

Impact of Altruism on Opportunistic Communications

Kuang Xu*, Pan Hui[‡], Victor O.K. Li*, Jon Crowcroft[†], Vito Latora[§], Pietro Lio[†]

*The University of Hong Kong [†]University of Cambridge

[‡]Deutsche Telekom Laboratories/TU-Berlin [§]Università di Catania and INFN

Abstract—Opportunistic networking largely relies on human-carried devices as relays to move data across the network. Human altruistic behavior is an important factor in the feasibility and performance of such a system. In this paper, we study the impact of different altruism distributions on the system throughput and delay of opportunistic communications. We evaluate the system performance by utilizing social network topological models and experimental human mobility traces with different communication patterns. We also discuss several observations from the result of the study. To the best of our knowledge, this is the first study of the impact of altruism on opportunistic communications.

Index Terms—Opportunistic communication, altruism, social network, community, simulation.

I. INTRODUCTION

Humans carry an increasing number of small wireless devices capable of generating, storing, and forwarding information. As we move through our environments, temporary opportunities arise for short range connections amongst these devices, without relying on a network infrastructure. Hence opportunistic communication largely relies on human-carried devices to relay messages for others. However, in wireless environments, especially for small hand-held devices, resources such as battery or memory is still a major concern, and relaying messages for others requires altruistic behavior. Altruism refers to behavior that is “costly to the organism performing the behavior but benefits other organisms” [20]. Since humans constitute the network, human altruistic behavior inevitably impacts opportunistic communication system. It is an important metric to assess the feasibility and performance of opportunistic communication. While in the literature a host of studies focus on exploiting social context [5] [11] [12] or mobility patterns [3] [15] [23] in opportunistic communication, we find the study of this topic conspicuously missing.

In this paper, we intend to give an elementary but systematic study of the impact of altruism on the communication network, particularly on opportunistic networks. We expect our result to be applicable to other social-related networks. In particular, we look at how robust an opportunistic network is with different distributions of altruism in the population, and how the different altruism distributions affect the system performance. Since human altruism is closely related to kinship or socialization (e.g. friends, colleagues, and acquaintances) [20], we generate social network topologies using several popular social network models [13] [19] [25], and simulate asynchronous messaging in the network. To also account for human mobility, we further

evaluate the system throughput using two experimental human mobility traces gathered in two research projects [1] [6]. We find that both information dissemination in social networks and data forwarding in human mobile networks are quite robust towards altruism distributions, with better performance in those altruism distributions biased to the class of nodes that are more responsible for maintaining their underlying networks.

For the communication model, we utilize different traffic patterns. We assume each node in the network can generate messages for others but with different probabilities. For an ecological community, correlated interaction means that “an organism of a given type might be more likely to interact with another organism of the same type than with a randomly chosen member of the population” [20]. This correlated interaction is also applicable to human communication patterns, and people in the same community usually talk to each other more often than to people in other communities [18]. We believe these community-biased communication patterns also limit the impact of altruism in the system, and we find that the community-biased communication could further increase the system performance.

We proceed in this paper as follows. Related work is described in Section II, methodology in Section III, altruism model and communication pattern in Section IV, results discussion in Section V, and conclusions in Section VI.

II. RELATED WORK

Altruistic behavior has been studied in the literature using the game theory approach. For example [14] and [16] model players’ payoff as their own altruism coefficients and also what they believe their opponents’ coefficients to be. In this paper we are interested in the system throughput and delay when the system reaches steady state. Therefore, we assume static distributions of the altruism for the study and will investigate the more transient gaming strategies in the future. [4] studies the price of anarchy of traffic routing under the assumption that users are partially altruistic. They show that if all users have a coefficient of altruism $\beta > 0$, then the price of anarchy is bounded by $1/\beta$. The altruism model in this paper follows an incentive (utility) approach as introduced by [14]. They assume the utility of each player is a linear combination of his own *a priori* payoff and the payoffs of other players. We do not focus on how a particular distribution of altruism is induced by a certain kind of incentive of the

nodes, but study the result of the system with the already formed altruism distribution. Altruism has also been studied in cooperative services. The authors in [2] present a general approach to construct cooperative services, which tolerate *Byzantine, Altruistic, and Rational* behavior. Altruistic nodes in the system are those that follow the suggested protocol exactly, and hence they are not a challenge in the protocol design. In P2P research, Piatek *et al.* have explored the altruistic behaviors in BitTorrent by measurement [22]. They found that high capacity peers send much faster than the minimal rate required for reciprocation, which is similar to one of our observations from the geometric altruism distribution that nodes with higher degrees are of more importance in maintaining the system performance. This paper focuses on evaluating the performance of information dissemination in opportunistic networks with nodes of different altruistic states, and we find this type of study conspicuously missing in the literature.

