Robust Reputation in Decentralised Markets

Henry Franks, Nathan Griffiths, and Arshad Jhumka

Department of Computer Science, University of Warwick, CV4 7AL, UK
{H.P.W.Franks, Nathan.Griffiths, H.A.Jhumka}@warwick.ac.uk

Abstract. Establishing cooperation and protecting individuals from selfish or malicious behaviour are key goals in open multi-agent systems. Incomplete information regarding potential interaction partners can undermine mechanisms such as trust and reputation, particularly in lightweight systems designed for individuals with significant resource constraints. In this paper, we (i) propose extending low cost reputation mechanisms with gossiping algorithms, (ii) introduce four simple aggregation strategies for incorporating gossiped information, and (iii) evaluate our model on a variety of synthetic and real-world topologies. We show that (i) gossiping can significantly reduce the potentially detrimental influence of incomplete information and network structure on lightweight reputation mechanisms, (ii) using the most recently received gossip to make decisions results in up to a 25% reduction in selfishness, and (iii) gossiping is particularly effective at aiding agents with little or no interaction history, such as when first entering a system.

1 Introduction

The emergence of decentralised open Multi-Agent Systems (MAS) has required the development of mechanisms that bias interactions towards cooperative individuals, in order to protect agents from selfish behaviour and increase the aggregate welfare of the population. Distributed systems such as BitTorrent or mobile networks demonstrate the potential for emergent marketplaces in which participating individuals cannot rely on a centralised authority to enforce compliance to cooperative norms.

Typical approaches to achieving cooperation without centralised control involve trust and reputation mechanisms, that use observations and individual experience to aid decision making. Reputation relies on *indirect reciprocity*, in which a cooperative agent can expect its behaviour to be subsequently rewarded by a third party. Many reputation mechanisms (e.g. [12, 21]) are highly robust but incur large overheads, rendering them unsuitable to domains with constrained computational resources.

A variety of low-cost solutions have been proposed but their accuracy can be reduced by environmental factors. In this paper, we empirically analyse when efficacy is reduced and demonstrate a low cost solution (i.e. gossiping) to support reputation in such situations. Specifically, we show that incomplete information, where an agent knows only a sub-set of the available information about a potential partner, can result in inaccurate reputation assessments that reduce cooperation, and that the underlying network structure significantly influences emergent behaviour based on indirect reciprocity (Section 5). We propose gossiping as a low cost solution, and introduce four alternative aggregation rules (Sections 2.3 and 3.2). We show that gossiping can reduce selfishness in the population by up to 25%, and is particularly effective on real-world network topologies (Section 5.1).

2 Background

While many reputation mechanisms have been proposed, they rarely fully address the challenges posed by decentralised MAS marketplaces. We focus on *image scoring*, which illustrates the problems that reputation mechanisms can suffer in challenging MAS domains. In this section we discuss the factors underlying mechanism efficacy, introduce the image scoring approach, and discuss gossiping algorithms as a low-cost extension to reputation.

2.1 Reputation, incomplete information, and network topology

Trust and reputation are established mechanisms for supporting cooperative behaviour in distributed and open MAS [4]. Such domains are characterised by heterogeneous agent ownership, lack of centralised authority, complex social network structures, and a low probability of repeat interactions between agents. Limited computational resources reduce the applicability of robust multi-dimensional mechanisms (e.g. FIRE [12] or ReGreT [21]). Low-cost solutions relying on observation of agent behaviour are feasible (e.g. [19]), but can be destabilised by incomplete information.

The underlying network structure that constrains agent interactions also influences the performance of the system. The networks in real-world open MAS are typically *complex*, exhibiting scale-free node degree distributions or logarithmically-bounded path lengths between nodes (i.e. *small-world* properties). Recent research has identified further structural properties currently not modelled by synthetic network generation algorithms (e.g. scale-free edge embeddedness distributions [24]). A wide variety of research has demonstrated the impact of underlying network structure (e.g. [5, 11]), but the implications of network topology on reputation has not been extensively investigated. In this paper, we use a combination of synthetic and real-world networks to demonstrate that the underlying structure has a vast influence on reputation mechanism efficacy.

