Social-Based Trust Mechanisms in Mobile Opportunistic Networks

Abderrahmen Mtibaa
School of Computer Science
Carnegie Mellon University
Email: amtibaa@cmu.edu

Khaled A. Harras
School of Computer Science
Carnegie Mellon University
Email: kharras@cs.cmu.edu

***Abstract*—The fundamental challenge in opportunistic networking, regardless of the application, is enabling node cooperation to forward a message. While node cooperation is considered as a fundamental property in such networks, ensuring such a property between two devices in mobile opportunistic networks remains largely unexplored. In this paper, we investigate the potential impact of the lack of trust on node cooperation. We adopt a real-trace driven approach to study and analyze the trade-off between trust and success delivery rates in opportunistic networks. We explore leveraging social information to establish trustworthy communication for mobile opportunistic networks. We propose six trust based filters that use three social-based estimators of trust including common interests, common friends, and the distance in the social graph, coupled with two major techniques of trust establishment including Relay-to-Relay, and Source-to-Relay. We show that our trust filters achieve a good trade-off between trust and success rate by achieving more than 35% success rate compared to an un-trusted environment where 10% of the nodes refuse to cooperate in the absence of trust.**

I. Introduction

The proliferation of mobile devices such as netbooks, smart phones, laptops, sensors, wireless headsets, etc. makes opportunistically connecting theses devices a challenging area of research for a large spectrum of applications. The most challenging problem in such heterogeneous and opportunistic networks remains the multi-hop data transfer decision-making between devices since they are often disconnected from each other. Typically, data is stored, carried, and forwarded over time in an opportunistic hop-by-hop manner [3], [10]. In all these scenarios, devices *must cooperate* in order to make end-to-end message delivery possible and efficient.

In mobile opportunistic networks, nodes' cooperation is fundamental for the message delivery process. Therefore, the lack of nodes' cooperation (*e.g.,* a node may refuse to act as a relay and settle for sending and receiving its own data) causes considerable delay degradation in the network. As we demonstrate in this paper, there is a large potential impact of excluding a few nodes (*e.g.,* they do not cooperate in the message delivery process) on the overall performance of the network. To deal with this issue, opportunistic networks ought to ensure node cooperation relying on two major strategies: (*i*) enable trust across communicating entities, and (*ii*) integrate incentives into the operation of opportunistic networks.

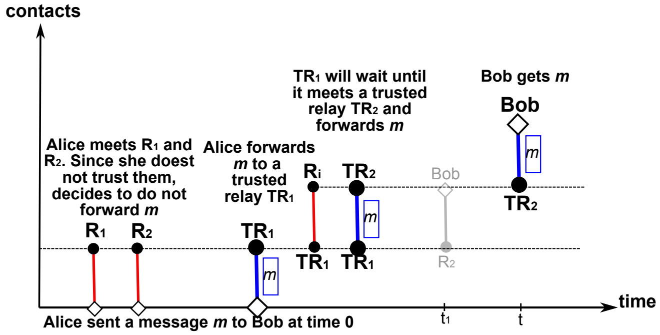
contacts
Alice meets R1 and R2. Since she doest not trust them, decides to do not forward m
Alice forwards m to a trusted relay TR1
TR1 will wait until it meets a trusted relay TR2 and forwards m
Bob gets m
Bob
R1
R2
TR1
Ri
TR2
TR1
TR1
Bob
R2
TR2
m
m
m
time
Alice sent a message m to Bob at time 0
t1
t

Fig. 1: Trust establishment in opportunistic communication

In this paper, we focus on trust establishment in opportunistic communication. The establishment of trust is more challenging in infrastructure-less networks such as opportunistic networks where no centralized mechanisms can be easily deployed. Fig. 1 shows an opportunistic communication scenario between Alice and Bob. In the absence of trust, Alice will try to maximize the probability of reaching Bob by sending her message m to all other relay nodes R_i in the network. To avoid unwanted communication and establish a trusted environment that increases nodes' cooperation, Alice will only forward her message m to relay nodes that she can trust (trusted relay TR_i). However, this trust-based communication environment may introduce additional delay by filtering out unwanted communication opportunities (*e.g.,* Bob could be reached at time $t_1 \ll t$ through R_2, but since Alice did not trust R_2, she will miss an opportunity to reach bob with shorter delay).