III. METHODOLOGY

We use two practical approaches to study the impact of altruism on opportunistic communication, based on social network topological models and human mobility traces.

A. Social Network Models

We first study information dissemination on three different topologies popularly used in modeling human social networks, referred to as Simple model [25], Newman model [19], and Kumpula model [13]. We consider this approach because, as is well known in sociology, altruism is closely dependent on kinship and social relationship. Whether people would share their resource to forward messages for others highly depends on the social links between the requesting and requested nodes. We use asynchronous messaging on the above generated topologies, which means the delivery of different messages in the system are independent. Besides, we set the Time-to-Live (TTL) value of all messages to the diameter of their underlying network topology¹. The diameter of a network is the longest shortest-path between any two nodes in the network. In our system, each node is assigned an altruism value according to the distributions we will introduce in Section IV. During the simulation, each source node broadcasts its message to its neighbors only once. If a node receives a message with live TTL, and it has not been forwarded by this node before, the node makes a decision of whether to broadcast this message according to its altruistic state, and updates the TTL accordingly. In broadcasting, a message may arrive at the same node multiple times with different TTLs. To focus on the effect of altruism on the networks, and isolate the influence of the message forwarding strategy, we simulate the “best effort” strategy, such that a message’s replica with greater TTL will arrive at a node first.

Table I summarizes the characteristics of the three models². In order for the topologies to be comparable with each

¹This guarantees that when all the nodes in the network are totally altruistic, the system throughput is 1, which serves as the upper bound of the system.

²For Kumpula model, a community detection algorithm [21] is applied to split the network into communities by using a 7-clique percolation.

Network model	Simple	Newman	Kumpula
Number of nodes	1000	996	1000
Number of edges	5001	5024	4954
Diameter	6	6	9
Clustering coefficient	0.264	0.092	0.473
Number of communities	50	50	43

Table I
CHARACTERISTICS OF THE THREE SOCIAL TOPOLOGIES

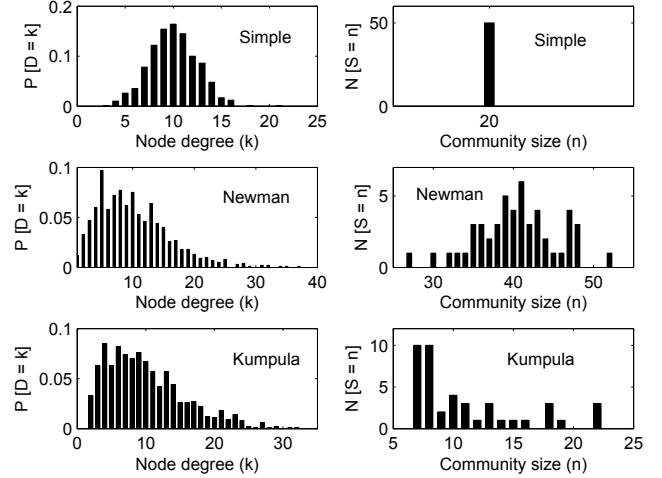


Figure 1. Degree and community distribution of the social topologies other, we keep the number of nodes approximately equal to 1,000 and the average number of neighbors of each node approximately equal to 1% of the total population. Thus, the difference in their structures results solely from structuring of links according to different social scenarios. Fig. 1 shows the degree and community distributions of all the topologies.

B. Human Mobility Traces

The first approach studies the effect of altruism on the spreading of messages on social networks, and gives an estimated performance, but it does not consider the mobility of nodes and cannot fully reflect the performance of a mobile social network. Then, we utilize human mobility traces to further validate our observations. We utilize two experimental datasets gathered in the Reality Mining Project [1] and the Huggle Project [6]. In these experiments, Bluetooth-enabled mobile devices ran software logging contacts with each other by doing Bluetooth device discovery periodically. Table II summarizes these experiments. The characteristics of these datasets, such as inter-contact and contact distributions, have been explored in several studies [8] [12].

