

Barter-based cooperation in delay-tolerant personal wireless networks

Levente Buttyán¹ László Dóra¹ Márk Félegyházi² István Vajda¹

¹ Laboratory of Cryptography and Systems Security (CrySyS)
Budapest University of Technology and Economics, Hungary
{buttyan, dora, vajda}@crysys.hu

² Laboratory for Computer Communications and Applications
EPFL — Switzerland, Lausanne
mark.felegyhazi@epfl.ch

Abstract

In this paper, we consider the application of delay-tolerant networks to personal wireless communications. In these networks, selfish nodes can exploit the services provided by other nodes by downloading messages that interest them, but refusing to store and distribute messages for the benefit of other nodes. We propose a mechanism to discourage selfish behavior based on the principles of barter. We develop a game-theoretic model in which we show that the proposed approach indeed stimulates cooperation of the nodes. In addition, the results show that the individually most beneficial behavior leads to the social optimum of the system.

1 Introduction

A delay-tolerant wireless network is a special type of wireless mobile ad-hoc network where the transfer of messages from their source to their destination is performed by the intermediate nodes in a *store-and-forward* manner. This means that the intermediate nodes carry the messages and pass them on to other intermediate nodes when they have a connection (e.g., when they are in vicinity).

In this paper, we consider the application of delay-tolerant networks for personal wireless communications. Such networks can complement traditional personal wireless communications systems, such as cellular networks, in applications where local information needs to be distributed to a set of nearby destinations based on their interest in the information.

As a motivating example, let us consider a touristic city,

such as Rome or Paris, where it would be beneficial for the tourists to be able to share information concerning the various touristic sights. A possible solution would be to set up an on-line bulletin board where tourists can post messages of potential interest for other tourists. However, this solution needs a service provider that runs the bulletin board service, and each tourist must have wireless Internet access for posting and downloading messages. The business model behind this solution would likely require the tourists to pay for both the service usage and the network access.

An alternative solution could benefit from the proliferation of Bluetooth capable personal devices such as mobile phones, PDAs, and MP3 players. These devices can communicate with each other when they are in vicinity even without any user intervention. Touristic information can then be distributed in a store-and-forward manner by using these devices and by exploiting the mobility of the tourists themselves. This would result in a city-wide delay-tolerant wireless network of personal devices.

A potential problem in delay-tolerant personal wireless networks is that the quality of the service provided by the system heavily depends on the users' willingness to cooperate. In particular, the users may act selfishly meaning that they download messages from other users that are interesting for them, but they deny storing and distributing messages for the benefit of other users. The motivation for such a selfish behavior is that personal devices are usually battery powered and have limited resources in terms of CPU and memory; hence, the users' interest is to save battery and other resources as much as possible. As shown in [1] if the majority of the users behave selfishly, then the message delivery rate decreases considerably, and damages the quality of the service provided by the network.

In this paper, we address the problem introduced above. Our main contributions are the following: (1) we propose a mechanism for stimulating cooperation in delay-tolerant personal wireless networks based on the principles of *barter*; (2) we develop a game-theoretic model in which the proposed mechanism can be studied; and (3) we consider the efficiency of the social optimum with respect to the Nash equilibria (i.e., the price of anarchy [7]). To the best of our knowledge, we are the first to propose and analyze a mechanism that stimulates cooperation in the context of delay-tolerant networks.

2 State-of-the-art

So far, the problem of selfish nodes has been addressed mainly in the context of mobile ad-hoc networks. The proposed solutions to stimulate cooperation can be broadly classified into two categories: reputation systems (see e.g., [8, 2]), and virtual payment based methods (see e.g., [3, 10]). Usually, these solutions require authentication (and related key management), and/or the presence of a trusted third party. In addition, the payment based solution also raises the problem of determining the price of different actions (see e.g., [4]).

Researchers have also studied under what conditions cooperation can emerge spontaneously among the nodes in ad-hoc networks (see e.g., [9, 5]).

The application of delay-tolerant networks for personal wireless communications is also considered in [6]. In particular, the authors show, by analytical tools and by means of simulations, that delay-tolerant networks can achieve a reasonably high throughput such that they can support various personal communication services.

To the best of our knowledge, [1] is the only paper so far that raises the problem of selfishness in delay-tolerant networks. The authors study the performance of three representative routing algorithms in the presence of some selfish nodes. They show that when the nodes behave selfishly, the performance decreases, in the sense that messages are delivered with a longer delay if they are delivered at all. However, the authors do not propose any mechanism to stimulate cooperation. The results presented in [1], can be viewed as a motivation for our work.

