ONSIDE-SELF: A Selfish Node Detection and Incentive Mechanism for Opportunistic Dissemination

Radu-Ioan Ciobanu, Radu-Corneliu Marin, Ciprian Dobre, Valentin Cristea

Abstract Opportunistic networks are formed of mobile devices (such as smartphones and tablets belonging to social users) that communicate using close-range protocols. These networks are based on the store-carry-and-forward paradigm, where contacts between nodes are used opportunistically to transport data from a source to a destination, even though the two nodes might never be in direct communication range. Data dissemination assumes that nodes do not send directed messages (i.e., from a source to a pre-set destination), instead using channels to perform communication. Nodes are able to subscribe to channels, which are represented by interests (e.g., a node interested in IT will need to receive all messages marked with that tag). The main requirement of opportunistic networks is that the participating nodes should be altruistic, since communication is performed with the help of other nodes. However, this might not always be the case, since selfish nodes might decide that they do not want to help others. Such nodes should be detected and not allowed to participate in the dissemination process. This way, their messages will not be delivered, so they will be forced to become altruistic if they want a good networking experience. In this chapter, we propose a method for detecting and punishing selfish nodes in opportunistic networks dissemination, using gossiping mechanisms over the dynamic social network. Nodes learn about the behavior of other nodes and,

Radu-Ioan Ciobanu
University Politehnica of Bucharest, Splaiul Independentei 313, Bucharest, Romania, e-mail: radu.ciobanu@cti.pub.ro

Radu-Corneliu Marin
University Politehnica of Bucharest, Splaiul Independentei 313, Bucharest, Romania, e-mail: radu.marin@cti.pub.ro

Ciprian Dobre
University Politehnica of Bucharest, Splaiul Independentei 313, Bucharest, Romania, e-mail: ciprian.dobre@cs.pub.ro

Valentin Cristea
University Politehnica of Bucharest, Splaiul Independentei 313, Bucharest, Romania, e-mail: valentin.cristea@cs.pub.ro

when a contact occurs, share this information with an encountered device. We apply this method to an existing social and interest-based dissemination algorithm (ONSIDE) and show that it correctly detects and punishes selfish nodes, thus increasing the network's behavior in terms of message delivery and congestion.

1 Introduction

Opportunistic networks (ONs) are decentralized networks formed mostly of mobile devices such as smartphones or tablets, where no direct routes exist between nodes, and disconnections are the norm. They are based on a paradigm entitled store-carry-and-forward, which assumes that nodes generate data, store it, carry it around the network, and then forward it when the destination is met, or when an encountered node is deemed to be able to deliver the data closer to the destination than the current carrier. ONs are thus based on the altruism of nodes, since there may be situations where a source node might never be in contact with the intended destination of the data it generates. This is especially true when dealing with dissemination, since nodes generate content marked with certain tags, and expect it to be delivered to all interested nodes. Because there are no brokers or other central entities like in regular publish/subscribe-based dissemination, ONs have to rely on the fact that the participating nodes are altruistic.

However, this may not always be the case. There are situations where nodes have no interest in helping others. Because the nodes in an opportunistic network generally belong to humans, their movement is based on the social relationships between their owners. This leads to cases where nodes may be altruistic towards nodes in their own community, and selfish towards non-connected nodes. Moreover, there are situations where, because of reduced storage/computing capabilities or low battery, they might not even be able to relay data for others. Thus, selfish nodes in ONs must be detected and punished, but a difference should be made between nodes that can't relay data, and nodes that don't want to relay data. For this reason, we proposed SENSE [4], a social collaboration-based selfish node detection and incentive mechanism for opportunistic networks, which is able to improve the overall behavior of routing in an ON (in terms of hit rate, delivery latency, congestion) when selfish nodes are present. Moreover, through incentives, SENSE is able to modify the behavior of selfish nodes and convince them to collaborate.

However, SENSE was built for routing in ONs, where nodes send directed unicast messages. When dealing with data dissemination, the situation is different, because nodes that generate data are not aware of the recipients. They simply publish messages marked with certain tags (or interests, since we are dealing with nodes belonging to social human users). In these situations, nodes don't group together based only on social connections, but also based on common interests. Therefore, in this chapter we propose an improved version of SENSE that functions for dissemination in ONs. We apply it on a social-aware and interest-based data dissemination algorithm called ONSIDE [5] (which has been shown to perform better than other exist-

ing ON dissemination methods), resulting in an algorithm we call ONSIDE-SELF (ONSIDE with SENSE-based selfish node detection mechanisms). We present it here, and test it on real-life mobility traces, showing that it is able to improve the behavior of an ON where selfish nodes are present, in terms of message delivery rate and congestion.

The rest of this chapter is structured as follows. Section 2 highlights the most important existing selfishness detection methods and trust mechanisms in opportunistic networks. Section 3 presents ONSIDE, a social and interest-based data dissemination mechanism for ONs which will be used to assess the performance of our proposed solution. Section 4 shows SENSE, a method for detecting and punishing selfish nodes in opportunistic routing, which is the basis for the algorithm proposed in this chapter. Then, Section 5 presents our proposal for ONSIDE-SELF, a solution for selfish node detection and incentivising in opportunistic data dissemination. Finally, Section 6 highlights the performance of ONSIDE-SELF in an ON, while Section 7 presents our conclusions.