We leverage social relationships between users to facilitate and enable trustful communication between users. We propose and study a set of social trust filters to identify the subset of contacts between nodes that are allowed in the forwarding path as shown in Fig. 1. We utilize explicit social information coupled with real human mobility trace to establish trustworthy communication between a particular node and: (*i*) a relay (Relay-to-Relay based trust), or (*ii*) the source node (Source-to-Relay based trust). Relay-to-Relay based trust uses a transitive trust approach to establish a trusted path between a source node S and a destination node D relying on trusted communication between every two successive relay nodes on this path. Source-to-Relay based trust requires a pre-establishment of trust between the source node and all relay nodes used in the forwarding path. We couple these two approaches with three social-based trust filters (common friends, common interests, the distance in the social graph) to introduce and study six forwarding trustworthy techniques to enable trustworthy communication.

We evaluate these six filters over real mobility datasets coupled with social information [12], [4]. Our results depict that simple social relationships between users can be utilized to ensure trust in opportunistic networks. While these results show a cost incurred for this trust manifested in additional end-to-end delays, this cost is justified by the fact that in the absence of trust, users may not cooperate and the performance then drops significantly. Our trust filters, therefore, yield a fair trade-off between trust and success rate by achieving more than 35% success rate compared to an untrusted environment where only 10% of the nodes refuse to cooperate in absence of trust.

II. Related Work

Trust establishment in communication has been a very challenging problem especially with the proliferation of communication devices. This problem is further exacerbated with online applications such as email, VoIP, P2P sharing systems, and online social networks. Different approaches were recently proposed to address this problem.

Content-based filtering is the most popular approach utilized to especially combat email spam [11], [8]. Moreover, many online content sharing systems such as YouTube and Flicker use content filtering based on users' rates; users rate the content item they have viewed to other users to identify relevant content and avoid unwanted content. Content filtering approach, however, suffer from false positives and false negatives concerns and relies on centralized servers which make it difficult to adopt in fully distributed systems such as mobile opportunistic networks. Another approach, which can be easily deployed in a distributed manner, is simply charging the sender of a message; we impose a cost on the sender. This approach attempts to copy the postal service model and later on SMS. They are based on the assumption that the cost will discourage people from widely distributing content. [1], [14] claim that it is surrealistic to adopt a decentralized payment system that charges the sender per message sent.

Most relevant to online social networks, there is the Whitelisting approach. This approach is mostly adopted in VoIP systems, P2P content sharing systems, and online Social networks. End-users must exchange a request invitation; if the invitation is accepted the user is added to a white-list and will be able to send messages. Hence, no communication is allowed if one party declines the request invitation and the user is blocked by adding him to a blacklist.

In the context of mobile Ad-Hoc networks and DTNs, researchers have proposed many approaches for trust establishment based on mobility and contact characteristics [9], or content-based reputation techniques [2], [15], [5]. In [5], authors have considered how incentives can be integrated into the operation of a mobile ad hoc network, using link and node characteristics such as bandwidth and power usage to determine prices in a distributed fashion. In DTNs, [15] shows that by considering strategies that take into consideration the nodes' cooperation, one can aid the effects of non-cooperative behavior in DTNs. Although these approaches uses contact properties and the number of messages transferred to determine cooperation metrics, we believe that explicit social information could be considered as a better estimation for trustful communication in mobile opportunistic networks.

