Optimal Forwarding Strategies for Delay Tolerant Networks with Multiple Classes and Beaconing Costs

Nicola Basilico, Nicola Gatti, Matteo Cesana

***Abstract*—Delay–tolerant networks (DTNs) are sparse and/or highly mobile wireless ad hoc networks which assure no continuous connectivity. One central problem in DTNs is is the routing of packets from a source towards the desired destinations. The literature provides a number of results for very restrictive scenarios in which mobile nodes have the same mobility characteristics and in which only forwarding costs are taken into account. In this work, we extend this result to more concrete situations in which mobile nodes have different mobility characteristics and can use different technologies and also beaconing costs are taken into account. In particular, we provide a greedy algorithm that find the source's optimal forwarding policy in polynomial time and we provide some illustrations of our algorithm.**

***Index Terms*—Delay–tolerant networks, epidemic routing, multi–class mobiles**

I. Introduction

DELAY–tolerant networks (DTNs) (also known as disruptive–tolerant networks) are sparse and/or highly mobile wireless ad hoc networks which assure no continuous connectivity. Examples of such networks are those operating in mobile or extreme terrestrial environments, or planned networks in space. Disruption may occur because of the limits of wireless radio range, sparsity of mobile nodes, energy resources, attack, and noise. One central problem in DTNs is is the routing of packets from a source towards the desired destinations. When no information is available *a priori* over the mobility pattern of the nodes, a common technique for overcoming lack of connectivity is to disseminate multiple copies of the packet in the network: this enhances the probability that at least one of them will reach the destination node within a given temporal deadline. This is referred to as *epidemic–style* forwarding, because, alike the spread of infectious diseases, each time a packet–carrying node encounters a new node not having a copy thereof, the carrier may infect this new node by passing on a packet copy; newly infected nodes, in turn, may behave similarly. The destination receives the packet when it meets an infected node.

The literature provides a number of results for very restrictive scenarios in which mobile nodes have the same mobility characteristics and in which only forwarding costs are taken into account [2]. In such situations the calculation of the optimal forwarding policies of the source is easy and can be accomplished in closed form. In this work, we extend this result to more concrete situations in which mobile nodes have different mobility characteristics and can use different technologies and also beaconing costs are taken into account. We briefly summarize our main contributions:

- we extend the model presented in [2] introducing multiple classes of mobile users, introducing beaconing costs, and different transmission technologies;
- we show that optimal policies cannot be found in closed form as instead it is with only one class and without beaconing costs;
- we provide a greedy algorithm that in polynomial time finds the optimal policies.

The report is organized as follows. In Section 2, we resume the main works presented in the literature. In Section 3, we introduce the model described in [2] and we extend it. In Section 4, we describe our algorithm. In Section 5, we show some experimental results produced with our algorithm. Finally, Section 6 concludes the report.

II. State of the art

A number of papers deal with the problem of studying forwarding strategies for DTNs. The basic model is the following. There is one or more packets that must flow from a source to a sink. Each packet is associated with a temporal deadline, beyond which the packet has no value. The source and each mobile node have a potentially different strategy in terms of packet forwarding. The main problems studied in the literature are the following.

Problem 1 (Source problem): The problem is to determine the optimal forwarding strategy of the source in terms of maximizing the probability to deliver the packet by the temporal deadline subject to energy constraints. A crucial aspect is the study of strategies that depend on several parameters (e.g., time, mobility class of the mobile nodes, community class of the mobile nodes) and the analysis of the cost of such parameters (e.g., beaconing costs or forwarding cost).

This problem is customarily modeled as an optimization problem.

Problem 2 (Mobile node problem): The problem is to determine the optimal forwarding strategy of the mobile nodes both when they are cooperative and when they are non–cooperative. When the nodes are cooperative, they act in order to maximize probability to deliver the packet by the temporal deadline under energy constraints. In the case the nodes are non–cooperative, they act in order to maximize their expected reward from delivering the packet under energy constraints.

This problem is customarily modeled as a cooperative/non–cooperative game theory problem.