2 Related Work

Several selfish node detection and incentive mechanisms for mobile networks have been proposed over the years. For example, the mechanism described in [8] uses a collaborative watchdog approach to detect selfish nodes in Delay-Tolerant Networks (DTNs) and Mobile Ad Hoc Networks (MANETs) and spread this information to others. A node's perception of another node in the network can have three states: unknown (no information available), selfish or altruistic. Thus, if a node A has no information about a node B and receives a "selfish" or "altruistic" message from a node C, then A sets its perception of B according to the received information. When A already has a perception about B and the opposite information is received (e.g., A thinks B is altruistic, but C states that it is selfish), the perception is reset to the no-data state. The main drawback of this method is that nodes can easily be fooled if the malicious entities in the network act in cooperation. All the attacker nodes have to do is to ensure they make contact with the attacked node on a regular basis, and this way they can manipulate its perceptions (make the malicious nodes seem altruistic and the regular nodes seem selfish).

Instead of marking nodes as either selfish or altruistic, we propose an approach that uses values between 0 and 1 for a node's altruism since it is more realistic. Our approach is somewhat similar to [10], where gossiping is used by nodes to spread their interpretation of the monitoring level, for a faster detection of selfish nodes in the network. Another method [25] splits selfish nodes into *free riders*, *black holes* and *novas*, and uses message path analysis to separate them from other nodes.

In [7], the authors propose an ontology-based trust model, where nodes' behavior is analyzed based on direct and indirect reputation. A node A's direct trust in another node B is thus given based on A's own experiences with B, in terms of information retrieval and connection service. Based on the collected data, A performs an average

(simple or exponential) over the last experiences with B. On the other hand, A's trust in B is also computed using indirect information obtained from other nodes in the network, where other nodes' opinions about B are weighted using A's opinion of said nodes. A fade factor is also introduced by the authors, so the information collected from other nodes is of actuality. Through an ontology, a node A is able to assign one of five trust values (Very Untrustworthy, Untrustworthy, No Opinion, Trustworthy and Very Trustworthy) to its information, and only nodes with a similar or higher trust value (from A's standpoint) are able to receive the requested data.

In [16], nodes' trust is social-based, since it is argued that they belong to an opportunistic network composed of people's devices (such as smartphones). Thus, socially-connected nodes have an intrinsic trust in each other, since they are likely to interact more often in good conditions. The authors propose two major techniques of establishing trust: Relay-to-Relay and Source-to-Relay. When using the former method, a node that is carrying a message computes an encountered peer's altruism based on the relationship between the two nodes, while the latter method assumes that candidate forwarders are analyzed based on their relationship with the message's source. For each of the two trust methods, three ways of computing a node's trust are proposed: common interests, common friends and social graph distance. ONSIDE-SELF is based on a similar idea, where a node's selfishness regarding a message is estimated based on its history of forwarding similar messages (in terms of message tags, i.e., interests).

A social-based trust solution is also proposed in [19] and [20], where the social network (with its pre-established friends), its structure and its dynamics are used to create a subset of trusted nodes in the network. Moreover, nodes that are frequently co-located (the so-called familiars), as well as nodes with common tastes, are employed as the basis of the trust algorithm. Two complementary approaches are employed for social trust establishment: explicit and implicit. The former is obtained directly from the social networks, where values of 1 are assigned to a node's direct friends, and decreasing values for one-hop friends, two-hop friends, etc. The implicit trust is obtained based on node similarity (namely the degree to which two nodes' familiars coincide) and familiarity (accumulated contact time). It is shown that the explicit social trust helps identify legitimate users, whereas the implicit social trust is more useful for getting valuable opinions from nodes with similar interests.

RADON [11] is a reputation-assisted data forwarding solution for ONs that is based on the notion of positive feedback messages (PFMs). These are special confirmation messages that help the reputation mechanism monitor the behavior of a forwarder. They contain the IDs of the nodes exchanging data, as well as the number of total encounters between the two nodes and the signature of the PFM creator, and are generated by a node upon receiving a message and disseminated epidemically in the network. They are used by nodes in the network to assess the reputation of other nodes, since counters are updated for each forwarder depending on whether the PFM arrived or not in time (or if it arrived with the expected content). The main problem of this solution is the risk of congesting the network with the PFMs, especially if their TTLs are high. On the other hand, low TTLs may lead to PFMs not

reaching the intended destinations, and thus wrong opinions being formed regarding the nodes in the network.

Incentive methods for selfish nodes in ONs have also been proposed, and one such example is IRONMAN [1], which uses pre-existing social network information to detect and punish selfish nodes, incentivising them to participate in the network. Each node stores a perceived altruism value for other nodes, initialized based on the social network layout. When a node *A* meets a node *B*, it checks its encounter history to see if *B* has ever created a message for *A* that has been relayed to another node *C*. If this is the case, and *A* has encountered *C* after *B* had given it the message but *A* didn't receive the message, then *C* is considered selfish, and *A*'s perceived altruism of *C* is decreased. Whenever a node *A* receives a message from a node *B* which is not the source, *A*'s perceived altruism of *B* is increased. As an incentive mechanism, IRONMAN nodes stop relaying data for peers that are considered selfish. Thus, selfish nodes might end up not being able to send their messages, unless they become altruistic.

Another method specifically designed for disseminating data in opportunistic networks using an incentive-driven publish/subscribe scheme is ConDis [24], which uses TFT (Tit-for-Tat) as the incentive method for dealing with selfish nodes. TFT requires that nodes exchange equal amounts of data, so (for example) if node *A* asks node *B* to relay five of its messages, then it must also be willing to relay up to five of *B*'s messages in turn. This way, if selfish nodes want their messages to reach the intended destinations, they need to become altruistic. ONSIDE-SELF uses a similar mechanism, where selfish nodes are not helped by other peers unless they start relaying data for others.