	CoNext07	CoNext08	Infocom06
duration	3 days	3 days	3 days
mobility patterns	Bluetooth contacts	Bluetooth contacts	Bluetooth contacts
social patterns	MobiClique app. [12]	Facebook + MobiClique app. [12]	Facebook
# connected nodes	27	19	47
# edges	115	102	219
average degree	9.5	9.2	9.3
social diameter	4	4	4
median inter-contact	10 min	10 min	15 min
median contact time	240s	180s	150s

TABLE I: Characteristics of our experimental data sets

III. Data sets & Methodology

We are interested in delivering data among a set of N mobile wireless nodes. Communication between two nodes is established when they are within radio range of each other. Data is forwarded from source to destination using these opportunistic *contacts*. We model the evolution of contacts in the network by a time varying graph $G(t) = (V, E(t))$ with $N = |V|$. We assume that the network starts at time t_0 and ends at time T (T can be infinite). We call this temporal network [7] the *contact graph*. Paths in such temporal network are constructed as a concatenation of contacts following a chronological property. Among these paths, a path from $s \in V$ to $d \in V$ starting at time t_0 is delay-optimal if it reaches the destination d in the earliest possible time. All the path construction rules we consider fits in the following general model: depending on the source s and the destination d, a rule defines a subset of contacts between pairs of nodes $(u \to v)$ which are allowed in forwarding path.

Our analysis relies on three datasets collected in conference environments [12], [4]. In addition to human mobility information, these datasets contain social relationships between the experimentalists. A summary of the corresponding parameters is given in Table I.

CoNext07 dataset [12] contains mobility information and information about the social relation between the participants during ACM CoNext 2007 conference. During the experiment, the social networking application indicated when a contact, or a contact of a contact, was in Bluetooth range/neighborhood. This connection neighborhood was then displayed on the user's device which in turn could add new connections or delete existing connections based on the physical interaction consequent to the application notification.

CoNext08 dataset [16] was performed at ACM CoNext 2008 conference using 22 smartphones. Social profiles were initialized based on the user's Facebook profile. The social network, then, evolved throughout the experiment as users could make new friends and discover (and create) new groups (*i.e.,* interest topics) and leave others. For the analysis we consider the collected contact trace and the final social graph of 19 devices (the rest of devices were not collecting data on each day of the experiment).

Infocom06 dataset [4] was collected with 78 participants during the IEEE Infocom 2006 conference. People were asked to carry an experimental device (called iMote) with them at all time. These devices were logging all contacts between participating devices. Questionnaires were given to participants to fill theirs nationalities, languages, countries, cities, academic affiliations and topic of interests. In this paper, we consider a social graph based on users Facebook friendship graph (obtained offline).

IV. Trust vs. Efficiency in Opportunistic Networks

In this section, we investigate the trade-off between trust and efficiency in opportunistic networks. We propose and study a set of social trust filters. These filters use explicit social relationships between nodes to ensure a trust communication in mobile opportunistic networks. We, first, motivate the use of trust filters in opportunistic networks. We next evaluate, based on real mobility traces, the performance of a trustful communication in opportunistic networks (using different sets of social trust filters).

A. Node Cooperation in Opportunistic Networks

The lack of trust between nodes may lead to potential dissatisfaction amongst them. This may cause a decrease of nodes' cooperation in the forwarding process leading to considerable delay degradation in the network. On the other hand, a pre-established trust environment may encourage nodes to cooperate, contribute to the forwarding process, and participate longer in the system.

We now demonstrate that excluding a small set of popular[1] nodes impact the overall message delivery performance in the network. These nodes, while popular, are more exposed than other to unwanted communication. They require (more so than other nodes) an establishment of a secure and trusted environment for opportunistic communication.