		TWO–HOP	MULTI–HOP
SINGLE–PACK		[1], [2], [4]	[3], [9]
MULTI–PACK	limited buffer	[5]	
	unlimited buffure	[6], [7]	

TABLE I
CLASSIFICATION OF THE STATE OF THE ART.

In the following, we propose a classification of the current state of the art. We identify three main dimensions:

- *source vs mobile*: we classify the results in terms of strategies for the source or for the mobile nodes;
- *single–packet vs. multi–packet*: we classify the results in terms of strategies applicable when there is only one packet or more packets (in the case of multi–packet, we distinguish between unlimited buffer and limited buffer);
- *two–hop vs. multi–hop*: we classify the results in terms of strategies applicable only for situations with two hops (from the source to a mobile and from the mobile to the sink) and strategies addressing multi–hop routing schemes.

We report in Table I the main results presented in the literature.

In [2], the authors study a scenario with a single source and a number of mobile nodes that have the same mobility capacities. The forwarding routing is two hops. The problem is the determination of the optimal strategy of the source when an energy constraint is present over the packet forwarding (no beaconing cost is taken into account). The authors show that the best strategy the source can have is based on temporal threshold such that the source transmits the packet at all the time points that precede the temporal deadline and does not transmit after. In [4], the authors focus on the mobile nodes. There is only one packet and the routing is two hop. The nodes are considered selfish/non–cooperative and the problem is to incentivize the mobile nodes to forward the packet. The first node that delivers the packet receives a unitary reward. The only cost of the mobile nodes is due to beaconing. In addition, all the nodes have the same mobility capabilities. The authors study the best strategies by using evolutionary game theory tools.

In [3], the authors focus on multi–hop routing scenarios in the attempt to design incentives for the mobile nodes to avoid that these nodes behave strategically in terms of edge hiding and edge insertion. The proposed model works only when beaconing costs are not present and each mobile node is subject to the same cost. A similar technological, but not technical, approach is proposed in [9].

In [5], the authors study different routing strategies when the nodes have a limited buffer and provide an experimental comparison of the strategies. A similar work is presented in [6] except that here the authors introduce social information about the nodes.

Other related works dealing with cooperative aspects are [7] and [8].

III. MODEL

The model we develop shares several assumptions with the work proposed in [2] and extends it by considering a multi–class scenario. In the following, we introduce the notation and we briefly summarize [2].

We consider a scenario with one source and one sink node. Mobile nodes populate the environment and are divided into C different classes. Each class encodes the capabilities of a node in terms of mobility and transmission technology. In particular, a class i is associated with the following parameters:

- N_i is the number of nodes belonging to class i;
- R_i is the communication range of nodes in class i; this parameter relates to the transmission technology for that class;
- v_i the speed of mobile nodes in class i; this parameter relates to the mobility features for that class.

We consider a discrete representation of time that develops in a sequence of slots of fixed duration Δ. The k-th time slot corresponds to the time interval $[k\Delta, (k+1)\Delta)$.

A packet has to be delivered to the sink. We assume that, initially, the packet is only held by the source node and that it must be delivered by time τ, where τ refers to the packet's time to live. In our discrete time representation, we denote the total number of useful time slots as $K = \lfloor \tau/\Delta \rfloor$. By arriving in the proximity of the source or sink, mobile nodes make contact with them and get an opportunity to exchange data. We focus on a 2–hop routing scheme. That is, upon reception of a message from the source, a mobile node will forward it only if a contact is subsequently made with the sink. Nodes arrivals at the source are modeled through a multi–class Poisson process, where the arrival rate for nodes of class i is denoted by λ_i. Arrival rates are computed as a function of class parameters as proposed in [2]:

$$\lambda_i = \frac{8wR_iv_i}{\pi L^2}$$

(w is a constant value set to 1.3693, see [2] for the details).

Every time a contact is made between the source and a mobile node that did not receive the packet already, the source has to decide whether to forward the message or not. The forwarding policy is defined as $\mu(k)$, indicating the forward probability at time slot k. When a packet is transmitted from the source to a mobile, the probability that it will eventually be delivered to the sink before time $\tau\Delta$ is, in general, increased. However, transmissions introduce costs and, consequently, budget constraints. We distinguish between two cost factors:

- *transmission cost*: the energy consumption incurred to transmit messages;
- *beaconing cost*: the energy consumption deriving from connection control and signaling transmissions.