We develop a event/contact driven simulator which can replay the collected mobility traces. To evaluate the forwarding performance under different distributions of altruism, we use a complete opportunistic flooding strategy that messages are duplicated and sent to all mobile nodes encountered, as in

Experimental dataset	Reality	Cambridge
Device	Phone	iMote
Duration (days)	246	11
Granularity (seconds)	300	600
No. of experimental devices	97	54
No. of internal contacts	54,667	10,873
Average no. of contacts/pair/day	0.024	0.345

Table II
CHARACTERISTICS OF THE TWO EXPERIMENTAL DATASETS

epidemic routing [24]. As in Section III-A, a node makes a decision about whether to relay a message for its neighbor depending solely on its altruistic state. In the simulation, we assume that each node asks another node it encounters for message forwarding. In addition, a node does not record whether a particular message has already been requested for forwarding before, unless it has relayed the message. This is similar to a simple memory-less Markov process. It makes sense in real life as a mobile device does not have to waste resource to store information about a message that it refused to relay.

IV. SYSTEM ASSUMPTIONS

In this paper, we utilize several different altruism models and communication patterns.

A. Altruism Models

Altruism is observed in many aspects of the modern societies and also exhibited by hunter-gatherers, who typically have dense networks of exchange relations, food-sharing, cooperative hunting, and collective warfare. Human altruism has been intensively studied [9] [10] [20]. In general, each node can have different altruism states for different other nodes. Here we assume several altruism distributions for the whole population. All the altruism values are distributed between 0 and 1, where 0 stands for totally selfish and 1 stands for totally altruistic. (In [4] [16], altruism values are also modeled as between -1 and 1 , where -1 stands for totally spiteful, but we do not consider spiteful behavior in this study.)

- *Degree-biased Distribution*, relates the altruism to one characteristic of the network topology, namely, node degree (in social science, it is also referred to as degree centrality), as follows:

$$a_i = \left(\frac{k_i - k_{min}}{k_{max} - k_{min}} \right)^\alpha \quad \text{with } \alpha > 0 \quad (1)$$

where k_{min} and k_{max} are respectively the smallest and the largest degrees in the network. With this formula, independently of the value of α , we have $a_i = 0$ for $k_i = k_{min}$, and $a_i = 1$ for $k_i = k_{max}$, that is the node with lowest degree always has $a = 0$, while the hub has $a = 1$. When $\alpha = 1$ the value of a grows linearly with the degree, while $\alpha > 1$ ($0 < \alpha < 1$) indicates a superlinear (sublinear) dependence of a from k . Notice that, when $\alpha = 0$ all nodes have the same altruism $a = 1$. The scenario for this model is that in the social network, people become popular and have many friends since they are more willing to help others. If we want to have $a_i = 1$ for $k_i = k_{min}$, and $a_i = 0$ for $k_i = k_{max}$, that is the node with the lowest degree always has $a = 1$, while the hub has $a = 0$ we need to use instead:

$$a_i = \left(\frac{k_{max} - k_i}{k_{max} - k_{min}} \right)^\alpha \quad \text{with } \alpha > 0 \quad (2)$$

The scenario for this is that if a person has too many friends, he may not have enough resource to help all of them, while a person with only one friend will probably be very glad to help out this friend.

- *Geometric Distribution*, relates the altruism to the distance of social relationship between people. As the simplest case, the altruistic values are calculated per pair of nodes, and each follows a distribution such that the probability decreases with the social hops, k , as in $P(X = k) = (1 - p)^{(1-k)} \cdot p$, where p can be considered as the altruism of a person towards his closest acquaintances (i.e. the first social hop). However, increasing p indicates the increase in the decaying rate of altruism of nodes farther away from it. The assumption is that all the acquaintances of a person share a fixed amount of resource from this person. In order to guarantee that the maximum altruism value is 1, we normalize the altruism values by the parameter p . In real life, altruism also decreases from kinship (first hop) to those people who are farther away on the social graph [17].
- *Uniform and Normal Distribution* dictates that the altruistic value of the whole population is uniformly distributed or follows normal³ distribution between 0 and 1. Uniform and normal distributions are popularly encountered in nature and society, and can be possibilities for altruism distribution.
- *Community-biased Distribution* accounts for the heterogeneity of altruism towards different people. It assumes that people in a community have greater incentives to carry messages for the other members in the same community and less incentives to carry messages for people outside the community. Hence, we model the altruism of each node using a tuple (a, e) , representing a node with probability a to carry data for intra-community nodes and e for inter-community nodes. The community-biased altruism has significant effect on the system performance of opportunistic communications. But due to space limitations, in this paper, we do not present the result in Section V, and leave the detailed study as future work.

