

Selfishness, Altruism and Message Spreading in Mobile Social Networks

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

* Deutsche Telekom Laboratories/TU-Berlin [†] University of Cambridge

[‡] University of Hong Kong [§] Università di Catania and INFN

Abstract—Many kinds of communication networks, in particular social and opportunistic networks, rely at least partly on humans to help move data across the network. Human altruistic behavior is an important factor determining the feasibility of such a system. In this paper, we study the impact of different distributions of altruism on the throughput and delay of mobile social communication system. We evaluate the system performance using four experimental human mobility traces with uniform and community-biased traffic patterns. We found that mobile social networks are very robust to the distributions of altruism due to the nature of multiple paths. We further confirm the results by simulations on two popular social network models. To the best of our knowledge, this is the first complete study of the impact of altruism on mobile social networks, including the impact of topologies and traffic patterns.

I. INTRODUCTION

We envision a future in which a multitude of devices carried by people are dynamically networked. We aim to build Pocket Switched Networks (PSN) [4] [11] for such environments. A PSN is a type of Delay Tolerant Networks (DTN) [8] which uses contact opportunities (hence the name “opportunistic network”) to allow humans to communicate without network infrastructure. PSNs rely on devices carried by humans to relay messages for others. In wireless environments, especially for small hand-held devices, battery life is a major concern and hence relaying messages for others may require altruistic behavior. Since humans constitute the network, human altruistic behavior impacts the throughput of the communication system. It is an important factor in the feasibility of a PSN, but we found that the study of this area was conspicuously missing in the literature to date.

In this paper, we aim to give a systematic study of the impact of altruism on communication, particularly on opportunistic networks, but we expect our results to be applicable to other social networks. In particular, we look at how robust an opportunistic network is under different distributions of altruism in the population. For the communication model, we use different traffic patterns. We assume each node on the network can create messages for any other, but with differing probabilities. For an ecological community, correlated interaction means that an organism of a given type is more likely to interact with another organism of the same type than with a randomly chosen member of the population [18]. This correlated interaction is also applicable to human communication patterns, and people in the same community may talk to each

other more often than to people in other communities. We believe these community-biased communication patterns also limit the impact of altruism (or selfishness) on the system, and hence we are interested in investigating it.

To account for human mobility, we evaluate the system throughput using four real traces gathered by two research projects. We find that both information dissemination on both these kind of mobile social networks is very robust towards different forms of altruism distribution, largely because of multiple-path nature. The system can have around 90 percent throughput of the original all-altruistic system, whether with uniform, normal, or degree-biased altruism distribution. We also find that community-biased traffic pattern can further increase the robustness of the system. To confirm the generality of this result, we established the same findings with two widely used social models that are available [21] [13].

The rest of the paper is organised as following: experimental datasets (Section II), altruism and traffic models (Section III), results and evaluations of dynamics networks (Section IV), evaluation using static social networks (Section V), related work (Section VI), and conclusions (Section VII).

II. MOBILE SOCIAL EXPERIMENTS

In this paper, we use three experimental datasets gathered by the Huggle Project ¹ over two years, referred to as *Cambridge*, *Infocom05*, and *Infocom06*; and one dataset from the MIT Reality Mining Project [7], referred to as *Reality*. In *Cambridge*, iMotes were distributed mainly to two groups of students from University of Cambridge Computer Laboratory, specifically undergraduate Year 1 and Year 2 students, and also some PhD and Masters students. In *Infocom05*, mobile devices were distributed to approximately fifty students attending the Infocom student workshop. In *Infocom06*, the scenario was very similar to *Infocom05* except that the scale is larger, with 80 participants. Participants were selected so that 34 out of 80 form four subgroups by academic affiliations. In *Reality*, 100 smart phones were deployed to students and staff at MIT over a period of 9 months. These phones ran software that logged contacts with other Bluetooth-enabled devices by doing Bluetooth device discovery every five minutes. The four experiments are summarised in Table I. The characteristics of these datasets, such as inter-contact and contact distribution, have been explored in several previous studies [4] [12]. We

* Pan Hui was in Cambridge when he conducted this work.