The transmission cost is directly proportional to the number of mobiles that get the message from the source. The upper bound on this number is denoted as ψ, and it represents a maximum budget for the forwarding policy. The beaconing cost paid per time slot is denoted as b_i. Since different classes might be characterized by different transmission technologies, this parameter is class–dependent.

Given this scenario, the objective is to compute the forwarding policy that maximizes the probability of delivering the message before time $K\Delta$ and that satisfies the budget constraint.

In the following, we provide a characterization of the optimal policy problem. First, we briefly summarize the contributions of [2] where an analytical method is provided for the single–class case. Then, we show how an analytical study is not possible for our multi–class extension and we provide an algorithm to deal with that.

A. The model proposed by Altman et al.

Altman *et al.* consider a setting with a unique class of mobile nodes. They define X_k as the random variable expressing the number of mobile nodes that have received the packet by time slot k and $Q_{k,k'}$ the probability that a mobile node does not receive any packet in time slots $k,\dots,k'$. The expected number of mobile nodes that receive a packet in time slots $0,\dots,k$ is:

$$\mathbb{E}[X_k] = N \cdot (1 - Q_{0,k})$$

where

$$Q_{k,k'} = e^{-\lambda\Delta \sum_{i=k}^{k'} \mu(i)}$$

An energy constraint over $\mathbb{E}[X_K]$ is posed as $\mathbb{E}[X_K] \leq \psi$, keeping into account only forwarding costs. The probability that a packet is delivered by k to the sink is:

$$F_D(k) = 1 - \prod_{i=0}^{k-1} X_i^* \cdot (\lambda\Delta)$$

where

$$X_k^*(s) = \mathbb{E}[e^{-sX_k}]$$

The authors show that $F_D(K)$ is strictly monotonically increasing in μ and that the optimal strategies μ, maximizing $F_D(K)$, are based on a threshold h such that:

$$\mu(k) = \begin{cases} 1 & k < h \\ \alpha & k = h \\ 0 & k > h \end{cases}$$

where $\alpha \in [0,1)$. The threshold h and the value α can be found in closed form as:

$$h = \left\lfloor -\frac{\log(1-\frac{\psi}{N})}{\lambda\Delta} \right\rfloor$$

$$\alpha = -\frac{\log(1-\frac{\psi}{N})}{\lambda\Delta} + \left\lfloor \frac{\log(1-\frac{\psi}{N})}{\lambda\Delta} \right\rfloor$$

B. Extension to multiple classes

Each class of mobile nodes will be characterized by a different λ_i and a different N_i. Similar as done by Altman *et al.*, we can define $X_{i,k}$ as the random variable expressing the number of mobile nodes of class i that have received the packet by time slot k and $Q_{i,k,k'}$ the probability that a mobile node of class i does not receive any packet in time slots $k,\dots,k'$. The source can have a different policy $\mu_i(k)$ for each different class of mobile nodes. As a result we have:

$$\mathbb{E}[X_{i,k}] = N_i \cdot (1 - Q_{i,0,k})$$

where

$$Q_{i,k,k'} = e^{-\lambda_i\Delta \sum_{j=k}^{k'} \mu_i(j)}$$

It can be easily seen that $F_D(k)$ is strictly monotonically increasing in μ_i and that the optimal strategies are based on a threshold as in the case with a single class except that the threshold h_i of each class can be different. More precisely, the probability that a packet is delivered by k to the sink is:

$$F_D(k) = 1 - \prod_{j=1}^{C} \prod_{i=0}^{k-1} X_{j,i}^* \cdot (\lambda\Delta)$$

where

$$X_{i,k}^*(s) = \mathbb{E}[e^{-sX_{j,kj}}]$$

For simplicity, we consider two extreme situations: the first in which all the classes of mobile nodes exploit the same technology and the second in which each class of mobile nodes exploits a different technology. In the first case, we need to search for the policy profile $(\mu_1,\dots,\mu_C)$ maximizing $F_D(K)$ and satisfying the constraint:

$$w \cdot \sum_i N_i \cdot (1 - Q_{i,0,K}) + \max_i \left\{ b \cdot \sum_{j=0}^{K} \mu_i(j) \right\} \leq \psi$$

where b is the beaconing cost per time slot and w is the cost of forwarding the packet. In the second case, the constraint to satisfy is

$$\sum_i w_i \cdot N_i \cdot (1 - Q_{i,0,K}) + \sum_i b_i \cdot \sum_{j=0}^{K} \mu_i(j) \leq \psi$$

where b_i is the beaconing cost per time slot of technology adopted by class i and w_i is the cost of forwarding the packet with the technology adopted by class i.