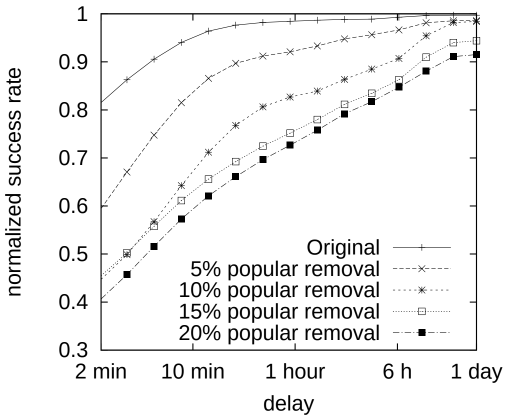

Fig. 2: Impact of excluding nodes on the overall performance

Social Filters / Trusted Entity	d-Distance	Common Friends	Common Interests
Relay-to-Relay	R2R: d-distance	R2R: CF	R2R: CI
Source-to-Relay	S2R: d-distance	S2R: CF	S2R: CI

TABLE II: Social-based trust filters in opportunistic networks

Fig. 2 shows the impact of the lack of cooperation of a small set of popular nodes on the overall performance in opportunistic networks. It plots the distribution of the success delivery rate with and without these popular nodes. We observe that the probability to reach a destination within 10 minutes drops from 96% to 83% when 5% of the popular nodes are excluded. Moreover, you may notice that the success rate regress is more significant when we remove the most popular nodes; while we show a 13% regress (from 96% to 83% within 10 minutes timescale) when only 5% of the most popular nodes are excluded, this regress is only by 3% when we exclude 5% of the remaining 90% popular nodes (*i.e.,* we compare 10% and 15% removal nodes curves).

To summarize, we observe that excluding a small set of popular nodes from the message delivery process leads to a major regression in the overall performance. It is therefore crucial to further satisfy these nodes by ensuring a trusted communication environment. In the following, we study the trade-off between trust and efficiency in mobile opportunistic networks using real mobility traces. We introduce six trust based filters, and compare the performance of the resulting trusted environment to the Epidemic routing where all nodes cooperate (*i.e.,* optimal, and idealistic environment).

B. Social-Based Trust Filters

We present a set of trust filters that uses social relationships between users to establish trustful communication between these users. Two users may trust each other if they share i common interests, have f common friends, are friends, or if they are both friends of the message's sender user, etc.

We introduce, in Table II, two major techniques for trust establishment: (*i*) Relay-to-Relay based trust, and (*ii*) Source-to-Relay based trust. Relay-to-Relay based trust uses a transitive trust approach to establish a trusted path between a source node S and a destination node D relying on trusted communication between every two successive relay nodes on this path. Source-to-Relay based trust requires a pre-establishment of trust communication between the source and all relay nodes used in the path. We couple these two techniques with three social estimators of trust based on the distance in the social graph, common interests, and common friends.

[1] Nodes' popularity computed based on the PeopleRank algorithm [13] that measures the relative importance of a node in a mobile opportunistic network.

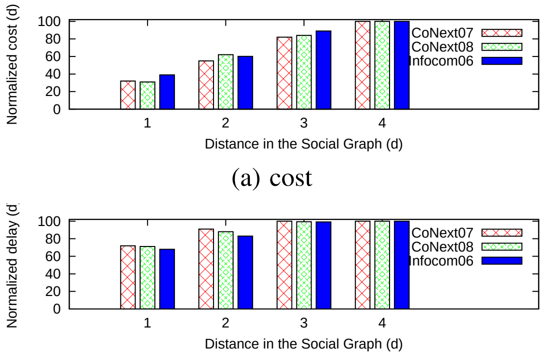

(a) cost

(b) Success delivery rate

Fig. 3: Performance evaluation of R2R: d-distance

We utilize two main metrics to evaluate the proposed trust filters: (*i*) the *normalized success rate within time t*: the probability of f to successfully deliver the message to its destination within time t normalized by the same probability given by epidemic forwarding algorithm [17] (idealistic scenario with no trust: optimal success rate within the same time t), and (*ii*) the *normalized cost*: the fraction of contacts (*i.e.,* number of replica copies) used by f normalized by the fraction of contacts used by epidemic forwarding (the most expensive).