¹<http://www.huggleproject.org>

Experimental data set	Cambridge	Infocom05	Infocom06	Reality
Device	iMote	iMote	iMote	Phone
Network type	Bluetooth	Bluetooth	Bluetooth	Bluetooth
Duration (days)	11	3	3	246
Granularity (seconds)	600	120	120	300
Number of experimental devices	54	41	98	97
Number of internal contacts	10,873	22,459	191,336	54,667
Average no. of contacts/pair/day	0.345	4.6	6.7	0.024

TABLE I
CHARACTERISTICS OF THE FOUR EXPERIMENTAL DATA SETS

believe these four datasets cover a rich diversity of environments from busy metropolitan city (*Infocom06* in Barcelona) to quiet university town (*Cambridge*), with an experimental period from several days (*Infocom05*) to almost one year (*Reality*).

For the simulation evaluation, we developed a simulator called *HaggleSim*, which can replay the collected mobility traces, and simulate different forwarding strategies on every contact event. This simulator is completely driven by contact events. The original trace file is divided into discrete sequential contact events, and fed into the simulator as input. In all the simulations in this work, we divided the trace into discrete contact events with a granularity of 100 seconds.

III. ALTRUISM AND TRAFFIC MODELS

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] [18]. In general, each node (or person) can have different altruism values with respect to other nodes (or people). Here we examine several altruism distributions, including a fixed percentage of selfish nodes, and uniform, normal, and geometric distributions of altruistic values for the whole population. All altruism values are distributed between 0 and 1, where 0 stands for totally selfish and 1 stands for totally altruistic. (In the literature [5][15], altruism values are usually modeled as between -1 and 1 , where -1 stands for totally spiteful², but we will not consider spiteful behavior in this study.)

- *Percentage of Selfishness*, the percentage of selfish nodes varies between 0 and 100, and the other nodes are totally altruistic. This is the simplest altruism distribution, but we know in reality a node will usually not be totally selfish or altruistic.
- *Uniform Distribution*, the altruistic value of the whole population is uniformly distributed between 0 and 1. Uniform and normal are distributions popularly encountered in nature and can be possible models for altruism.
- *Normal Distribution*, the altruistic value of the whole population follows the normal distribution with the values normalised to between 0 and 1. Since the range of a normal distribution is from negative infinity to positive infinity, in order to be able to normalise, we adopted 5% and 95% of its Cumulative Distribution Function (CDF)

value as two cutoffs in the negative and the positive directions of its distribution function. For example, if the value of the cutoffs at the 5% point is -10 and at the 95% point is 10 , all results will be incremented by 10 and then divided by $10 - (-10) = 20$.

- *Geometric Distribution*, the altruistic values are calculated per-node pair, and each follows a distribution such that the probability decreases with the social hop-distance, k , in the following way: $P(X = k) = (1 - p)^{(1-k)} * p$, where p is the altruism value for the first hop. In order to guarantee that the maximum altruism value is 1, we normalise the altruism values against the parameter p . In real life, altruism also decreases from kinship (first hop) to farther away of the social graph [16].
- *Degree-biased Distribution*, relates the altruism to the node degree as follows:

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

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 respectively 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 because they are more willing to help other people. 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:

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

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

- *Community-biased Distribution*, considers also 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 on each node using two variables, a and e , for intra-community and inter-community altruism, respectively. For convenient presentation later, we would use $(0.7, 0.1)$ to represent 0.7 probability to carry data for intra-community and 0.1 for carrying data for inter-community.

²Sometimes termed *Byzantine* in systems literature




Fig. 1. Delivery ratio of four datasets with percentages of selfish nodes

B. Traffic Models

Here we assume asynchronous communication for the application scenario. In principle, each node on the network can create messages for any other nodes, but we want to see the impact of different communication patterns. At the beginning of the simulation, each node will pick its list of destination nodes for communication using some communication pattern distribution. We analyse uniform and community-biased traffic from a source for the communication pattern.