3 System description

3.1 The barter-based approach

Our approach to stimulate the cooperation of nodes in a delay-tolerant personal wireless network falls neither in the class of reputation based schemes, nor in the class of rewarding schemes. Instead, it is based on the principles of

barter. More specifically, we require that when two nearby nodes establish a connection, they first send the description of the messages that they currently store to each other, and then they agree on which subset of the messages they want to download from each other. In order to ensure fairness, the selected subsets must have the same size, and the messages are exchanged in a message-by-message manner, in preference order. If any party cheats, the exchange can be disrupted, and the honest party does not suffer any major disadvantage (i.e., the number of messages downloaded by the honest party is at most one less than the number of messages downloaded by the misbehaving party).

Note that it is entirely up to the nodes to decide which messages they want to download from each other. They may behave selfishly by downloading only those messages that are of primary interest for them. However, selfish behavior may not be beneficial in the long run. In particular, the idea is that a message that is not interesting for a node A may be interesting for another node B , and A may use it in the exchange protocol described above to obtain a message from B that is indeed interesting for A . In other words, the messages that are not directly interesting for a node still represent a *barter value* for the node, and hence, it may be worth downloading and carrying them. The purpose of our analysis later in this paper is to verify whether this statement holds or not.

3.2 System model

Our system model relies on the following assumptions:

- If two nodes establish a connection, then the lifetime of this connection is sufficiently long such that the two nodes can fully execute their message exchange.
- Every message has approximately the same size, and therefore, the cost of downloading a message over the wireless link is the same for every message.
- The communication cost of a message is much higher than its storage cost. Consequently, we assume that storage has no cost, and storage space is not limited in the nodes. We intend to relax this assumption in our future work.
- Messages lose their value over time. This is true for the primary value of a message as well as for its barter value. The primary value of a message decreases in time, because in the kind of applications for which delay-tolerant personal networks are envisioned, this is usually the case (e.g., touristic information, news, advertisements, announces, blogs, etc.). A message loses its barter value in time, because as it is disseminated in the network, the demand for that message decreases which pushes its value down.

In our model, the mobile nodes carry and exchange messages. These messages are generated by special nodes that are assumed not to be selfish, therefore the mobile nodes can get messages without requiring to have one in return. The message generating nodes are static and they generate new messages with a fixed average rate ρ . Each message generating node stores only the most recently generated message, which can be downloaded at the cost of communication by any mobile node that passes by the message generating node.

Each message has a type for each node. For simplicity, we distinguish only two types: primary messages and secondary messages. A message is a primary message for a given node, if the node is interested in the content of the message. In addition, primary messages also have some barter value. In contrast to this, the content of a secondary message is not directly interesting for the given node, it has only barter value. The nodes may decide to download, carry, and distribute secondary messages to exchange them for primary messages later. Note that a message may have different types for different nodes, as different nodes are interested in different contents.

Each node is characterized by its *scope of interest* $0 < p \leq 1$ that represents the likelihood that a randomly selected message in the network is a primary message for that node. A small p means that the node is interested only in a small fraction of messages, whereas a large p means that the node is interested in a large fraction of the messages.

Each message has some value for each node. The value of a message is determined by its type and its age. For simplicity, we assume that primary messages of the same age have the same value for the node. Similarly, secondary messages of the same age have the same value for the node. Without loss of generality, we assume that the value of a primary message at the time of its generation is one unit, and this is discounted in time such that the value of the same message after t time units is only δ^t , where $0 < \delta < 1$ is the discounting factor, which is the same for all nodes. The value of a secondary message at the time of its generation depends on how the node values secondary messages with respect to primary messages, and it is discounted in the same way as primary messages. In other words, if for a node, secondary messages are worth s units for some $0 \leq s \leq 1$ at the time of their generation, then the value of a secondary message after t time units is $s \cdot \delta^t$. Note that in general, the value of a secondary message cannot be larger than the value of a primary message of the same age (i.e., $s \leq 1$), because the primary message has the same barter value as the secondary message, and in addition, the node is interested in its content.

3.3 Barter-based message exchange

As we described above each node u decides which messages it wants to download from another node v in vicinity, and in what order. The node's behavior depends on two parameters:

- the ratio s_u between the value of a secondary message and the value of a primary message (of the same age) for the node, and
- a threshold value h_u , below which the node does not download secondary messages from other players.

We call the first parameter *secondary/primary ratio*, and the second parameter *secondary value threshold*.