C. Relay-to-Relay based Trust

The basic filters to estimate trust communication between two nodes i and j utilize simple social properties between these two nodes i and j. Following, a set of social estimation of trusted communication between these two nodes.

1) **R2R: d-Distance:** The distance in the social graph between two nodes can be a good trust estimation metric. The shorter the social distance they have from each others, the more they trust each others; friends ($dist = 1$) trust each others more than friends-of-friends ($dist = 2$). Where $dist(i, j)$ measures the shortest distance between two nodes in a graph.

We introduce the *d-distance trust filter* ds such that only the contacts $< i, j >$ allowed in the forwarding process satisfies $d\text{–}distance(i, j) = dist(i, j) \geq d$, where $d = 1..D$ is an integer no greater than the diameter D of the social graph.

We plot in Fig. 3 the success delivery rate and the cost of ds filter approach. Our results show that trusted communication may not be efficient, and the more trust a node expects the less success delivery rate it gets. For example, if we only allow direct friends to be involved in the forwarding process (*i.e.,* $d = 1$), we reduce the cost to less than 40% compared to flooding. However the success rate within a 10 minutes timescale decreases to roughly 70%. Moreover, if we relax such assumption and allow friends-of-friends to be involved in the forwarding process we increase the success rate and achieve more than 80% while using only 50% of the total contacts used in Epidemic forwarding.

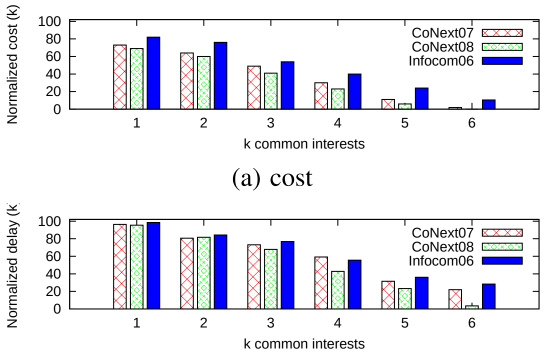

(a) cost

(b) Success delivery rate

Fig. 4: Performance evaluation of R2R: common interests

2) **R2R: Common Interests:** Common interests between two nodes has been largely considered as a good approximation of social similarities [12], [6]. People sharing one or many common interests tend to go to similar locations, events, etc. Note that using common interests as an approximation of social similarity between two users does not require an explicit confirmation from both users. We consider the common interest based filter ci such that $< i, j >$ contacts are allowed in the forwarding process $\iff ci(i, j) = \sum_{I} 1_{i,j \in s(I)} \geq K$, where $s(I)$ is the set of users that subscribe to the interest I.

Fig. 4 evaluates the common interest based filter performance. All the data sets show similar results; the more common interests the filter requires the worse the success rates the system achieves; with roughly 50% cost (no less than 3 common interests) the overall success rate is no more than 80%. Moreover, if only users that share no less than 5 interests could participate, the performance drops to less than 30%. Besides the fact that this filter is using implicit social information where there is no explicit social connection between the end users, this filter gives poor results compared to the friendship filter ($d - distance$).

3) **R2R: Common Friends:** Friendship is usually considered as a good approximation of a strong social bond between people. However, a typical friendship graph may contain strong relationships such as family members or best friends, and also "wayward" friends or family members that someone may not communicate with at all. In social network theory, social relationships are viewed in terms of nodes and ties. Researchers differentiate then between weak and strong ties in social networks. In this paper, we focus more on strong social ties since people may not trust friends if they do not communicate often. Common friends was largely used by online social networks such as Facebook, LinkedIn, etc in their content recommendation systems.

We introduce the common friends based filter which allows only $< i, j >$ contacts in the forwarding process $\iff$ $cf(i, j) = \sum_{n} 1_{n \in F(i) \cap F(j)} \geq K$, where $F(i)$ is the set of node i's friends.