- *Uniform Traffic*: 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 Traffic*: Each node tend to communicate more with people inside their community and less with people in other communities [17]. This traffic pattern is determined by two parameters P_{intra} and P_{inter} , where P_{intra} is the probability of generating traffic for community members and P_{inter} is the probability for inter-community communication. $P_{intra} + P_{inter} = 1$.

Uniform traffic assumes an equal communication opportunity for each node pair, but in real life a person will not talk to the whole population with the same probability. Community-biased traffic is more realistic in that it assumes heterogeneity in communication pattern. We also analysed *Geometric Traffic* (each node creates messages for other nodes with a probability decreasing geometrically from the source), but since the results are heavily biased toward the first hop and will not involve any altruism distribution so we will not include the results in this paper.

IV. RESULTS AND EVALUATIONS

The evaluation in this section is mainly focused on the robustness of the real human mobility traces under different altruism distributions. We also look at the possible impact of the traffic patterns on overall system throughput. Since not all the datasets we can access have *a priori* community information, we concentrate here on the robustness of the mobile networks under different altruism distributions. We study community-biased traffic in the next section with static social network models.

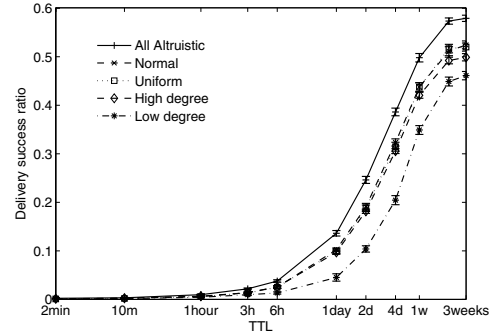


Fig. 2. Delivery ratio with uniform traffic on Reality data

Figure 1 shows the delivery ratio of the four datasets with the percentage of selfish nodes in the system. It seems that the delivery ratio decreases linearly with the percentage of selfish nodes. It is less tolerant to selfish nodes than the static network cases, but it makes sense here that the delivery ratio of a mobile network is proportional to the number of users participating in the system. Figure 2 and Figure 3 show the delivery ratio of the Reality and Cambridge datasets, respectively. Each graph shows the delivery ratio under all-altruistic, normal, uniform, high-degree-biased, and low-degree-biased altruism distribution. The x-axis shows the TTL of each message, which is the time each message can stay in the system (messages will be removed if they stay in the system over their specific TTL) and hence also indirectly reflects the delays of the deliveries. From both figures, we can see that normal, uniform, and high-degree-biased distributions yield very similar delivery ratios (their confidence intervals are almost overlapping). In addition the delivery ratios of these distributions are very close to the performance of the “All Altruistic” case, especially for the Cambridge datasets.³ The low-degree-biased case has lower delivery ratio in both cases. It seems that high degree nodes are more important for improving the system throughput. This matches the previous study that the node betweenness centrality values are highly correlated with the node degree in these mobility traces [12]. But even for this worst case, the delivery is still no less than half of “All Altruistic” for low TTL case and 68% for the Reality and 87% for the Cambridge data. Similar results also observed in other datasets.

Figure 4 shows the delivery ratio for the Cambridge and MIT Reality datasets with varying intra- and inter-community altruism. For Cambridge, we have *a priori* information about the grouping; for Reality, we use community detection algorithm to cluster them, which may not be ideal but can still provide some insights. We set the TTL to one day for Cambridge and one week for Reality, respectively, to produce reasonable delivery ratios (see Fig. 3 and Fig. 2). We increase the altruism value at 0.1 per step. From Figure 4 we can

³The 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 very sparse. Since the dataset lasts for 9 months (the longest available dataset) it still makes a good subject for social network study, and we include it here.




Fig. 4. Delivery ratio of Cambridge and Reality data with varying intra- and inter-community altruism

see that there is no impact from intra-community altruism but only from inter-community altruism (see the horizontal stripes on the figures), and the increase in delivery ratio is not very significant (from 0.58 to 0.62 for Cambridge, and 0.34 to 0.40 for Reality). The main reason is that a large fraction of deliveries are for nodes within the same community, so that even the intra-community altruism value is low; because of many intra-community contacts, the delivery ratio is still reasonable high with 0.1 altruism. Further increase in intra-community altruism value will not improve the delivery within the same communities with that large value of time-TTL, but increase in inter-community altruism will increase the chance of spreading the message to other communities. We have plotted the graph when time-TTL equal to three hours for the Cambridge case. The intra-community altruism is dominant in this case and we observed vertical stripes, but the differences in delivery ratio from altruism equal to 0.1 to 1.0 are still not very significant (from 0.22 to 0.27). These results further show the robustness of the opportunistic network because of the presence of multiple forwarding paths.