When two nodes get in the vicinity of each other, they interact in the following way:

1. The nodes exchange the list of the messages that they carry. The exchanged lists contain only the short descriptions of the messages (including their time of generation) rather than the messages themselves.
2. Each node u removes from the list $L_v^{(0)}$ received from v the messages that u already stores in memory, and thereby obtains the list $L_v^{(1)}$.
3. Each node u determines the value of the messages listed in $L_v^{(1)}$ based on their types, their ages, and the secondary/primary ratio s_u as described above. Then, u removes those secondary messages from $L_v^{(1)}$ whose value is below the secondary value threshold h_u , and it also removes those primary messages from $L_v^{(1)}$ whose value is below a small constant c representing the communication cost of downloading a message. The list obtained in this way is denoted by $L_v^{(2)}$.
4. Each node u orders the messages contained in $L_v^{(2)}$ by their value in descending order. The resulting ordered list $L_v^{(3)}$ is the list of messages that u wishes to download from v .
5. The nodes exchange the first $\ell = \min(|L_u^{(3)}|, |L_v^{(3)}|)$ messages from the beginning of their lists on a message-by-message manner, where $|L|$ denotes the length of the list L . Thus, the number of exchanged messages is determined by the length of the shorter list, which is in accordance with the barter principle described in Section 3.1.

4 Game model

As described above, in our proposed barter-based scheme, when two nodes establish a connection, they decide which messages they download from each other. The

behavior of the nodes is determined by two parameters: the secondary/primary ratio, and the secondary value threshold. The nodes make their choice in a selfish manner, to maximize their own benefit. Therefore, it is convenient to model the unfolding of the system in a game-theoretic framework. In this section, we describe the elements of this framework: the players, the strategy space, and the payoffs.

A natural approach would be to model each node as an individual player, however, we refrain from doing that in order to control the complexity of the analysis. Instead, we make the reasonable assumption that in practice, the majority of the nodes follow some pre-programmed protocol (whatever it is), and there may be a small minority of nodes that deviate from this protocol. Therefore, we define a two-player game, where each player is a group of nodes. The first player is called the *crowd*, and it represents the majority of the nodes. The second player is called the *deviators*, and it represents the small group of nodes that deviate from the default program.

For simplicity, we assume that all nodes that belong to the same group behave in the same way (i.e., they choose the same secondary/primary ratio and secondary value threshold). However, the scope of interest of the nodes that belong to the same group may be different.

The strategy of each player can be represented by a pair of real numbers $(s, h) \in [0, 1] \times [0, 1]$. The first number s is the secondary/primary ratio of the nodes represented by the player, and the second number h is their secondary value threshold.

Note that when $h \geq s$, the nodes do not download any secondary messages. Therefore, all strategies where $h \geq s$ are equivalent. We will represent this equivalence class with a single strategy where $h = 1$ and $s = 0$. Besides this strategy, we only consider strategies where $h < s$.

The nodes receive a *score* after each interaction and they accumulate these scores to obtain their total score at the end. The payoffs obtained by the players in the game are defined as the average total score of the nodes in the respective groups.

The score r_u received by node u after an interaction is composed of two parts: a gain and a loss. The gain is determined by the total value of the primary messages downloaded in the interaction and the scope of interest of the node. The loss is determined by the total number of exchanged messages in the interaction. The formula of the score computation is the following:

$$r_u = \left(\sum_i \frac{\delta^{t_i}}{p_u} \right) - \ell \cdot c \quad (1)$$

where δ is the system wide discounting parameter, t_i is the age of the i -th primary message downloaded in the interaction, p_u is the scope of interest of node u , ℓ is the number

of messages exchanged in the interaction, and c is the cost of a single message exchange. Note that the values of the primary messages are weighted with $\frac{1}{p_u}$ in the computation of the gain. This means that the relative value of a primary message is higher for those nodes whose scope is smaller (at least when determining the gain).

5 Simulations

We want to determine if our barter-based scheme stimulates the nodes to cooperate or not. For this purpose, we analyze the game defined in the previous section by means of simulations. We search for the Nash equilibria and study how far these Nash equilibria are from the socially optimal behavior of the players, where the social optimum is reached when the total payoff of the players is maximal.

Simulation settings. In our simulations, the crowds consists of 90% of the mobile nodes; the remaining 10% constitutes the deviators. The only parameter that characterizes the nodes is the scope of interest. This parameter is set as system parameter in each group (player) in the following way: 10% of the nodes have scope of interest $p = 0.01$, 80% of the nodes have $p = 0.1$, and the remaining 10% of the nodes have $p = 0.5$. This means that only a few nodes are interested in a very small or a very large number of messages.