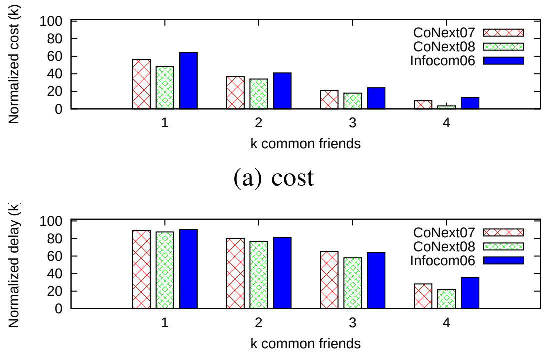

(a) cost

(b) Success delivery rate

Fig. 5: Performance evaluation of R2R: common friends

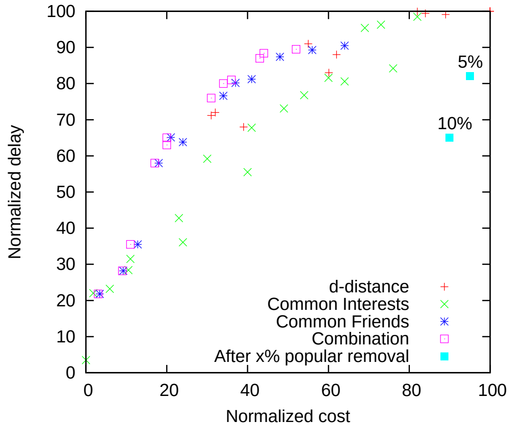

Fig. 6: Comparison of Relay-to-Relay based trust techniques

We show in Fig. 5 that by only allowing nodes that share no less than 2 friends we reduce the cost by a factor of 3 and we achieve roughly 80% of the optimal success rate. Moreover, if $k = 3$, the normalized success rate is no less than 60% with no more than 20% of the total number of contact used in Epidemic forwarding.

4) **R2R: Combination:** Inspired by the previous results that show that friendship based filters (*i.e.,* the d-distance or the common friends) give better cost/success rate performance compared to implicit social based filters (*i.e.,* common interests), we propose a combination of the two previously described friendship based filters. Note that $< i, j >$ contacts selected by the common friends filter implicitly verify the following $dist(i, j) \leq 2$. We propose a combination of 1-distance and common friends based filters.

In order to compare the combination filter to all previous approaches, we aggregate all results from all datasets in only one plot (Fig. 6). The x-axis is the normalized cost and on the y-axis we have the corresponding normalized success rate. We observe that the closer the dots are to the top left corner, the better the performance achieved; only a few contacts were efficiently used to forward messages to all the destinations. We show that the success rate of Relay-to-Relay filters exponentially increase with the cost; a near optimal performance is achieved within a cost of 70%. This implies that this technique while filtering out many communication opportunities is able to deliver a message with a near-optimal success rate performance. We show also that the combination filter outperforms all other filters proposed and achieves the best cost success rate trade-off.

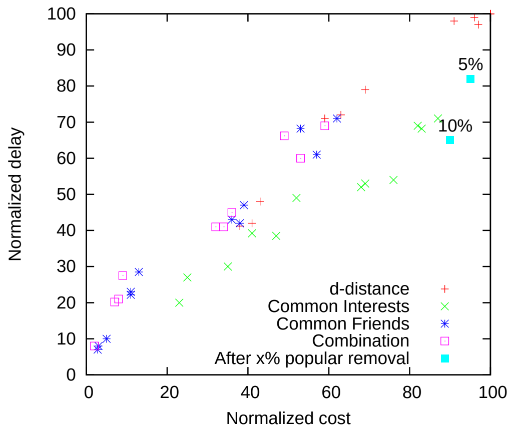

Fig. 7: Comparison of Source-to-Relay based trust techniques

In the figure, we use “x% untrusted environment” to denote the network configuration where x% of the most popular nodes are dissatisfied and decide to not participate in the forwarding process in an epidemic setting. We compare the performance of our filters to the performance of 5% and 10% untrusted environment. We observe that it is possible to ensure trust with almost an optimal success delivery rate within a 10 minutes timescale, and outperform the 10% untrusted environment by roughly 35% (with the same cost). Moreover, our filters achieve the same success rate given by 10% untrusted environment with 50% less contacts.

