compute the cost, we sum up all transmissions operated (including replicas) and normalize this value with the number of bundles delivered in the network. Fig. 2(b) shows that the difference between Price and flooding is about 89. The difference is lower when we consider Prophet (42) and Bubble (7.6), but big enough to consider that only Price and Bubble-C are the most cost efficient approaches here (but recall that Bubble-C achieves much lower delivery ratios).

4 Conclusion

The combination of different forwarding algorithms is important to make circulating data bundles in disruption tolerant networks. We argued and showed that the combination of geographic and contact-based forwarding strategies is efficient in such context. In a future work, we plan to develop an analytical model to evaluate the importance of each parameters.

References

- [Boc09] M. Boc. Profile of mobility: user-centric networking. *Thesis*, November 2009.
- [DH07] E.M. Daly and M. Haahr. Social network analysis for routing in disconnected delay-tolerant MANETs. In *ACM Mobihoc*, pages 32–40, Montreal, Quebec, Canada, 2007.
- [HCY08] P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: social-based forwarding in delay tolerant networks. In *ACM Mobihoc*, pages 241–250, Hong Kong, China, May 2008.
- [KKK06] M. Kim, D. Kotz, and S. Kim. Extracting a mobility model from real user traces. In *IEEE Infocom*, Barcelona, Spain, April 2006.
- [LDS03] A. Lindgren, A. Doria, and O. Schelén. Probabilistic routing in intermittently connected networks. *Proceedings of ACM Sigmobile*, 7(3):19–20, 2003.
- [LFC06] J. Leguay, T. Friedman, and V. Conan. Evaluating mobility pattern space routing for dtms. In *IEEE Infocom*, pages 1–10, Barcelona, Spain, April 2006.