A very thorough survey of security and trust management in ONs is presented in [22]. Among a multitude of security and privacy-related issues, the authors also tackle the problem of managing trust in opportunistic networks, in terms of having confidence that a node that is relayed a message will successfully deliver it towards the intended destination. The authors present existing trust solutions for mobile networks, and split them into several categories, depending on the type of trust establishment: reputation-based trust [12], social trust [20, 16], environmental trust [19] and data-centric trust.

3 Data Dissemination in Opportunistic Networks (ONSIDE)

As previously stated, in this chapter we attempt to improve an existing data dissemination method for ONs with selfishness detection and incentive mechanisms. The chosen algorithm is ONSIDE, which we proposed in [5]. ONSIDE (OpportuNistic Socially-aware and Interest-based DissEmination) is a dissemination strategy that leverages information about a node's social connections, interests and contact history, in order to decrease network overhead and congestion, while not affecting the network's hit rate and delivery latency. This is done by carefully selecting the nodes that act as forwarders, instead of simply flooding every node. ONSIDE

is based on two assumptions. Firstly, it takes advantage of the fact that nodes that have common interests (i.e., that are subscribed to the same channels) tend to meet each other more often than nodes that do not. This happens because humans generally form communities based on similar tastes and preferences (as shown in previous works [15, 14, 6]), since people sharing common interests are more likely to bond together. Because of this, we believe that data dissemination in opportunistic networks can be improved in terms of bandwidth usage and congestion by leveraging interest information when performing dissemination decisions. The second assumption made by ONSIDE is that online social network connections (such as Facebook friendships, Google+ circles or LinkedIn endorsements) are respected in an ON node's encounters. We have shown in [3] that a node encounters other socially-connected nodes with a high probability. Not only is this true, but there is also a high chance that a node encounters a second-degree neighbor.

When two ONSIDE nodes meet, each node analyzes the other's messages and decides which of them should be downloaded. This analysis is performed by calling a function for every message in the encountered node's data memory, which returns a boolean value that specifies whether the analyzed message is of interest to the current node. In this case, a downloaded message is not necessarily one that the current node is interested in, but also that other nodes that are tightly connected to the current node (either through frequent contacts, a strong social connection, or common topics) are interested in. This way, congestion is avoided by not flooding the entire network with all the messages, and the network's hit rate is increased by leveraging altruistic nodes for a quicker dissemination.

The function used by a node A to analyze a message M from a node B and to decide whether it should be downloaded is:

$$\begin{aligned} exchange(A,B,M) = (&common\_interests(A,B) \geq 1) \\ &\wedge (interested(A,M.topic) \\ &\vee (interested\_friends(A,M.topic) \geq thr_f) \\ &\vee (interests\_encountered(A,M.topic) \geq thr_i)) \end{aligned} \tag{1}$$

In the formula above, *common_interests* returns the number of topics that both A and B are interested in, so data transfers are only performed between nodes with at least one common interest. The second component of the *exchange* function is $interested$, which returns $true$ if node A is subscribed to the channel that generated message M (i.e., if it is interested in M's topic). By using this function, a node will not only download a message for itself and then drop it after use, but will also store it for others, since it is highly likely to encounter other nodes that have similar interests to its own. The $interested\_friends$ function returns the number of online social network friends of node A that are subscribed to the channel that generated M. This component has the role of further reducing the amount of messages exchanged in the network, by only requesting a message if a node's social network friends are also interested in it. This not only reduces the congestion, but also has the role of speeding up the message's delivery. thr_f is a threshold that can be varied according to the ON's and social network's densities. Finally, $interests\_encountered$ is computed

based on node A's history of encounters. It returns the percentage of encounters with nodes that were interested in messages similar to M. This function is based on the assumption that a node's behavior in an ON is predictable (as shown in [2]), so that if it encountered many nodes subscribed to a certain channel, it is likely to encounter others in the future as well. thr_i is a threshold between 0 and 1 that can be varied depending on the number of channels in the network.

A detailed analysis of ONSIDE can be found in [5], but it should be noted that results show a decrease in an ON's congestion and overhead when using ONSIDE as opposed to other data dissemination techniques for opportunistic networks (such as ML-SOR [18]).

4 Selfishness in Opportunistic Routing (SENSE)

Data dissemination and routing algorithms in ONs assume that nodes are altruistic and thus willing to help transport each other's data, for the overall benefit of the network. This assumption is at the basis of opportunistic networks, where no central entity that governs communication exists. Moreover, since there is a high degree of mobility in ONs and there are no established routes between nodes, the communication is performed opportunistically: whenever a node encounters a peer and sees an opportunity for its data to be brought closer to the intended recipients, it forwards the data to the encountered node. Therefore, it can easily be seen that the presence of selfish nodes in ONs might drastically affect the overall performance of such networks, because nodes would send their data to encountered peers and assume that it will reach the intended destinations, but the encountered nodes might be selfish and drop that data. Since ON nodes can't be notified when their data has reached the destination, information might be lost because of selfish nodes.

Thus, we proposed a method for detecting and punishing selfish nodes in opportunistic networks, called SENSE [4]. It bases its analysis on the current context, such as social knowledge or information about the device's battery. We used social information because nodes tend to interact more and be more altruistic towards members of their own community. In terms of altruism modeling, we used the community-biased model [23], which assumes that people in a community have greater incentives to carry messages for other members of the same community. Thus, altruism is modeled using an intra- and an inter-community altruism level (both between 0 and 1), with the first value being higher.