D. Source-to-Relay Based Trust

The Relay-to-Relay based trust technique implicitly assumes that contacts are independent from each others, and a node j should only trust the node i from which it will receive the message. Opposed to Relay-to-Relay based trust which uses hop-by-hop trust establishment, source based filters use an “end-to-end” trust approach. It estimates trust between the actual node and the source S. We compare again the three trust filters; distance in the social graph, common friends and common interests as shown in Table. II. While receiving a message, a relay node filters out all messages generated by an untrusted source node S (filtering untrusted contacts between the source and the relay node R_i that receives the message).

We compare in Fig. 7 the distribution of both cost and success rate of the Source-to-Relay filters. Overall, we observe that Source-to-Relay filters yield a poor performance; as opposed to Relay-to-Relay technique Source-to-Relay success rate increase linearly with the cost. The distance filter achieve no more than 42% success rate when d=1 (*i.e.,* only the friends of the source node S are allowed to forward messages).

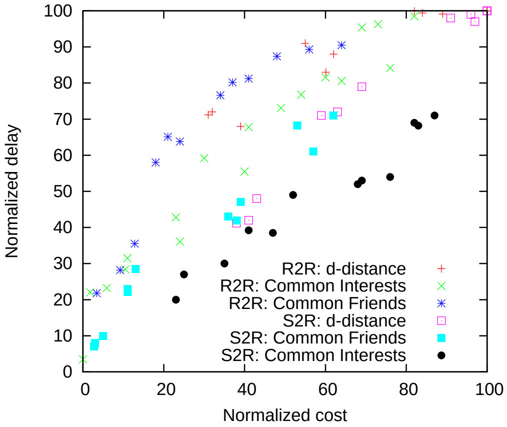

Fig. 8: Social filters comparison

Moreover, the common friends technique outperforms the two other techniques. However, it achieves only 50% success rate with 50% of the contacts. Finally, we also show that our filters still outperform the 5% and 10% untrusted environment performance by 5% to 12%.

E. Comparison

In order to compare our two proposed trust approaches, we aggregate all results from all datasets in only one plot (Fig. 8). We show that the common friends-based filter outperforms all the other studied approaches. It achieves one of the best cost/success rate trade-offs with more than 60% of success rate using no more than 30% of the contacts. Moreover, the S2R: d-distance gives the worst performance and achieves only 50% of the optimal success rate using roughly 50% of the contacts. However, since it is hard to measure trust quantitatively, this plot cannot show the best trust estimation approach. On the other hand, we show that Source-to-Relay based trust – which can be considered as a good approximation of trust – leads to a poor cost/success rate performance.

V. Conclusion and Future Work

The proliferation of online social network platforms and applications such as Facebook, Orkut, or MySpace, makes information about the social interaction of users easily accessible. In opportunistic networks, such information can be utilized for many purposes such as predicting future encounters of participating devices. In this paper, we have studied trust establishment in opportunistic networks via multiple social-based trust filters. Our work highlights the trust/efficiency trade-off in mobile opportunistic networks where social relationships between people can be considered to establish trust in opportunistic networks. We have shown that we can ensure trusted communication in opportunistic networks but this may cause additional end-to-end delay (*i.e.,* cost). Our trust filters, however, achieve a good trade-off between trust and success rate by achieving more than 35% success rate compared to an untrusted environment when 10% of the nodes refuse to cooperate in the absence of trust.

This work represents the very first step towards the idea of leveraging social networking information for enabling node cooperation in mobile opportunistic networks. There are several venues we plan to pursue in our future work. First, in this paper data sets represent a mobility trace within a single and small community. In a future work, it is important to investigate the performance of these trust filters in large scale networks. Second, an important research direction is to indicate whether these trust filters can be efficiently estimated and implemented using a distributed algorithm running with local information at the nodes. Finally, detailed mechanisms for sharing and transmitting this social information upon which trust is based is crucial.