2.2 Image scoring

We adopt a simplified model of reputation called *image scoring*, first introduced by Nowak and Sigmund (N&S) [17]. Agents maintain a set of reputation assessments (*image scores*) of others based on observation of interactions, and use these to determine whether to cooperate or not. Image scoring promotes cooperation without requiring reciprocity in the form of subsequent interactions between the same individuals [18]. It has low overheads making it applicable to open decentralised systems.

We consider image scoring in the *donation scenario* introduced by N&S [17], in which agents are randomly paired for interactions as *donor* (the agent who is assessed on their behaviour in the interaction) or *recipient*. The donor, based on its perception of the recipient's image score, chooses either to confer a benefit on the recipient at personal cost, or do nothing. For the purposes of discussion, we refer to the former as cooperation and the latter as defection. A subset of the participants' neighbours observe the donor's choice, and adjust their perception of the donor's image score depending on the action chosen. After a number of interactions, the best performing strategies are reproduced.

2.3 Gossiping

Gossiping algorithms, initially introduced by Frieze and Grimmet [9], perform data aggregation and spreading in distributed systems. Loosely modelled on the dynamics of human gossip, they are effective at spreading information, and have low space and time complexity and minimal bandwidth requirements compared to alternative spreading mechanisms [8, 13]. They have previously been applied to constrained trust and reputation problems [1, 20, 26], and can efficiently aggregate trust values without complex underlying data structures.

Given the low overheads, gossiping is an attractive solution to the problems arising from local perception of information by agents. Sommerfeld *et al.* [23] show that gossiping is an effective substitute for direct observation in human interactions. Sommerfeld *et al.*'s subsequent work [22] demonstrates that gossip is robust to the propagation of inaccurate information, and they conclude that humans use a majority rule: if the majority of gossips are positive, then individuals form a positive opinion of the subject. The low overheads, high robustness to inaccurate information, and ability to efficiently spread and aggregate information in decentralised domains make gossiping highly applicable to MAS.

3 Incorporating gossiping into image scoring

3.1 Image scoring model

We adopt the setup used by N&S as follows: each agent i is associated with a strategy k_i, chosen uniformly at random in the range $[-5, 6]$. Each agent maintains image scores I_i for each agent i it has observed interacting. Image scores are initialised at 0 and constrained to the range $[-5, 5]$. Each round m pairs of agents are randomly chosen from a population of n agents, with one being designated as the donor and the other as the recipient. If the donor's strategy is less than or equal to its perception of the image score of the recipient, $k_{donor} \leq I_{recipient}$, then it confers a benefit b on the recipient at a cost c to itself ($b = 1, c = 0.1$).

An agent assumes an image score of 0 if it has no data on the recipient. If the donor donates (cooperates), then observers of the interaction increment their perception of the donor's image score (the recipient's image score remains the same). If the donor does not cooperate, the perceived image score of the donor as held by the observers is decremented. An agent's strategy k_i thus represents the degree of selfishness of potential interaction partners that the agent is willing to cooperate with.

N&S consider both *complete* and *partial* observability of interactions. In the partial observability setting, N&S allow 10 randomly chosen agents to observe each interaction. Our formulation differs from N&S in that we consider an underlying network topology that constrains interactions. We model partial observability using a parameter, o, in the range $[0, 1]$, as the probability of each neighbour observing an interaction. If N_i denotes the set of neighbours for a given agent i, then, on average, $o \times |N_{donor} \cup N_{recipient}|$ observers are randomly selected for each interaction. We assume that interactions are observed without noise. The parameters $n = 100$ and $o = 0.1$ with a completely connected topology are equivalent to the original setup of N&S. Observability, in static connection topologies as investigated in this paper, can be viewed as a simple abstract model of resource constraints, or hardware or communications failure.

After m interactions offspring are generated proportionally to agents' final payoffs. If agent i has *fitness* f_i, where f_i is equal to its net benefit (the sum of the costs incurred and benefits received), then F is the net *population benefit* such that $F = \sum_{i=0}^{n} f_i$. An agent i will produce $\frac{n \times f_i}{F}$ offspring. The offspring's strategy is an exact copy of its parent's, with a small probability μ of mutation, such that the strategy is set to a random value (we adopt the value of $\mu = 0.001$ used by N&S). N&S found that strategies do not converge to a single value except for $o = 1$ and $\mu = 0$, but instead go through cycles as selfish agents become dominated by conditionally cooperative agents (called *discriminators* by N&S), who only help other cooperative individuals. These agents are then superseded by unconditionally cooperative agents (called *altruists*), who are subsequently invaded by selfish agents (called *defectors*).