V. EVALUATION USING SOCIAL NETWORK MODELS

Human mobility traces are ideal for mobile network studies, but due to the limitations of mobility experiments (e.g. high equipment cost and difficulty in recruiting volunteers), there are very few traces available. In order to draw a more general conclusions, and to be able to study community-biased traffic, we choose to use static social network topologies and simulate asynchronous messaging in the network. This approach is very popularly used in evolutionary biology to understand the impact of the underlying network topology. We consider this approach here because, as it is well known in sociology [18], where altruism is closely dependent on kinship and social relationship. Whether people would sacrifice their batteries to

relay messages for others is likely to depend on the social links between the requesters and the requested nodes. Hence, by using social network models, we can simulate message spreading in the system with different altruism distributions and communication traffic patterns.

Here we study information dissemination on two different topologies popularly used in modeling human social networks, namely, the Caveman model [21], and the Kumpula model [13]. In the simulation, we use asynchronous messaging on these generated topologies. We set the Time-to-Live (TTL) value of all messages to the diameter of the current network. The diameter of a network is the longest shortest-path between any two vertices in the network. Each source node broadcasts its message to its neighbours only once. If a node receives a message of zero TTL value, it will drop the message irrespective of the node's current altruistic state. If the TTL of the received message is greater than zero, and the message has not been relayed by this node before, the node will decide whether to relay this message to its neighbours according to its altruistic state. A message may arrive at the same node multiple times. This is similar to epidemic routing [20] but with a limit on spread. We are interested in know how the choice of altruism distribution changes the message dissemination.

A. Social Network Models

1) *Caveman Model*: The caveman model is a way to construct graphs with communities proposed by Watts [21]. A graph is built starting from K fully connected graphs (representing communities living in isolation, like primitive men in caves). According to this model, every edge of the initial network is re-wired to point to a node of another cave with a certain probability p . The re-wiring process is used to represent random interconnections between the communities. Individuals of one cave are closely connected, whereas populations belonging to different caves are sparsely connected. Therefore, the social networks generated using this model are characterized by a high clustering coefficient and low average path length. It has been proved that this model is able to reproduce social structures very close to real ones [21].

2) *Kumpula Model*: Kumpula *et al.* [13] present a network model where the weights are generated dynamically and they shape the developing topology. The Kumpula model is characterised by two attachment schemes, the local attachment (LA) where links are added between neighbour nodes, and the global attachment (GA) where links are added between nodes pairs further away on the topology. The local attachment is based on the concepts that people have closer relationship in social life have higher chance of socialising with each other again and also he fact of mutual friends/encounters (e.g. birthday parties) in social life. The global attachment is based on the concept of meeting random people outside the current social circle in the daily life. The Kumpula model also implements the node deletion (ND) mechanism, where some nodes are removed and replace with new nodes, which is similar to the forgotten process of social life (some friends

Community generation model	Caveman	Kumpula
Number of nodes	1000	1000
Number of edges	4500	4511
Diameter	10	7
Clustering coefficient	0.7399	0.2965
Max degree	13	33
Number of communities	48	17
Average size of communities	20.8333	58.8235

TABLE II
CHARACTERISTICS OF THE SOCIAL TOPOLOGIES

leave the current social circle and new friends join in). Details can be found in [13].

Table II summarises the characteristics of the two models we used in this paper. In order for the topologies to be comparable to each other, we keep the number of nodes to be approximately equal to 1,000 and the number of edges to be approximately equal to 4,500 with an average of 4.5 friends for each node in a total population of 1,000.

After the generation of the topologies, a community detection algorithm proposed by Clauset *et al.* is applied to split the topologies into communities, using modularity optimization [6].