When two nodes A and B running SENSE meet, each of them starts by computing an altruism value for the other node and, based on that value, decides if it will forward data for the other node. If the two nodes decide that they are altruistic towards one another, they exchange lists of past forwards O and past receives I. When a node receives one of these lists, it updates its own list with the newly received information. This way, a node can have a more informed view of the behavior of various nodes in the network, through gossiping.

Based on the lists of past encounters, each node computes a perceived altruism value for the other node with regard to the messages stored in its own data memory. If this value is within certain limits, the communication continues and the desired algorithm is applied. If not, then the encountered node is considered selfish and is notified that its messages won't be relayed. This functions as an incentive mechanism, because if a node wants its messages to be routed by other nodes, it shouldn't be selfish towards them. Therefore, every time a node is notified that it is selfish in regard to a certain message, it increases its altruism value. If there is a social connection between the node considered selfish and the source of the message, then the inter-community altruism is increased. If the two nodes aren't socially connected, the intra-community altruism value grows.

The formula for computing perceived altruism values for a node N and a message M based on the lists of past forwards (O) and receives (I) is:

$$altruism(N,M) = \sum_{o \in O, i \in I, o.m=i.m}^{N.id=o.d, N.id=i.s} type(M.id, o.m) \times thr(o.b) \quad (2)$$

In the formula presented above, a past encounter x has a field $x.m$ which specifies the ID of the message that was sent or received, $x.s$ is the source of the transfer, $x.d$ is the destination and $x.b$ is the battery level of the source. $type$ is a function that returns 1 if the types of the two messages received as parameters are the same (in terms of communities, priorities, etc.), and 0 otherwise, while thr returns 1 if the value received as parameter is higher than a preset threshold, and 0 if it is not. Basically, the altruism computation function counts how many messages of the same type as M have been forwarded with the help of node N, when N's battery was at an acceptable level.

A detailed analysis of the effects of SENSE on opportunistic routing can be found in [4]. There, we show that SENSE performs much better than existing solutions (such as IRONMAN [1]) in terms of message delivery, latency and congestion, when selfish nodes are present in the network. Moreover, we present a battery-aware scenario which shows that SENSE can distinguish between nodes that won't disseminate data for others because they are selfish, and nodes that are low on battery and can't help with dissemination. We also show that SENSE's selfish node detection accuracy can be as high as 70%.

5 ONSIDE-SELF

We are now interested in adding the selfish node detection and incentive mechanisms used by SENSE to data dissemination (and, in particular, to ONSIDE). As previously stated, nodes may be selfish for many reasons, such as low battery, insufficient memory, lack of incentives, etc. Generally, in an interest-based dissemination environment, nodes have no reason for acting as relays for messages that are not of interest to them or to the members of their interest community. However, for

the overall effectiveness of the network, nodes from separate interest communities should help each other, because even if they don't meet very often, they might be able to deliver messages to other interested nodes that aren't encountered by the data publishers.

When describing SENSE [4], we stated that we used the community-biased altruism model to compute an altruism level for each node in the network. This model assumes that each node has two altruism values: one for nodes in its own social community (intra-community altruism), and one for nodes outside its community (inter-community altruism). The intra-community value is higher than the inter-community one, since nodes have a greater interest in helping members of their own community. When dealing with data dissemination in an interest-based environment, the situation changes. Nodes are no longer split into communities based on their social relationships, but according to common interests. Therefore, instead of having different altruism values according to types of nodes, we propose having altruism values for a message's type. Since a message generated by a publisher in an interest-based environment is tagged with a topic, we propose that each node should have an altruism value between 0 and 1 for messages tagged with a topic that the node is interested in, and another one for messages tagged with a topic that is of no interest to the computing node. This way, a node is more likely to help deliver messages that it is interested in, since this also means that it will be of interest to the encountered nodes. We propose calling these two values the common-interest altruism and the no-interest altruism, with the former being naturally higher.

By using these two altruism values, a node is able to decide whether it will accept to forward a message that it receives from an encountered node. Thus, when a node A meets a node B, it runs a data dissemination algorithm (such as Epidemic [21] or ONSIDE) to decide what messages should be relayed to B, so that it can move them forward. Then, for each message it receives, B decides what the altruism level towards its tagged interest is, and thus whether it will accept it or not. If B decides that it has no interest in carrying a certain message, it drops it. Thus, having such selfish nodes in the network clearly affects its overall efficiency, both in terms of hit rate, as well as latency. Moreover, the delivery cost may suffer significantly, since messages are being sent, but they end up being dropped by the uninterested node. For this reason, we propose improving ONSIDE with selfish node detection and incentive mechanisms.

The ONSIDE-SELF selfishness detection mechanism is similar to SENSE [4]. Namely, each node stores a history of message exchanges, split into past forwards (O) and past receives (I). These two lists are updated through gossiping at every encounter with another node, and are used to decide whether a potential relayer is suitable for receiving the data (i.e., whether the encountered node is selfish towards a certain type of message or not). The decision is made by comparing the percentage of messages of the same type with the target message (i.e., that have the same tag) that have been successfully relayed by the encountered node, with a pre-established altruism threshold. If the computed value exceeds the threshold, then that node is not considered selfish, so the message will be relayed to it. If it is selfish, then the message is not sent, and the node is notified that it is considered selfish. Nodes

marked as selfish will stop receiving help from the other nodes in the network, as a punishment, until they stop being selfish. This is the incentive mechanism used for ONSIDE, which assumes that selfish nodes wish to become unselfish in order to convince other nodes to forward their messages. Nodes stop being selfish by increasing the altruism levels until they are not considered selfish any more. The common-interest altruism level is increased if a node is considered selfish towards a message that is tagged with one of its interests, while the no-interest altruism is increased otherwise.