N&S characterise the strategy space as: $k \leq 0$ denotes cooperation, since agents will interact with most other agents, and $k > 0$ denotes defection (called *selfish* by N&S). We further divide the cooperative strategy space into *unconditionally cooperative* ($-5 \leq k \leq -2$) and *conditionally cooperative* ($-2 < k \leq 0$). We describe interaction choices as follows. Interactions in which an agent cooperated based on its *perceived* image score of the recipient, when it should have defected based on the *actual* image score, or vice-versa, are *misclassified interactions*. An interaction is called *incorrect cooperation* if an agent cooperates when it should have defected, and an *incorrect defection* is where an agent defects when it should have cooperated. Incorrect defections are undesirable since they reduce the donor's image score, leading to fewer subsequent donations to the donor. Incorrect cooperations are undesirable since they allow selfish agents to gain higher payoff, and increase their chance of reproduction. In our simulations we maintain the absolute value of an agent's image score, to allow calculation of misclassified interactions, and this includes any incorrect cooperations or defections, i.e. it is the result of an agent's actual actions rather than how they should have acted given complete information.

3.2 Gossiping mechanism

Our simple gossip mechanism spreads perceived image scores as follows: each agent maintains a queue of received gossips, which are processed in a separate gossip phase. After an interaction, each observer starts a gossip with probability *ogp* (observer gossip probability) by sending a gossip packet to a randomly chosen neighbour. The probability of any given agent starting a gossip thus depends both on o, the probability it is chosen as an observer, and on *ogp*, the probability that an observer starts a gossip. Each gossip packet contains the image score of the donor, as perceived by the gossip starter, the unique ID of the donor, the unique ID of the gossip starter, and a time to live (TTL).

Every *gossipRate* interactions, there is a gossip phase. Each agent in turn updates their image scores for each agent that they have received gossips about, and if $TTL > 0$ propagates the gossip with $TTL_{t+1} = TTL_t - 1$ to a single randomly chosen neighbour that does not yet have the gossip. It is assumed that an agent can check if a neighbour has received a gossip already.

We propose four update rules for incorporating received gossip information.

1. *Aggregate Average (AA):* The agent replaces its perceived image score for agent i with the average of its previous perceived score for i and the values contained in all the received gossips concerning i.
2. *Average Replace (AR):* The agent replaces its perceived image score for agent i with the average of the values contained in all received gossips concerning i.
3. *Majority Replace (MR):* The agent replaces its perceived image score for i with the median value contained in all received gossips concerning i. As noted above, this is thought to approximate how humans process gossip [22].
4. *Most Recent (MRec):* The agent replaces its perceived image score for i with the most recent value received concerning i.

In practical terms, an agent using one of these update rules is updating it's *perception* of the recipient's probable strategy (i.e. selfish or not), based on the average or majority perception (for AA, AR, and MR update rules), or simply using the most recently received opinion with which to make a decision (i.e. using the MRec rule).

4 Experimental setup

We model two primary situations in which incomplete information may undermine the efficacy of reputation: (i) when there is a very low probability of having observed any interactions, such as when first entering a system, and (ii) when there is a very low probability of observing a complete set of interactions. We model (i) using a low ratio of interaction rate to population size, and (ii) using a very high ratio of interaction rate to population size. N&S used parameters of $n = \{20, 50, 100\}$ and $m = \{125, 200, 300, 500, 1000\}$, sufficient for (i) but limited for (ii). To investigate (ii), we consider $m = \{1000, 5000, 10000, 20000, 50000\}$ for $n = 100$ (i.e. a maximum ratio of $m/n = 500$). We use $o = 0.1, \mu = 0.001, b = 1$ and $c = 0.1$. Unless otherwise stated, we use an observer gossip probability $ogp = 1.0$ and $gossipRate = 1$. Since the diameter of the networks in our simulations is typically less than 5 we use a TTL of 5. We also performed simulations scaling the population to $n = 1000$ to test the effects of group size.