B. Communication Patterns

In this paper we assume asynchronous communication. In principle, each node in the network can create messages for any other nodes, but we study the effect of different communication patterns. At the beginning of the simulation, each node picks a list of destination nodes to send messages based on certain communication pattern. We analyze uniform and community-biased traffic from a source for the communication scenarios.

- *Uniform pattern*: The source-destination pairs are uniformly distributed throughout the whole population. In this case, each node has the same probability to communicate with all other nodes, and there is no bias for the traffic.
- *Community-biased pattern* dictates that each node tends to communicate with more nodes inside its community than those in other communities [18]. This traffic pattern is determined by two parameters, P_{in} and P_{out} , which

³We adopted 5% and 95% of the Cumulative Distribution Function (CDF) value as two cutoffs in the negative and the positive directions of the normal distribution function.

are the probability of generating messages for community members and that for inter-community communication. Note that $P_{in} + P_{out} = 1$.

Uniform traffic assumes an equal communication opportunity for each node pair, but in reality a person normally talks to different people with different frequencies. Community-biased traffic is more realistic in that it assumes heterogeneity in communication patterns.

V. RESULT AND EVALUATION

The evaluations in this paper mainly focus on: 1) the robustness of opportunistic communications under different altruism distributions, 2) the comparison of effects of different altruism models on the overall system performance, and 3) the effect of communication patterns on the system with different altruism models. We summarize the most general observations at the end of this section (Section V-D).

A. Performance Metrics

We utilize several metrics for the above analyses:

- *System throughput*, or delivery ratio, which is the proportion of the number of delivered messages out of the total number of created messages in the system.
- *System delay*, or delivery delay, which only takes into consideration the delivered messages. For the study based on social topological networks, it is the expected number of hops that delivered messages traverse divided by the TTL (set to the diameter of the network). For the study based on mobility traces, messages in the system are given a finite life-time, after which they are no longer considered of interest. By varying this life-time value, the delivery ratio also reflects the delay of the system.
- *System cost*, or delivery cost, is the expected number of message replicas the system takes to deliver a message. It is an important metric for our study since forwarding duplicated messages consumes resources, such as battery and memory. But for brevity of this paper, we leave its investigation as future work.

The results in this paper are obtained as averages of 20 simulations. We also compute the 95th percentile with Student's t-distribution, but we do not plot it for the results based on social topological networks since the confidence intervals are quite trivial.

B. Social topological Networks

Fig. 2 shows the system delivery ratio with the parameter p for geometric altruism distribution when the communication pattern is uniform. One interesting observation is that Simple is less tolerant to p than Newman or Kumpula. This is because there is no high-degree node (with degree above 20) in Simple as in the other two networks (see Fig. 1), and nodes with higher degree are more favorable to spread messages under the broadcasting case. We observe that for altruism of geometric distribution, high-degree nodes are more important to maintain the system performance. We will use this observation for later analyses. In addition, for the robustness, we can see that if the




Figure 2. Delivery ratio with p using uniform communication pattern




Figure 3. Delivery ratio with α using uniform communication pattern

traditional six degree of separation holds in the real world, the system can still provide over 85% throughput even if each node shares most of its resource (60%) with its neighbors and only a small portion (40%) with the rest (on average 99%) of the population (see Fig. 2 for the Newman model).

Fig. 3 shows the system delivery ratio with the parameter α for degree-biased altruism distribution when the communication pattern is uniform. We can see that the system performs satisfactorily when $\alpha < 1$. The plot also clearly conveys that for all three social topologies, the system delivery ratios under lower-degree biased altruism distribution are greater than those under high-degree biased scenario. From Fig. 1, we see that the social topologies we consider comprise more low-degree nodes than high-degree ones. We conclude that the system performs better when the altruism distribution favors the class of nodes that are more responsible for maintaining the network. The result in the next section will confirm this conclusion.