B. Results

We first evaluate the delivery ratio of Caveman and Kumpula models with uniform and normal altruism distributions. With uniform and normal altruism distributions, the delivery ratios for Kumpula model are all above 0.90 irrespective of the percentages of biased traffic within the same communities. Caveman model has slightly different property due to its large number of communities, the delivery ratio increase slightly from 0.75 to 0.91 in the uniform case and from 0.80 to 0.92 in the normal case when the percentage of community traffic increases from 0% to 100%.

To investigate the different impacts of intra-community altruism and inter-community altruism, we vary the intra-community altruistic value a and inter-community altruistic value e and observe the change in the delivery ratio of the system. We plot the contour graphs of the delivery ratio for all values of a and e between 0 and 1. Figure 5 shows the contour graphs for the Caveman model with four different community-biased traffic models, with community traffic (P_c) equal to 0, 0.3, 0.7, and 1.0, respectively. When $P_c = 0$, it is obvious that inter-community altruism has larger impact on the overall delivery of the system (we can see from the dominance of the horizontal stripes). Intra-community altruism starts to help improve the performance of the system when the inter-community altruism grows above 0.4. A combination of around (0.4, 0.6) can almost guarantee 100% delivery. The $P_c = 0.3$ case is quite similar to the $P_c = 0$ case except that intra-community altruism starts to show its effect sooner, so we see some changes in the low delivery region. The high delivery region (above 0.55) remains more or less the same. The empty area (delivery above 95%) remains similar. When $P_c = 0.7$, the high delivery region (delivery ratio about 0.55) is much bigger than the previous two cases. The minimum




Fig. 5. Delivery ratio of the Caveman model with varying intra- and inter-community altruism




Fig. 6. Delivery ratio of the Kumpula model with varying intra- and inter-community altruism

delivery ratio is now 0.35 instead of 0.05 in the $P_c = 0$ case. Intra-community altruism becomes much more dominant in the $P_c = 1$ case (we can see the vertical stripes dominate the entire graph). Figure 6 shows the same cases for the Kumpula model. It follows a similar trend as the Caveman model, but we can say that the Kumpula topology is less sensitive to the traffic models since the changes are not so significant compared with the Caveman model. This can be explained by the fact that the Caveman model has much more communities than the Kumpula model.

To compare with non-social graphs, We also evaluate message spreading on the Erdős-Rényi (ER) [3] and scale-free (SF) models [2], which are two popular random graph models. Both models have minimal response to the change of traffic pattern, and there is almost no change from $P_c = 0$ to $P_c = 0.7$. This is because they are actually not good social network models.

A conclusion we can draw from the results of the static models is that information dissemination on the social network topologies is quite robust toward the altruism distribution. Topology has little influence on delivery ratio if the altruism follows uniform and normal distributions. Traffic pattern

has significant impact if the altruism distribution is also community-biased. In this case, the topologies are even more robust towards altruism distribution.

VI. RELATED WORK

Altruistic behavior has been well studied in the literature using a game theory approach. For example [15] and [14] model players' payoff as the altruism coefficients of themselves and also what they believe their opponents' coefficient to be. We do not model the gaming behavior of nodes in our system; instead, we are more interested in the system throughput when the system reached its steady state distribution. Therefore, we assumed static distributions of altruism for our study and will look at the more dynamic gaming strategies in future. [5] 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. Again, 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 distribution. Altruism has also been studied in cooperative services. [1] describes a general approach to construct cooperative services, which tolerate *Byzantine, Altruistic, and Rational* behavior. Altruistic nodes in the system are the nodes which follow the suggested protocol exactly, hence they are not a challenging point for the protocol design. In P2P research, Piatek *et al.* have explored altruistic behavior in BitTorrent empirically [19]. They found that high capacity peers send much faster than the minimal rate required for reciprocation, similarly to our high degree-biased altruism modeling. Our work focuses on evaluating the performance of information dissemination in social networks and mobile networks with selfish or partially altruistic nodes, and we could not find this type of study in the prior literature.