We situate agents on a variety of network structures. We replicate N&S' completely connected topology, and implement random (such that each pair of nodes is connected with probability p), scale-free and small-world synthetic networks[1]. Scale-free networks are generated using the Eppstein and Wang [7] algorithm and small-world networks using Kleinberg's generation algorithm [14]. We further use 8 networks sampled from the Enron email dataset and the arXiv general relativity section collaboration network[2] to corroborate our results on networks that are structurally closer to those found in the real-world. We sampled four sub-graphs of around 1,000 vertices from each dataset using a breadth-first search (BFS) approach. We acknowledge that BFS can introduce structural biases into the sampled sub-networks [10]. Many other sampling techniques (e.g. Snowball Sampling or Forest Fire) also introduce biases [15], although unbiased techniques have been proposed [10]. In future work we aim to investigate the use of an unbiased sampling approach.

Our investigation focused on two main metrics: the strategy distribution for the population and the number of misclassified interactions. The results given are averaged over 20 runs for each configuration (giving a standard deviation ranging from 1–14%). We used $t = 10000$ generations of evolution. Due to the cyclic nature of strategies identified by N&S [17], we present results averaged over the course of the simulation, rather than taking the final state.

5 Results and discussion

Figure 1a shows the strategy distribution with $n = 100$ and $m = \{125, 300\}$ (i.e. a ratio of $m/n = \{1.25, 3\}$), across a variety of synthetic networks. Figure 1b plots results from varying m on a completely-connected topology and Table 1b summarises the results for each synthetic topology class. Finally, Table 1a shows the results from simulations with $n = 1000$ on scale-free and small-world networks.

These results show significant levels of selfishness across a variety of configurations. We conclude that there are three primary influences on levels of selfishness:

[1] Generated using the Java Universal Network/Graph Framework http://jung.sourceforge.net/

[2] The Enron and arXiv datasets are taken from http://snap.stanford.edu/data/

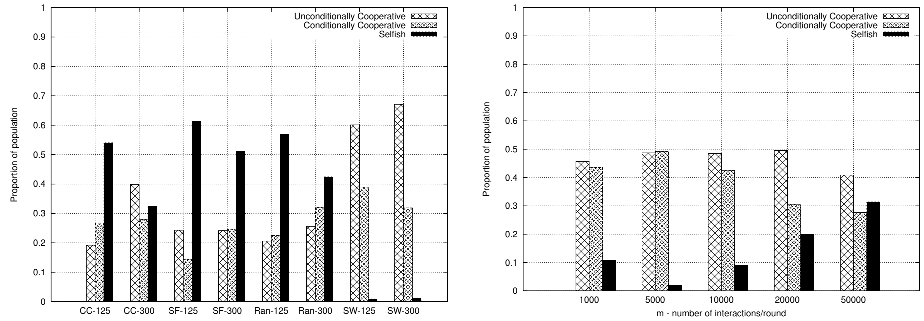

(a) (b)

Fig. 1: (a) Comparison of population strategy distribution for Completely Connected (CC), Scale-Free (SF), Random (Ran), and Small World (SW) topologies, using $m = \{125, 300\}$ and $n = 100$, (b) the population strategy distribution in a completely connected topology, and varying m.

1. **Underlying network topology**
   Selfishness is most evident in scale-free networks (structurally closest to the real-world), and small-world networks are particularly supportive of cooperative behaviour. Small-world networks have low geodesic path lengths and high clustering, implying a higher probability of connection between observers of an interaction and potential interaction pairs[3].
2. **Interaction rate**
   At very low rates ($m = 125$), there is not time for indirect reciprocity to take hold before selfishness increases (i.e. image scoring suffers from a *cold start* problem). As m increases selfishness is slightly reduced (down to 1.07% at $m = 1000$), but again rises as we approach $m = 50000$ (up to 31.4%). At low and high values of m, there is increased probability of agents having insufficient information to make an accurate assessments. As a result, the efficacy of image scoring is drastically reduced, and selfishness rises. These represent vulnerable configurations for reputation mechanisms.
3. **Population strategy distribution**
   A population with an equal strategy distribution increases the effect of incomplete information by increasing the uncertainty about a potential partner's strategy, making a decision based on incomplete information more likely to be incorrect.