While with only one class of mobile nodes the optimal policy can be found in closed form, with multiple classes it is not possible even in the case without beaconing costs. Consider a simplified case with two classes, no beaconing costs, and policies are such that $\mu_1 = \mu_2 = \mu$. We can find the optimal policy by solving the equation:

$$N_1 \cdot (1 - e^{-\lambda_1\Delta \sum_{i=k}^{k'} \mu(i)}) + N_2 \cdot (1 - e^{-\lambda_2\Delta \sum_{i=k}^{k'} \mu(i)}) = \psi$$

However, such an equation is transcendent and cannot be solved exactly. When the constraint $\mu_1 = \mu_2 = \mu$ is relaxed, allowing μ_1 and μ_2 to be any, the problem is harder. Indeed, we would have one transcendent equation with two variables and many solutions would fulfill such an equation. Among all the solutions, we need to maximize the non–linear objective function $F_D(K)$.

IV. An algorithm for the multi–class scenario

In this section, for simplicity, we provide an algorithm for the case in which there is no beaconing cost and only one technology is used. We discuss below how the algorithm can be extended to the general case.

Let us consider a given set of threshold policies $\mu_1, \mu_2, \dots, \mu_C$ and, in particular, one of such policies denoted as μ_i. We indicate with μ_j any other different policy. Also, let us recall that h_i is the time slot such that for every $k > h_i$ it holds that $\mu_i(k) < 1$. We want to compute the number of additional time slots in which μ_i could keep transmitting without violating the budget constraint. If we call this number $\hat{\mu}_i$, this means that by setting $h_i = h_i + \hat{\mu}_i$ we will consume all the ψ budget units. We can compute it as:

$$\hat{\mu}_i = -\frac{\Phi_i}{\lambda_i \Delta} - \sum_{k=0}^{K} \mu_i(k)$$

where

$$\Phi_i = \log\left(1 - \frac{\psi - \sum_j N_j\left(1 - e^{-\lambda_j \Delta \sum\limits_{k=0}^{K} \mu_j(k)}\right)}{N_i}\right)$$

Starting from the initial state where all the policies are set to 0 for every slot, Algorithm 1 iteratively updates a policy i^*. Such a policy is selected at each iteration as the one for which the marginal gain m_i is maximum. This value quantifies the improvement we obtain by updating i and keeping the other policies unchanged. The update, performed at instruction 5, amounts to append to policy μ_i a probability value which is comprised between 0 and 1, depending on the current residual budget. Instruction 6 stores the marginal gain and instruction 7 undoes the policy update. The algorithm ends when all the budget units have been spent.

Algorithm 1 Find optimal policy

1: $\mu_1(k), \ldots, \mu_C(k) \leftarrow 0 \; \forall k$
2: **while** $\sum_i \mathbb{E}[X_{i,k}] \leq \psi$ **do**
3: $\quad m_1, \ldots, m_C \leftarrow 0$
4: $\quad$ **for** every $i \in C$ **do**
5: $\quad\quad \mu_i(h_i + 1) \leftarrow \min\{1, \hat{\mu}_i\}$
6: $\quad\quad m_i \leftarrow F_D(K)$
7: $\quad\quad \mu_i(h_i + 1) \leftarrow 0$
8: $\quad$ **end for**
9: $\quad i^* \leftarrow \arg\max_{i \in C} \{m_i\}$
10: $\quad \mu_{i^*}(h_{i^*} + 1) \leftarrow \min\{1, \hat{\mu}_{i^*}\}$
11: **end while**

The algorithm is greedy and its optimality follows from the fact that the objective function F_D is submodular. Furthermore, it can be observed that the computational time of the algorithm is $O(C \cdot K)$ and therefore it can be solved efficiently in polynomial time.