In order to make the simulation feasible, we quantize the strategy space such that we restrict the possible values for s and h of each player to the set $\{0, 0.25, 0.5, 0.75, 1\}$. In addition, recall that we consider only strategies where $h < s$ and the single strategy where $h = 1$ and $s = 0$ (representing all strategies where $h \geq s$). The resulting strategy space is illustrated in Figure 1.

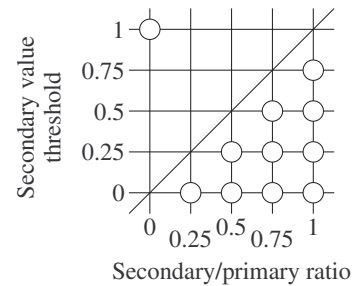


Figure 1. Quantized strategy space

In our simulations, the fixed-number nodes move in discrete time steps according to one of the two mobility models: the random walk and the grid-based random waypoint models.

In the random walk model, the nodes move on a grid of size 15×15 . In each time step, a node can move to one

of the four neighboring grid point, or stay where it is. The probability of each of these actions is 0.2. The nodes that happen to be at the same grid point in the same time step execute the message exchange protocol in such a way that each node interacts with each other node in a random order.

In the grid-based random waypoint model, the nodes move on a grid of size 20×20 . A subset of the grid points is chosen at random; these are called meeting points. Each node selects a meeting point randomly, and moves towards this meeting point with a fixed speed on the grid. When the meeting point is reached, the node stops and stays for randomly chosen time. Then it chooses another meeting point and begins to move again. The nodes that happen to be at the same meeting point in the same time step execute the message exchange protocol similarly to the random walk model.

Recall that in our system model, the messages are injected into the network by special message generating nodes that are static. In the random walk model, the message generating nodes are placed at a randomly selected subset of the grid points. In the grid-based random waypoint model, the message generating nodes reside in the meeting points.

The values of the parameters of the simulations and those of the mobility models are summarized in Tables 1 and 2.

Table 1. Parameter values for the simulations

Parameter	Value
Number of mobile nodes	300
Number of msg generating nodes	100
Message generation rate ρ	0.01
Discount parameter δ	0.995
Communication cost c	0.1
Simulation length (time steps)	1000

Table 2. Parameter values of the mobility models

Random walk	
Simulation area	15×15 grid
Prob. of staying	0.2
Prob. of each direction	0.2
Grid-based random waypoint	
Simulation area	20×20 grid
Velocity (grid/time step)	1
Number of meeting points	100
Probability of leaving a meeting point	0.1

Simulation results. Figure 2(a) and 2(b) show the payoffs received by the crowds and by the deviators, respectively, in the case of the random walk model. Figure 3(a)

and 3(b) show the results in the case of the grid-based random waypoint model. Note that for presentation purposes, we illustrate the payoff values with colors (gray scale) instead of numbers. The semantics is that the darker a cell is the higher the payoff achieved. The social optima and the Nash equilibria are marked with the symbol “|” and “-”, respectively (together denoted by “+”).

As one can see, in each case, there is a single Nash equilibrium, which consists of the strategy pair where the secondary/primary ratio is 1 and the secondary value threshold is 0.75 for both players. This means that in the Nash equilibrium, both players value secondary messages in the same way as primary messages, and as a consequence, both players download and carry secondary messages. Thus, our barter-based approach indeed stimulates cooperation. In addition, the relatively high value of the secondary value threshold means that the players download only fresh messages.

Moreover, as we can see, the Nash equilibria coincide with the social optima. This means that the behavior which is individually the most beneficial results in a socially optimal behavior. This is a strong result, which essentially means that the *price of anarchy* in our barter-based system is 1. Moreover this result is achieved without any overhead of introducing specific reputation-based or micropayment-based schemes.

Finally, the fact that in the Nash equilibrium both players play the same strategy means that it is safe to program the nodes with this strategy, and the nodes *truthfully* follow the pre-defined protocol.

6 Conclusion

In this paper, we addressed the problem of selfishness in delay-tolerant networks used for personal wireless communications. We proposed a mechanism to discourage selfishness based on the principles of *barter*. We developed a game-theoretic model, and analyzed our barter-based approach in this model. Our simulation results show that the proposed approach indeed stimulates cooperation of the nodes. In addition, the results show that the individually most beneficial behavior leads to the social optimum of the system.