The last two conclusions are corroborated by Figures 2a and 2b, which show the proportions of incorrect decisions for simulation runs shown in Figure 1b. Figure 2a shows a significant increase in the number of misclassified interactions as m (and population selfishness) increases, up to 3.5%, or an average of 1750 per round. Indirect reciprocity promotes cooperation through feedback effects, such that a cooperative action induces further cooperative actions, and so forth.

---

[3] Recall that although observers may be connected with the recipient of an interaction, they only update the score of the donor. For the observation to be of use, the observer must then also interact with the donor.

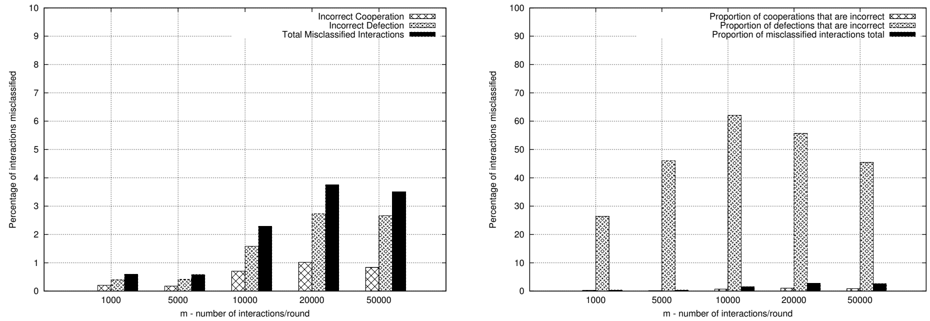

(a) (b)

Fig. 2: (a) The percentage of all interactions that were misclassified (y-axis shows 0–10% for clarity) in Figure 1b ($n = 100$), for differing values of m, (b) number of interactions misclassified as proportion of interaction type for simulation runs in Figure 1b.

		Proportion of population		
Network	Parameter	UC	CC	S
Eppstein	1000 edges	0.33	0.20	0.46
Eppstein	5000 edges	0.36	0.23	0.39
Eppstein	10000 edges	0.41	0.25	0.32
Kleinberg	CE1	0.42	0.26	0.31
Kleinberg	CE5	0.43	0.20	0.36
Kleinberg	CE10	0.34	0.30	0.35

(a)

m	CC		SF		SW		Ran	
	S	IP	S	IP	S	IP	S	IP
1000	0.13	3.71	0.91	1.2	0.01	0.1	0.03	2.0
5000	0.10	3.76	0.005	1.8	0.02	0.3	0.01	0.4
10000	0.21	5.71	0.004	1.1	0.03	0.4	0.02	0.3
20000	0.30	6.94	0.006	1.1	0.06	0.6	0.02	0.2
50000	0.39	5.93	0.01	2.2	0.11	0.6	0.04	0.2

(b)

Table 1: (a) Strategy distribution for a selection of scale-free and small-world networks. UC is Unconditionally Cooperative, CC is Conditionally Cooperative, and S is Selfish. CE is Clustering Exponent, $n = 1000$ and $m = 1000$ unless otherwise stated. (b) Selfish proportion of population (S) and Percentage of Incorrect interactions (IP) for Completely Connected (CC), Scale Free (SF), Small World (SW) and Random (Ran) networks while varying m, with $n = 100$. All other parameters as Figure 1a for both.

A decision to defect when an agent should have cooperated can thus significantly impact which strategies are reproduced. Figure 2b shows the proportion of interactions that were misclassified *as a proportion of that interaction type*. While only a negligible number of decisions to cooperate were incorrect, the number of incorrect defections was large, peaking at 62% for $m = 10000$. This further demonstrates the effect of population strategy diversity, since as we move past $m = 10000$ the proportion of incorrect decisions to defect reduces — there are more selfish agents, so a decision to defect is more likely to be correct. Reducing uncertainty over a recipient's strategy with supplementary mechanisms is therefore key to aiding reputation mechanism efficacy.