The formula for deciding if a node N is selfish towards a message M based on the lists of past forwards (O) and receives (I) is:

$$altruism(N,M) = \frac{\sum_{o\in O, i\in I, o.m=i.m}^{N.id=o.d, N.id=i.s} type(M.id, o.m) \times time(o,i) \times thr(o.b)}{\sum_{o\in O}^{N.id=o.d} type(M.id, o.m) \times thr(o.b)} \quad (3)$$

The formula above counts the number of message sent by any node to N which were successfully delivered by N to an interested node or to another carrier, and divides it by the total number of messages sent to N. Each node N has an ID ($N.id$), while a past encounter x has a field $x.m$ which specifies the message that was sent or received, $x.s$ is the source of the transfer, $x.d$ is the destination, and $x.b$ is the battery level of the source. Similar to SENSE, $type$ is a function that returns 1 if the two messages received as parameters have the same tag (and 0 otherwise), $time$ is 1 if the timestamp of the first parameter is lower than the timestamp of the second parameter (in this case, if the message was received by another node from N after N has received the message), and thr returns 1 if the battery value received is higher than a threshold. thr is used to remove the risk of considering a node selfish when it couldn't deliver a message due to its low battery, for example. In Section 6, we will present the results of running ONSIDE-SELF on several real-life traces, and compare its output to the one obtained by basic ONSIDE where selfish nodes are present in the network.

6 Experimental Results

This section presents our experimental setup, along with the results obtained by running ONSIDE-SELF on real-life mobility traces.

6.1 Experimental Setup

For our analysis, we used three real-life mobility traces, collected in different types of environments: Infocom 2006 [9], Sigcomm 2009 [17] and UPB 2012 [13]. Table 1 presents additional information about each trace. We ran these traces in

Table 1 Information about the mobility traces used.

Trace	Nodes	Duration	Type	Topics	Topics per node
Infocom 2006	98	4 days	Conference	27	14.53
Sigcomm 2009	76	4 days	Conference	154	15.61
UPB 2012	66	64 days	Academic	5	3.51

MobEmu [3], an opportunistic network emulator that is able to replay a trace and apply a desired algorithm when two nodes meet.

There are two metrics that we use for analyzing the obtained results. Firstly, the hit rate is defined as the ratio between the number of messages that have successfully arrived at nodes subscribed to the corresponding channels, and the total number of messages generated, multiplied by the number of subscribers to the channel. The second metric is the delivery cost, defined as the ratio between the total number of messages exchanged, and the number of generated messages multiplied by the number of corresponding channel subscribers.

Data is generated through channels that nodes are able to subscribe to. When a node is subscribed to a channel, it is interested in any data generated by that channel that it hasn't received yet. Because all three traces used for testing have interest information available, we consider that a channel is represented by a topic. Nodes can only generate data on the channels corresponding to their interests. Each node that has at least one interest generates 30 messages per day. A node that is interested in multiple topics is able to randomly choose which tag will be used for a message out of its interests. Therefore, there are 27 channels for Infocom 2006, 154 for Sigcomm 2009 and 5 for UPB 2012.

In order to highlight the benefits of ONSIDE-SELF, we compare it to basic ONSIDE, to Limited Epidemic, and to a dissemination-modified version of ML-SOR [18] (as described in [5]). Limited Epidemic is a modified version of the Epidemic algorithm [21] which doesn't assume that a node's data memory is unlimited. Instead, if the memory fills and a new message must be stored, the oldest one is replaced. In order to analyze how the various algorithms we test with fare in different conditions, we also vary a node's data memory size. Thus, a node is able to store either 20, 100, 500 or 4500 messages at once in its memory. The sizes of I and O are set to 100 each. The two ONSIDE thresholds (thr_f and thr_i) are the ones presented in [5].

6.2 Results

In this section, we present the results of running ONSIDE with selfish nodes detection and incentive mechanisms. We begin by showing the effects of adding selfishness to nodes in the Sigcomm 2009 trace, as seen in Figure 1. We ran all our tests with each node's two altruism levels (common-interest and no-interest) dis-

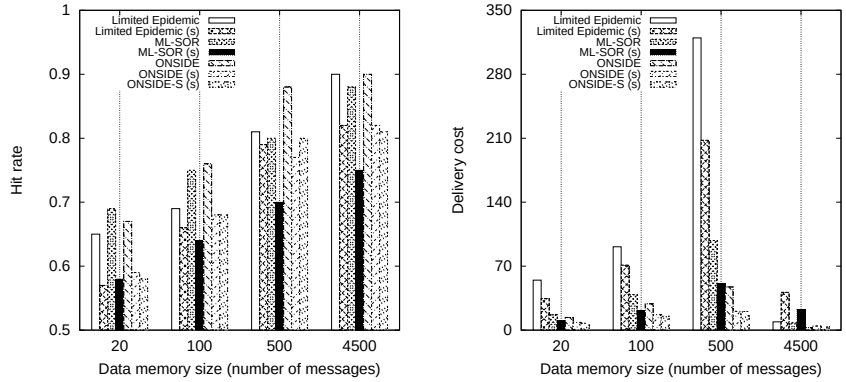

(a) Hit rate (b) Delivery cost

Fig. 1 Results for the Sigcomm 2009 trace ("s" means there are selfish nodes present).