Acknowledgements

The work presented in this paper has been partially supported by the BIONETS Project (www.bionets.org) and the Hungarian Scientific Research Funds (OTKA T046664). The authors are thankful to Eleonora Borgia, Julien Freudiger, and Michal Piorkowski for their useful comments on an early version of this paper.

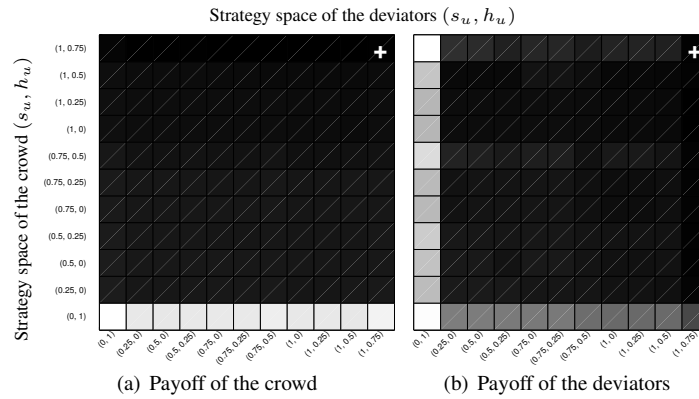


Figure 2. Simulation results in the random walk model (darker means higher payoff)

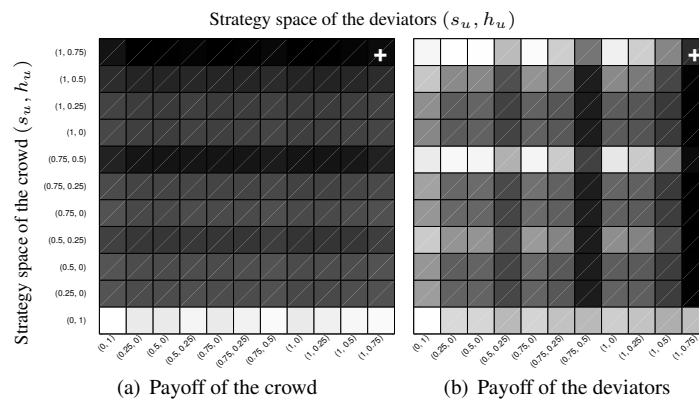


Figure 3. Simulation results in the grid-based random waypoint model (darker means higher payoff)

References

- [1] A. V. Antonis Panagakis and I. Stavrakakis. On the effects of cooperation in DTNs. In *Proc. of The Second IEEE/Create-Net/ICST International Conference on COMmunication System softWare and MiddlewaRE (COMSWARE)*, January 7-12 2007.
- [2] S. Buchegger and J.-Y. L. Boudec. Performance analysis of the CONFIDANT protocol (cooperation of nodes-fairness in dynamic ad-hoc networks). In *Proc. of the 3rd ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc'02)*, pages 80-91, June 9-11 2002.
- [3] L. Buttyán and J.-P. Hubaux. Stimulating cooperation in self-organizing mobile ad hoc networks. *ACM/Kluwer Mobile Networks and Applications (MONET) Special Issue on Mobile Ad Hoc Networks*, 8(5), October 2003.
- [4] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring. Modelling incentives for collaboration in mobile ad hoc networks. In *Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt'03)*, March 3-5 2003.
- [5] M. Félegyházi, J.-P. Hubaux, and L. Buttyán. Nash equilibria of packet forwarding strategies in wireless ad hoc networks. *IEEE Transactions on Mobile Computing*, 5(5), May, 2006.
- [6] G. Karlsson, V. Lenders, and M. May. Delay-tolerant broadcasting. In *CHANTS '06: Proceedings of the 2006 SIGCOMM workshop on Challenged networks*, pages 197-204, New York, NY, USA, 2006. ACM Press.
- [7] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS'99)*, March 1999.
- [8] P. Michiardi and R. Molva. Core: A Collaborative REputation mechanism to enforce node cooperation in mobile ad hoc networks. In *Proc. of Communication and Multimedia Security 2002*, Sep. 26-27 2002.
- [9] V. Srinivasan, P. Nuggehalli, C. F. Chiasserini, and R. R. Rao. Cooperation in wireless ad hoc networks. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM '03)*, March 30 - April 3 2003.
- [10] S. Zhong, Y. R. Yang, and J. Chen. Sprite: A simple, cheat-proof, credit-based system for mobile ad hoc networks. In *Proceedings of the IEEE Conference on Computer Communications (INFOCOM '03)*, March 30 - April 3 2003.