Scaling up the population to 1,000 agents demonstrates a smoothing effect. The influence of incomplete information is slightly reduced, but the populations are more evenly distributed with selfishness remaining significant. Interestingly, the support that a small-world topology displayed for cooperative behaviour in a small population ($n = 100$)is no longer present (Table 1a), and selfishness levels are similar to scale-free networks. Increasing the number of edges in scale-free networks slightly reduces the level of selfishness, corroborating our hypothesis regarding visibility of agent interactions.

5.1 Supporting reputation with gossiping

In this section, we present results from implementing gossiping with our four aggregation rules. Table 2 compares levels of selfishness in the population for the same configuration as Figure 1a, except that agents gossip and use the Average Replace update rule. Figure 3a shows the strategy distribution using the different update rules on a scale-free topology with $m = 1000$. Finally, Table 3b shows the results from using gossiping on the real-world networks. From these results, we can conclude the following:

1. **Gossiping significantly reduces levels of selfishness.**
   On average, gossiping reduces levels of selfishness by around 10% in the synthetic networks and around 18% in the real-world networks.
2. **There is no clear relationship between the number of interactions and the reduction in selfishness, whereas there is a clear link between network structure and gossiping efficacy.**
   The real-world networks and scale-free synthetic networks in particular show significant reductions in selfish behaviour. Given that real-world domains often have scale-free properties, these results suggest that gossiping can be practically applied. Random networks are less conducive to gossiping than other network classes, which may be a consequence of their reduced clustering. As argued above, clustering increases the probability of observations being of use, and gossips are simply a substitute for direct observation.
3. **All update rules except Aggregate Average show a statistically significant decrease in selfishness ($\alpha = 0.05$).**
   In the synthetic networks Aggregate Average performs worse than the other update rules, which perform fairly equally. In the real-world networks Most Recent performs consistently and gives the largest reduction in selfish behaviour, while the other update rules occasionally increase selfish behaviour.

On average across the 4 update rules, 331.7 million gossips were started, with 1.436 billion gossip packets sent over 10 million interactions, or 143 packets per interaction. Agents adopted a new image score for a given individual 496.4 million times. On average, a single gossip causes 1.50 image score changes. Aggregate Average is the only rule to incorporate the agent's current image score of the gossip subject, whereas the other three rules only consider the received gossips. A gossip using the Aggregate Average rule causes, on average, 1.09 image score changes, which is less than the other rules (1.67 for Average Replace, 1.63 for Most Recent, and 1.60 for Majority Replace). Update rules that do not incorporate the *current* perception of the subject's image score perform better. That Most Recent performs as well as the others suggests that many of the updates are for when agents have no information (i.e. they assume an image score of 0), and the gossip provides initial data for decisions. Aggregate Average incorporates the assumption of an image score of 0, biasing the resultant value. These

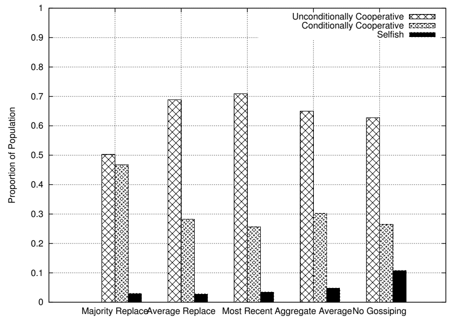

(a)

	Selfishness					
Topology	$ogp = 0$	$ogp = 1$				max. Diff.
		AA	AR	MR	MRec	
Enron-1	0.29	0.33	0.20	0.33	0.04	**0.25**
Enron-2	0.31	0.35	0.22	0.33	0.19	**0.12**
Enron-3	0.33	0.38	0.25	0.18	0.21	**0.15**
Enron-4	0.39	0.20	0.25	0.21	0.16	**0.23**
arXiv-1	0.33	0.39	0.21	0.19	0.15	**0.18**
arXiv-2	0.34	0.50	0.10	0.30	0.15	**0.24**
arXiv-3	0.34	0.40	0.19	0.35	0.22	**0.15**
arXiv-4	0.27	0.21	0.16	0.33	0.19	**0.11**

(b)

Fig. 3: (a) Strategy distribution using gossiping with $m = 1000$, on a scale-free network with 1000 edges, varying the update rule, and (b) the effect of gossiping on selfishness in real-world networks with each update rule (Aggregate Average (AA), Average Replace (AR), Majority Replace (MR), Most Recent (Mrec)) ($ogp = 1$). $n = 1000, m = 1000$, and all other parameters as Figure 1a.