tributed normally in the network with a mean of 0.4 for no-interest altruism and 0.6 for common-interest altruism. When there are selfish nodes in the network, the network's average hit rate drops considerably, especially for higher data memory sizes. Thus, the hit rate when nodes are selfish drops by 8.10% for Limited Epidemic, 12.54% for ML-SOR and 7.92% for ONSIDE, for a data memory of 4500. The delivery cost is affected even more, increasing by 32 for Limited Epidemic, 14 for ML-SOR and 2 for ONSIDE, when selfish nodes are present. This shows that, when nodes are not altruistic towards each other, the overall effectiveness of the network is affected, and thus even the selfish nodes suffer the consequences. This is where adding selfish node detection mechanisms can lead to improvements.

We ran the ONSIDE-SELF algorithm while varying the altruism threshold from 0.1 to 0.9, in increments of 0.1. The results in Figure 1 are for the threshold that yielded the best results, which in this case was 0.7. Figure 1(a) shows that ONSIDE-SELF is able to increase the hit rate of ONSIDE in certain situations (for data memory sizes of 100 and 500), even if the network is affected by selfish nodes. It does this by not sending messages to selfish nodes, keeping them for forwarding to more suitable sources instead. Moreover, the incentive mechanism convinces the selfish nodes to become more altruistic if they want to have their messages delivered. Figure 1(b) shows that ONSIDE-SELF is also able to decrease the delivery cost of the network, thus also decreasing the ON congestion. This happens because fewer messages are being sent in the network, since nodes that are considered selfish are ignored and no data is sent to them (data which, if the selfishness detection mechanism were correct, would end up being dropped anyway). Furthermore, ONSIDE-SELF's delivery cost is still much lower than the one obtained by Limited Epidemic and ML-SOR.

Figure 2, which presents the results for UPB 2012, shows that the situation for this trace is different. Firstly, it can be seen that the hit rate is not affected very much, even though selfish nodes are present in the network. This happens because this par-

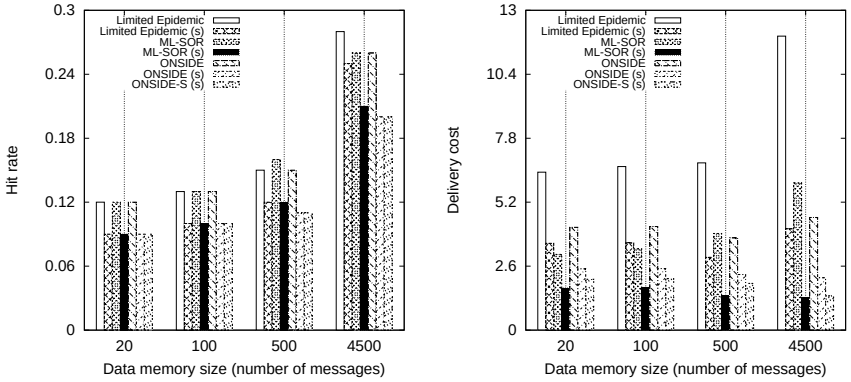

(a) Hit rate (b) Delivery cost

Fig. 2 Results for the UPB 2012 trace ("s" means there are selfish nodes present).

ticular trace doesn't have many nodes, but they encounter each other many times, since the trace duration is high, and all nodes are students or teachers at the same faculty. Thus, a node might be selfish on one encounter towards another node, and altruistic on the next encounter, depending on the altruism level. Moreover, there are very few defined interests in this trace (only five), so there is a very high chance that two nodes that meet have at least one common interest, so the common-interest altruism level (which is higher than the no-interest altruism) is used most of the time. The delivery cost of the three algorithms when nodes are selfish is lower than for the non-selfish situation, because some of the messages are dropped by the selfish nodes, so they don't get the chance to be relayed onwards, thus decreasing the number of data transfers. It can be seen in Figure 2 that ONSIDE-SELF (with an empirically-chosen altruism threshold of 0.3) doesn't bring any improvements to hit rate for UPB 2012, but it does manage to decrease the delivery cost, when compared to ONSIDE. Thus, the network becomes less congested.

Finally, Figure 3 shows the results for the Infocom 2006 trace. In terms of comparison with the results obtained when there are no selfish nodes in the network, the situation is similar to UPB 2012: the hit rate is not affected very much, and the delivery cost is decreased. However, ONSIDE-SELF (with an altruism threshold of 0.5) is able to increase the hit rate obtained by ONSIDE, while also decreasing the delivery cost by as much as 6.5 for a data memory of 4500 messages.