When also beaconing costs are present, the algorithm can be easily extended. The crucial issue in this case is that $\hat{\mu}_i$ is not given by a linear equation, but it requires the resolution of a transcendent equation with a single variable. This can be (approximately) accomplished by using the Newton algorithm.

V. Experimental evaluation

We performed a set of experiments with Algorithm 1 where we employed two different classes of mobile nodes. We have $N_1 = 5$ nodes belonging to *Class 1*, moving with a relatively slow speed of 1.5 m/s. *Class 2* is populated by $N_2 = 5$ nodes too, moving at a higher speed of 5 m/s. We fixed some other parameters: $L = 2500$, $R_1 = R_2 = 10$, $\Delta = 10$, $K = 10$. In Figure 1, we report the optimal policies returned by Algorithm 1 with different budget constraints. As expected, when the budget is small ($\psi = 2$ and $\psi = 4$), Class 2 is preferred since nodes in this class move faster and, consequently, exhibit a higher arrival rate at the source. When the budged limit is large ($\psi = 8$), the source can afford to transmit during the whole time period for both classes, in order to maximize the delivery probability. An interesting result is obtained in the intermediate situation, i.e., when $\phi = 6$. In this case, the optimal solution returns two different thresholds for the two classes, spending more budget for Class 2.

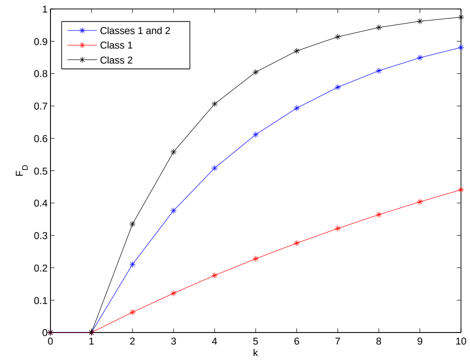

Fig. 2. Multi–class versus single–class cases

In Figure 2, we compare the trend of F_D obtained by the optimal multi–class policy with the same trend when the whole nodes population is concentrated in just one class. The reported results consider the case when $\phi = 6$. As it can be seen, combining nodes of these two classes gives an intermediate performance (in terms of delivery probability) which is comprised between the better performance of fast moving nodes and the worse performance of slow moving ones.

VI. Conclusions and future works

Delay–tolerant networks constitute a challenging scenario in which mobile nodes can communicate without continuous connectivity. One crucial problem is the development of forwarding policies when, as usual, no information is available on the mobility patterns of the mobile nodes. The literature developed epidemic approaches according which a packet is replicated and sent to many nodes by the source. However, the literature focuses on very narrow scenarios in which all the mobile nodes have the same mobility characteristics, no beaconing cost is present, and all the nodes adopt the same transmission technology.

In our work, we extended the state of the art, allowing mobile nodes to have different mobility characteristics, including beaconing cost, and allowing the nodes to adopt different transmission technologies. We showed that, differently from the cases studied in the state of the art, the optimal source's policies cannot be found in closed form, but searching algorithm are necessary. We showed in addition, that the optimization problem is not linear. However, a greedy algorithm can be used to find the optimal solution that requires polynomial time in the size of the problem. Therefore, optimal