We have shown in Fig. 2 that high-degree nodes are more important for maintaining the system performance when the altruism distribution is geometric. Then we compare the system performance under geometric and high-degree biased altruism distribution in Fig. 4 and Fig. 5, when the communication pattern is community-biased. The x-axis represents the percentage of traffic targeting the same community with the source. In order for the two altruism distributions to be comparable with each other, we choose the values of p and α




Figure 4. Delivery ratio with community biased communication pattern




Figure 5. Delivery delay with community biased communication pattern

where the delivery ratios are same in Fig. 2 and Fig. 3⁴. We can see that for all three underlying networks, the system has greater throughput with the altruism of geometric distribution than with high-degree biased scenario, and the former case also incurs lower system delay than the latter. From the observation we conclude that for the altruism distributions that favor the class of nodes with higher degrees, the system prefers the altruism distribution based on social distance rather than that based on social degree centrality. The plots also show that the community-biased communication pattern can further increase the system performance, and the system throughput and delay increases and decreases, respectively, with the increase in percentage of intra-community traffic.

Fig. 6 presents the system delivery ratio for altruism of Uniform and Normal distributions when the communication pattern is community-biased. It clearly shows that Uniform and Normal altruism distribution have quite similar (or the same) impact on the system performance. Again, we can observe that the system is quite robust towards any pattern of communication (with system delivery ratios above 90%).

C. Human Mobility Traces

In this section, we look at whether the major results from the above simulations can also be observed on the real human mobile network. Since not all the datasets we can access have

⁴We choose $p = 0.4$ and $\alpha = 1$ for Simple, $p = 0.7$ and $\alpha = 1.6$ for Newman, and $p = 0.5$ and $\alpha = 0.9$ for Kumpula.




Figure 6. Delivery ratio with uniform and normal altruism distribution

a priori community information, which changes dynamically by nature [26], we leave the study of community-biased communication as future work.

Fig. 7 and Fig. 8 show the system throughput on the Reality and Cambridge datasets, respectively. Each graph shows the throughput under totally altruistic, normal, uniform, and degree-biased altruism distributions. The x-axis shows the TTL⁵ of each message in natural logarithmic scale, which is the time each message can stay in the system, and hence also indirectly reflects the delays of the deliveries. From both figures, we can observe that normal and uniform altruism distributions yield quite similar delivery ratios (their confidence intervals are almost overlapping), which coincide with the observation in Section V-B. One interesting observation is that the system yields greater delivery ratio under altruism of high-degree biased distribution than that under high-degree biased scenario. This seems to contradict the observation in Section V-B that the low-degree biased altruism distribution favors the system performance more. But the study in [12] shows that the node betweenness⁶ are highly correlated with the node degree in these mobility traces, which indicates that for these datasets, the nodes with high degree are more responsible for maintaining the network connectivity (since they move more actively than other nodes). This is in agreement with the generalized observation in Section V-B that the system performs better when the altruism distribution favors the class of nodes that are more responsible for maintaining the network. In addition we can see that the system throughput under these altruism distributions are very close to that of the “Totally altruistic” case, especially for the Cambridge datasets⁷. And even for this worst case, the system throughput is still no less than half of that under “Totally altruistic” for TTL with small values (68% for the Reality dataset and 87% for the Cambridge dataset). These results further confirm the robustness of opportunistic communications towards distributions of altruism.

⁵Here, the TTL is different from that in Section V-B, which is the number of forwarding hops a message can last.

⁶Betweenness [7] of a node is the total number of shortest paths between all possible pairs of nodes that pass through this node.

⁷The delivery ratio of the Reality dataset is generally low even with TTL up to one week. This is because many participants switched off their Bluetooth transceivers, which makes the network quite sparse.




Figure 7. Delivery ratio with uniform traffic on Reality data




Figure 8. Delivery ratio with uniform traffic on Cambridge data

D. Discussion

Finally we summarize the most general observations from the above studies both on social network topological models and experimental human mobility traces.

- Opportunistic communication is quite robust when the altruism of participating nodes is considered.
- Opportunistic communication prefers the altruism model biased to the class of nodes that are more responsible for maintaining its underlying network.
- The uniform and the normal distribution of altruism have almost the same impact on the system throughput of opportunistic communications.
- Community-biased communication pattern can further increase the performance of opportunistic communication system with different altruism models.