The results presented above show that, although ONSIDE-SELF might be able to successfully detect selfish nodes, there are situations where it can't do much about the hit rate or the delivery cost. This happens because selfishness in data dissemination is a somewhat different problem than selfishness in routing, mainly because, when we are talking about data dissemination and publish/subscribe, there isn't a single destination for a message. Instead, many nodes are interested in a message, so the message spreads much easier through the network. Although a node might be selfish, it still has to download the data it is interested in, which might be of interest

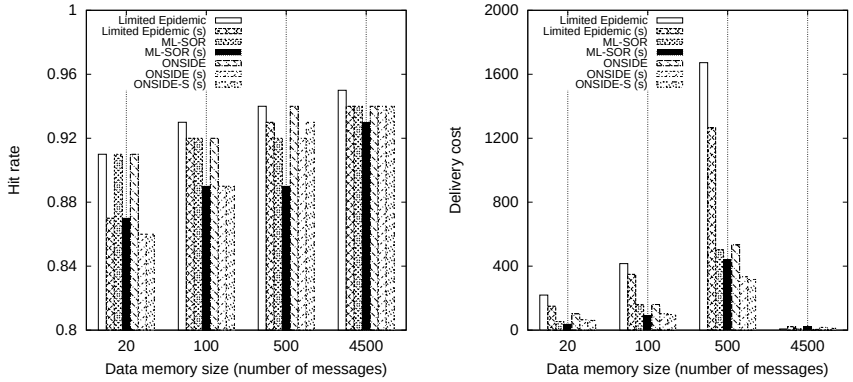

(a) Hit rate (b) Delivery cost

Fig. 3 Results for the Infocom 2006 trace ("s" means there are selfish nodes present).

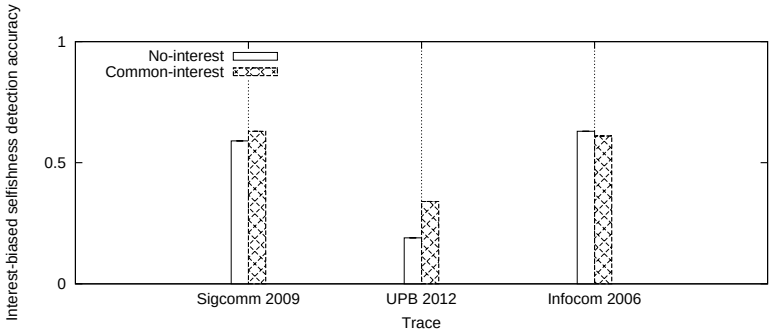

Fig. 4 Interest-biased selfishness detection accuracy results.

to many other nodes, so it invariably delivers it to them eventually. Moreover, since more nodes are interested in a message, there are more paths a message can take to reach an interested subscriber.

In order to verify the success of selfish nodes detection, we measured the interest-biased detection accuracy, which represents the percentage of nodes that end up with an altruism (common-interest or no-interest) value of 1 due to the incentive mechanisms. Having an altruism of 1 means that these nodes have been recognized by most nodes in the network as being selfish, and have thus been avoided until their altruism value increased. The results are shown in Figure 4, and it can be seen that, for the Sigcomm 2009 and Infocom 2006 traces, the detection accuracy values are around 60%, so more than half of the nodes in the ON that are selfish are recognized as such, and convinced by the incentive mechanism to be more altruistic. However, both the no-interest and the common-interest detection accuracy values for UPB 2012 are very low. The cause of this is that, as stated before, there

are very few interests in this trace, so generally nodes, although they are selfish, have an interest in delivering a message, since it is most likely of interest to them as well.

7 Conclusions

In this chapter, we have presented ONSIDE-SELF, a modified version of the SENSE selfish node detection and incentive mechanism for opportunistic networks, applied to the ONSIDE dissemination algorithm. By testing with real-life opportunistic mobility traces, we have shown that ONSIDE-SELF is able to detect a high percentage of selfish nodes, while at the same time incentivising them to become altruistic. We have shown this to lead to an overall improvement in the hit rate and the congestion of the network.

Acknowledgements The presented work is co-funded by project MobiWay, PN-II-PT-PCCA-2013-4-0321.

References

1. Bigwood, G., Henderson, T.: IRONMAN: Using social networks to add incentives and reputation to opportunistic networks. In: SocialCom/PASSAT, pp. 65–72. IEEE (2011)
2. Ciobanu, R.I., Dobre, C.: Predicting encounters in opportunistic networks. In: Proceedings of the 1st ACM Workshop on High Performance Mobile Opportunistic Systems, HP-MOSys '12, pp. 9–14. ACM, New York, NY, USA (2012). DOI 10.1145/2386980.2386983. URL http://doi.acm.org/10.1145/2386980.2386983
3. Ciobanu, R.I., Dobre, C., Cristea, V.: Social aspects to support opportunistic networks in an academic environment. In: Proceedings of the 11th international conference on Ad-hoc, Mobile, and Wireless Networks, ADHOC-NOW'12, pp. 69–82. Springer-Verlag, Berlin, Heidelberg (2012). DOI 10.1007/978-3-642-31638-8_6
4. Ciobanu, R.I., Dobre, C., Dascalu, M., Trausan-Matu, S., Cristea, V.: SENSE: A collaborative selfish node detection and incentive mechanism for opportunistic networks. Journal of Network and Computer Applications **41**(0), 240 – 249 (2014). DOI http://dx.doi.org/10.1016/j.jnca.2014.01.009
5. Ciobanu, R.I., Marin, R.C., Dobre, C., Cristea, V.: Interest-awareness in data dissemination for opportunistic networks. Ad Hoc Networks (2014). DOI http://dx.doi.org/10.1016/j.adhoc.2014.07.004
6. Costa, P., Mascolo, C., Musolesi, M., Picco, G.P.: Socially-aware routing for publish-subscribe in delay-tolerant mobile ad hoc networks. Selected Areas in Communications, IEEE Journal on **26**(5), 748–760 (2008). DOI 10.1109/JSAC.2008.080602
7. Goncalves, M.R.P., dos Santos Moreira, E., Martimiano, L.A.F.: Trust management in opportunistic networks. In: Networks (ICN), 2010 Ninth International Conference on, pp. 209–214 (2010). DOI 10.1109/ICN.2010.41
8. Hernández-Orallo, E., Serrat Olmos, M.D., Cano, J.C., Calafate, C.T., Manzoni, P.: Evaluation of collaborative selfish node detection in MANETs and DTNs. In: Proceedings of the 15th ACM international conference on Modeling, analysis and simulation of wireless