References

[1] M. Abadi, A. Birrell, M. Burrows, F. Dabek, and T. Wobber. Bankable postage for network services. In *In Proc. Asian Computing Science Conference*, pages 72–90, 2003.

[2] S. Buchegger and J. Y. Le Boudec. A robust reputation system for P2P and mobile ad-hoc networks. In *Proceedings of the Second Workshop on the Economics of Peer-to-Peer Systems*, 2004.

[3] J. Burgess, B. Gallagher, D. Jensen, and B. Levine. Maxprop: Routing for vehicle-based disruption-tolerant networking. In *Proceedings of IEEE Infocom*, 2006.

[4] A. Chaintreau, A. Mtibaa, L. Massoulié, and C. Diot. The diameter of opportunistic mobile networks. In *Proceedings of ACM CoNext*, 2007.

[5] J. Crowcroft, R. Gibbens, F. Kelly, and S. Ostring. Modelling incentives for collaboration in mobile ad hoc networks. In *Proc. of Workshop Modeling and Optimization in Mobile, Ad Hoc and Wireless Networks (WiOpt)*, 2003.

[6] E. M. Daly and M. Haahr. Social network analysis for routing in disconnected delay-tolerant manets. In *MobiHoc '07: Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*, pages 32–40, New York, NY, USA, 2007. ACM.

[7] D. Kempe, J. Kleinberg, and A. Kumar. Connectivity and inference problems for temporal networks. In *Proceedings of ACM STOC*, 2000.

[8] S. Kumar, S. Dharmapurikar, F. Yu, P. Crowley, and J. Turner. Algorithms to accelerate multiple regular expressions matching for deep packet inspection. In *Proceedings of the conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '06, pages 339–350, New York, NY, USA, 2006. ACM.

[9] U. Kumar, G. S. Thakur, and A. Helmy. Protect: Proximity-based trust-advisor using encounters for mobile societies. *CoRR*, abs/1004.4326, 2010.

[10] A. Lindgren, A. Doria, and O. Schelén. Probabilistic routing in intermittently connected networks. *SIGMOBILE Mob. Comput. Commun. Rev.*, 7(3):19–20, 2003.

[11] G. Mishne, D. Carmel, and R. Lempel. Blocking blog spam with language model disagreement. In *In Proceedings of the First International Workshop on Adversarial Information Retrieval on the Web (AIRWeb*, 2005.

[12] A. Mtibaa, A. Chaintreau, J. LeBrun, E. Oliver, A.-K. Pietilainen, and C. Diot. Are you moved by your social network application? In *WOSN'08: Proceedings of the first workshop on Online social networks*, pages 67–72, New York, NY, USA, 2008. ACM.

[13] A. Mtibaa, M. May, M. Ammar, and C. Diot. Peoplerank: Social opportunistic forwarding. In *Proceedings of IEEE INFOCOM (mini conference)*. IEEE, 2010.

[14] A. M. Odlyzko. The case against micropayments. In *Financial Cryptography*, pages 77–83, 2003.

[15] A. Panagakis, A. Vaios, and I. Stavrakakis. On the effects of cooperation in DTNs. In *Proceedings of the 2nd International Conference on Communication Systems Software and Middleware, 2007. COMSWARE 2007*, pages 1–6. IEEE, Jan. 2007.

[16] A.-K. Pietiläinen, E. Oliver, J. LeBrun, G. Varghese, and C. Diot. Mobiclique: Middleware for mobile social networking. In *WOSN'09: Proceedings of ACM SIGCOMM Workshop on Online Social Networks*, August 2009.

[17] A. Vahdat and D. Becker. Epidemic routing for partially-connected ad hoc networks. Technical Report CS-2000-06, UCSD, 2000.