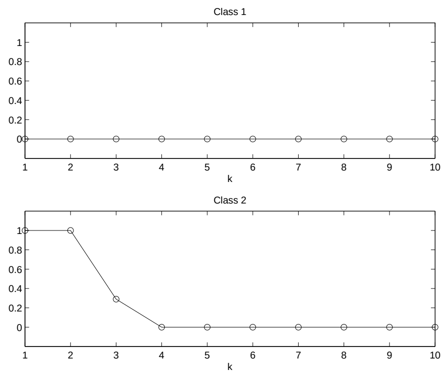

(a) $\psi = 2$

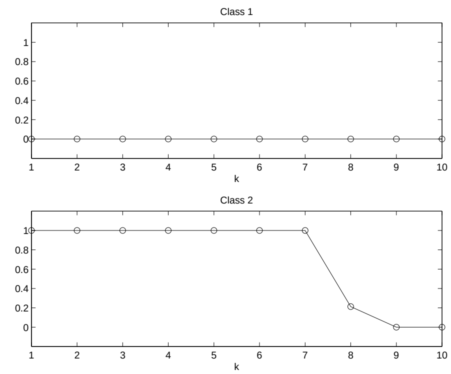

(b) $\psi = 4$

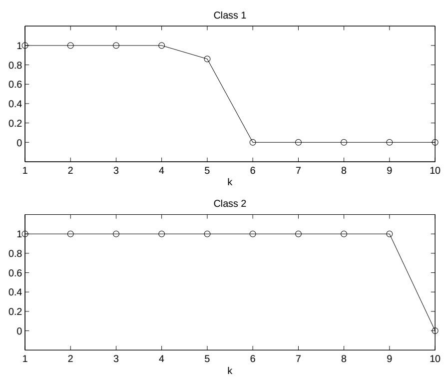

(c) $\psi = 6$

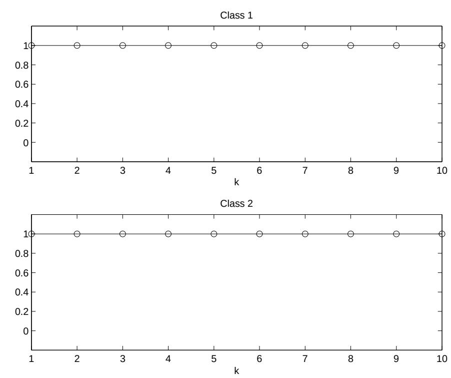

(d) $\psi = 8$

Fig. 1. Optimal policies in the multi–class case with different budget constraints; each graph reports μ_1 and μ_2.

policies can be found efficiently. Finally, we provided some experimental results produced by our algorithm.

In future, we will extend our work along many directions. We briefly list the main ones.

- We will complete the extension of the algorithm to the case with beaconing costs and different technologies, formally proving the optimality of the algorithm, and we will produce a thorough experimental evaluations.
- We will introduce different forms of constraints, e.g., the possibility to adopt simultaneously only one technology.
- We will study the policies of the mobile nodes when multiple mobility classes are present and how their policies affect the source's policies.
- We will develop an experimental settings with real mobile nodes and we will evaluate the optimal policies derived with our model.

Acknowledgment

The work is founded by the project PRIN 2009 People–Net.

References

[1] E. Altman. Competition and cooperation between nodes in delay tolerant networks with two hop routing. *Network Control and Optimization*, pages 264–278, 2009.

[2] E. Altman, G. Neglia, F. De Pellegrini, and D. Miorandi. Decentralized stochastic control of delay tolerant networks. In *INFOCOM 2009, IEEE*, pages 1134–1142. IEEE, 2009.

[3] B.B. Chen and M.C. Chan. Mobicent: a credit-based incentive system for disruption tolerant network. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. IEEE, 2010.

[4] R. El-Azouzi, F. De Pellegrini, and V. Kamble. Evolutionary forwarding games in delay tolerant networks. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt), 2010 Proceedings of the 8th International Symposium on*, pages 76–84. IEEE, 2010.

[5] S. Jain, K. Fall, and R. Patra. *Routing in a delay tolerant network*, volume 34. ACM, 2004.

[6] Q. Li, S. Zhu, and G. Cao. Routing in socially selfish delay tolerant networks. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9. IEEE, 2010.

[7] T. Ning, Z. Yang, X. Xie, and H. Wu. Incentive-aware data dissemination in delay-tolerant mobile networks. SECON, 2011.

[8] D. Niyato, P. Wang, W. Saad, and A. Hjorungnes. Coalition formation games for improving data delivery in delay tolerant networks. In *GLOBECOM 2010, 2010 IEEE Global Telecommunications Conference*, pages 1–5. IEEE, 2010.

[9] H. Zhu, X. Lin, R. Lu, Y. Fan, and X. Shen. Smart: A secure multilayer credit-based incentive scheme for delay-tolerant networks. *Vehicular Technology, IEEE Transactions on*, 58(8):4628–4639, 2009.