and mobile systems, MSWiM '12, pp. 159–166. ACM, New York, NY, USA (2012). DOI 10.1145/2387238.2387266
9. Hui, P., Crowcroft, J.: Bubble Rap: forwarding in small world DTNs in ever decreasing circles. Tech. Rep. UCAM-CL-TR-684, University of Cambridge Computer Laboratory (2007)
10. Lavinia, A., Dobre, C., Pop, F., Cristea, V.: A failure detection system for large scale distributed systems. In: Complex, Intelligent and Software Intensive Systems (CISIS), 2010 International Conference on, pp. 482 –489 (2010). DOI 10.1109/CISIS.2010.29
11. Li, N., Das, S.K.: RADON: Reputation-assisted data forwarding in opportunistic networks. In: Proceedings of the Second International Workshop on Mobile Opportunistic Networking, MobiOpp '10, pp. 8–14. ACM, New York, NY, USA (2010). DOI 10.1145/1755743.1755746. URL http://doi.acm.org/10.1145/1755743.1755746
12. Li, N., Das, S.K.: A trust-based framework for data forwarding in opportunistic networks. Ad Hoc Netw. **11**(4), 1497–1509 (2013). DOI 10.1016/j.adhoc.2011.01.018. URL http://dx.doi.org/10.1016/j.adhoc.2011.01.018
13. Marin, R.C., Dobre, C., Xhafa, F.: Exploring predictability in mobile interaction. In: Emerging Intelligent Data and Web Technologies (EIDWT), 2012 Third International Conference on, pp. 133–139. IEEE (2012). DOI 10.1109/EIDWT.2012.29
14. Mei, A., Morabito, G., Santi, P., Stefa, J.: Social-aware stateless forwarding in pocket switched networks. In: INFOCOM, 2011 Proceedings IEEE, pp. 251–255 (2011). DOI 10.1109/INFCOM.2011.5935076
15. Moghadam, A., Schulzrinne, H.: Interest-aware content distribution protocol for mobile disruption-tolerant networks. In: World of Wireless, Mobile and Multimedia Networks Workshops, 2009. WoWMoM 2009. IEEE International Symposium on a, pp. 1–7 (2009). DOI 10.1109/WOWMOM.2009.5282479
16. Mtibaa, A., Harras, K.A.: Social-based trust in mobile opportunistic networks. In: Computer Communications and Networks (ICCCN), 2011 Proceedings of 20th International Conference on, pp. 1–6 (2011). DOI 10.1109/ICCCN.2011.6006047
17. Pietiläinen, A.K., Oliver, E., LeBrun, J., Varghese, G., Diot, C.: MobiClique: middleware for mobile social networking. In: Proceedings of the 2nd ACM workshop on Online social networks, WOSN '09, pp. 49–54. ACM, New York, NY, USA (2009). DOI 10.1145/1592665.1592678
18. Socievole, A., Yoneki, E., De Rango, F., Crowcroft, J.: Opportunistic message routing using multi-layer social networks. In: Proceedings of the 2Nd ACM Workshop on High Performance Mobile Opportunistic Systems, HP-MOSys '13, pp. 39–46. ACM, New York, NY, USA (2013). DOI 10.1145/2507908.2507923
19. Trifunovic, S., Legendre, F.: Trust in opportunistic networks (2009)
20. Trifunovic, S., Legendre, F., Anastasiades, C.: Social trust in opportunistic networks. In: INFOCOM IEEE Conference on Computer Communications Workshops , 2010, pp. 1–6 (2010). DOI 10.1109/INFCOMW.2010.5466696
21. Vahdat, A., Becker, D.: Epidemic routing for partially connected ad hoc networks (2000)
22. Wu, Y., Zhao, Y., Riguidel, M., Wang, G., Yi, P.: Security and trust management in opportunistic networks: a survey. Security and Communication Networks **8**(9), 1812–1827 (2015). DOI 10.1002/sec.1116. URL http://dx.doi.org/10.1002/sec.1116
23. Xu, K., Hui, P., Li, V.O., Crowcroft, J., Latora, V., Lio, P.: Impact of altruism on opportunistic communications. In: Proceedings of the first international conference on Ubiquitous and future networks, ICUFN'09, pp. 153–158. IEEE Press, Piscataway, NJ, USA (2009)
24. Zhou, H., Wu, J., Zhao, H., Tang, S., Chen, C., Chen, J.: Incentive-driven and freshness-aware content dissemination in selfish opportunistic mobile networks. In: Mobile Ad-Hoc and Sensor Systems (MASS), 2013 IEEE 10th International Conference on, pp. 333–341 (2013). DOI 10.1109/MASS.2013.54
25. Zhou, Q., Ying, J., Wu, M.: A detection method for uncooperative nodes in opportunistic networks. In: Network Infrastructure and Digital Content, 2010 2nd IEEE International Conference on, pp. 835 –838 (2010). DOI 10.1109/ICNIDC.2010.5657987