VI. CONCLUSION AND FUTURE WORK

While proposing that human altruism has impact on opportunistic communications, this paper presents a primary investigation on this topic. We elaborate several observations from the study. In this paper, we only consider the altruism distribution at the steady state, and it would be interesting to study the altruism resulting from gaming strategies. Currently, we assume each user does not keep its rejection history of message delivery for other nodes. In the future, we can study variations of this. For example, limiting the number of delivery requests for users to incur lower system cost. It would also be

interesting to study how much resource a node can save by being selfish or only helpful within the community.

ACKNOWLEDGMENT

This research is supported in part by the University of Hong Kong Strategic Research Theme on Information Technology, and the EC IST SOCIALNETS.

REFERENCES

- [1] Reality mining project. <http://reality.media.mit.edu/>.
- [2] A. S. Aiyer et al. Bar fault tolerance for cooperative services. *SIGOPS Oper. Syst. Rev.*, 39(5):45–58, 2005.
- [3] A. Chaintreau et al. Impact of human mobility on the design of opportunistic forwarding algorithms. In *Proc. INFOCOM*, April 2006.
- [4] P.-A. Chen and D. Kempe. Altruism, selfishness, and spite in traffic routing. In *Proc. of the 9th ACM conf. on electronic commerce*, 2008.
- [5] E. Daly and M. Haahr. Social network analysis for routing in disconnected delay-tolerant manets. In *Proceedings of ACM MobiHoc*, 2007.
- [6] C. Diot et al. Huggle project. <http://www.huggleproject.org>, 2004.
- [7] S. N. Dorogovtsev and J. F. F. Mendes. *Evolution of networks: from biological nets to the Internet and WWW*. Oxford University Press, 2003.
- [8] N. Eagle and A. Pentland. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, 10(4):255–268, 2006.
- [9] E. Fehr and U. Fischbacher. The nature of human altruism. *Nature*, 425(6960):785–791, October 2003.
- [10] E. Fehr and B. Rockenbach. Detrimental effects of sanctions on human altruism. *Microeconomics 0305007*, EconWPA, May 2003.
- [11] F. Guidic and Y. Maheo. Opportunistic content-based dissemination in disconnected mobile ad hoc networks. In *Proc. of UBIKOMM '07*, 2007.
- [12] P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: Social-based forwarding in delay tolerant networks. In *Proc. of MobiHoc*, 2008.
- [13] J. M. Kumpula, J. P. Onnela, J. Saramaki, K. Kaski, and J. Kertesz. Emergence of communities in weighted networks. 2007.
- [14] J. O. Ledyard. Public goods: A survey of experimental research. *Public Economics 9405003*, EconWPA, May 1994.
- [15] J. Leguay, T. Friedman, et al. Evaluating mobility pattern space routing for DTNs. In *Proc. INFOCOM*, 2006.
- [16] D. K. Levine. Modeling altruism and spitefulness in experiment. *Review of Economic Dynamics*, 1(3):593–622, July 1998.
- [17] E. A. Madsen et al. Kinship and altruism: A cross-cultural experimental study. *British Journal of Psychology*, 98:339–359, May 2007.
- [18] A. Mislove et al. Ostra: Leveraging trust to thwart unwanted communication. In *Proc. of NSDI '08*, 2008.
- [19] M. E. J. Newman and J. Park. Why social networks are different from other types of networks. *Phys. Rev. Lett. E*, 68:036122, May 2003.
- [20] S. Okasha. Altruism, group selection and correlated interaction. *British Journal for the Philosophy of Science*, 56(4):703–725, December 2005.
- [21] G. Palla et al. Uncovering the overlapping community structure of complex networks in nature and society. *Nature*, 435(7043):814–818, 2005.
- [22] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani. Do incentives build robustness in bittorrent? In *NSDI '07*, 2007.
- [23] N. Sarafjanovic-Djukic et al. Island hopping: Efficient mobility-assisted forwarding in partitioned networks. In *IEEE SECON*, 2006.
- [24] A. Vahdat and D. Becker. Epidemic routing for partially connected ad hoc networks. Technical Report CS-200006, Duke University, 2000.
- [25] D. J. Watts. *Small worlds: the dynamics of networks between order and randomness*. Princeton Studies on Complexity. Princeton University Press, 1999.
- [26] K. Xu, G.-H. Yang, V. O.-K. Li, and S.-Y. Chan. Detecting dynamic communities in opportunistic networks. In *Proc. of IEEE ICUFN '09*, 2009.