VII. CONCLUSION AND FUTURE WORK

Opportunistic communication and information dissemination in social networks are surprisingly robust toward the form of altruism distribution, largely due to their multiple forwarding paths. This supports the case for building opportunistic networking applications in terms of one human factor. In this work, we only consider the static altruism distribution at the steady state. It would be interesting to study the altruism values resulting from gaming strategies and feedback from prior delivery history to the future. Currently, we assume each user will not memorise its own previous rejection history for message delivery of other nodes. In the future, we can study variations of this. For example, limit the number of delivery requests for a single user to avoid excessive requests, which can result in higher delivery. One more thing that would be interesting to look at would be the power consumption of

nodes, and how it is affected by altruism, for example to see how much power a node might save by choosing to be selfish compared with being only altruistic within its community.

VIII. ACKNOWLEDGEMENTS

This work is funded in part by the EC IST SOCIALNETS - grant agreement number 217141. Vito Latora was supported by the Engineering and Physical Sciences Research Council - grant number EP/F013442/1.

REFERENCES

- [1] A. S. Aiyer, L. Alvisi, A. Clement, M. Dahlin, J.-P. Martin, and C. Porth. Bar fault tolerance for cooperative services. *SIGOPS Oper. Syst. Rev.*, 39(5):45–58, 2005.
- [2] A. L. Barabasi and R. Albert. Emergence of scaling in random networks. *Science*, 286(5439):509–512, October 1999.
- [3] B. Bollobas. Random graphs. 2001.
- [4] A. Chaintreau, P. Hui, et al. Impact of human mobility on the design of opportunistic forwarding algorithms. In *Proc. INFOCOM*, April 2006.
- [5] P.-A. Chen and D. Kempe. Altruism, selfishness, and spite in traffic routing. In *EC '08: Proceedings of the 9th ACM conference on Electronic commerce*, pages 140–149, New York, NY, USA, 2008. ACM.
- [6] A. Clauset, M. E. J. Newman, and C. Moore. Finding community structure in very large networks. *Physical Review E*, 70:066111, 2004.
- [7] N. Eagle and A. Pentland. Reality mining: sensing complex social systems. *Personal and Ubiquitous Computing*, V10(4):255–268, May 2006.
- [8] K. Fall. A delay-tolerant network architecture for challenged internets. In *Proc. SIGCOMM*, 2003.
- [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] P. Hui, A. Chaintreau, et al. Pocket switched networks and human mobility in conference environments. In *Proc. WDTN*, 2005.
- [12] P. Hui, J. Crowcroft, and E. Yoneki. Bubble rap: Social-based forwarding in delay tolerant networks. In *MobiHoc 08: Proceedings of the 9th ACM international symposium on Mobile ad hoc networking & computing*, May 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] D. K. Levine. Modeling altruism and spitefulness in experiment. *Review of Economic Dynamics*, 1(3):593–622, July 1998.
- [16] E. A. Madsen, R. J. Tunney, G. Fieldman, H. Plotkin, R. Dunbar, J. Richardson, and D. McFarland. Kinship and altruism: A cross-cultural experimental study. *British Journal of Psychology*, 98:339–359, May 2007.
- [17] A. Mislove, A. Post, P. Druschel, and K. P. Gummadi. Ostra: leveraging trust to thwart unwanted communication. In *NSDI'08: Proceedings of the 5th USENIX Symposium on Networked Systems Design and Implementation*, pages 15–30, Berkeley, CA, USA, 2008. USENIX Association.
- [18] S. Okasha. Altruism, group selection and correlated interaction. *British Journal for the Philosophy of Science*, 56(4):703–725, December 2005.
- [19] M. Piatek, T. Isdal, T. Anderson, A. Krishnamurthy, and A. Venkataramani. Do incentives build robustness in bittorrent? In *NSDI'07*, Cambridge, MA, April 2007.
- [20] A. Vahdat and D. Becker. Epidemic routing for partially-connected ad-hoc networks. Technical report, Duke University, 2000.
- [21] D. J. Watts. *Small Worlds The Dynamics of Networks between Order and Randomness*. Princeton Studies on Complexity. Princeton University Press, 1999.