		Selfishness		
Topology	m	$ogp = 0$	$ogp = 1$	Diff.
Completely-Connected	125	0.540	0.418	**0.122**
Completely-Connected	300	0.324	0.221	**0.103**
Scale-free	125	0.613	0.479	**0.134**
Scale-free	300	0.512	0.330	**0.182**
Random	125	0.569	0.527	**0.042**
Random	300	0.424	0.256	**0.168**

Table 2: The average proportion of selfish agents in the population for the runs in Figure 1a ($ogp = 0$) compared with runs using the same configuration except that agents gossip using the Average Replace update rule ($ogp = 1$).

results suggest that gossiping is a useful mechanism by which new entrants to a system can start interacting quickly without having to observe the population to gain sufficient information.

In the real-world networks, Aggregate Average still performs poorly, but Most Recent gives the most consistently beneficial results. These results are given for $m = 1000$, meaning that agents are likely to have very little or no information on potential interaction partners. The Most Recent rule is equivalent to allowing each of the gossip recipients to act as an observer of the interaction being gossiped about, and thus reduces the number of interactions necessary for indirect reciprocity to take hold. This corroborates our conclusion that gossiping is a particularly useful supplement to reputation for new entrants to a system, or in systems characterised by high levels of population churn. While the benefit of gossiping in the real-world networks is generally larger than in the synthetic networks, the introduction of gossiping occasionally results in an increase in selfishness (particularly with Aggregate Average, but never with Most Recent). This requires further investigation, and these results imply careful consideration must be given to how agents incorporate information attained through gossiping.

6 Summary and Further Work

In this paper, we have shown that (i) incomplete information can significantly undermine lightweight reputation mechanisms, with up to 62% of defection actions (in the completely connected topology scenario) taken incorrectly, (ii) the underlying network topology has a significant influence on levels of selfishness in the population, (iii) gossiping can reduce levels of selfishness by up to 25%, with the biggest gains found on real-world topologies, and (iv) using the most recently gossiped information about a potential partner results in the most consistent benefits, suggesting that gossiping may be particularly useful for agents first entering a system. In future work, we intend to extend our investigation on real-world networks (using an unbiased sampling algorithm) and determine the topological properties that are most conducive to supporting cooperative behaviour.

References

1. Y. Bachrach, A. Parnes, A. D. Procaccia, and J. S. Rosenschein. Gossip-based aggregation of trust in decentralized reputation systems. *Autonomous Agents and Multi-Agent Systems*, 19(2):153–172, December 2008.
2. B. Bako, I. Rikanovic, F. Kargl, and E. Schoch. Adaptive Topology Based Gossiping in VANETs Using Position Information. *Proceedings of the 3rd international conference on Mobile ad-hoc and sensor networks*, pages 66–78, 2007.
3. S. Buchegger. Coping with Misbehavior in Mobile Ad-hoc Networks. PhD Thesis, *Ecole Polytechnique Federale de Lausanne*. 2005.
4. C. Castelfranchi and R. Falcone. Principles of trust for MAS: Cognitive anatomy, social importance, and quantification. In *Proceedings of the 3rd International Conference on Multi Agent Systems*, pages 72–79, 1998.
5. M. Cohen,R. Riolo, and R. Axelrod. The Role of Social Structure in the Maintenance of Cooperative Regimes *Rationality and Society*, 13(1):5–32, 2001.
6. P. Costa, D. Gavidia, B. Koldehofe, H. Miranda, M. Musolesi, and O. Riva. When cars start gossiping. *Proceedings of the 6th workshop on Middleware for network eccentric and mobile applications*, pages 1–4, 2008.
7. D. Eppstein and J. Wang. A steady state model for graph power laws. In *Proceedings of the 2nd International Workshop on Web Dynamics*, 2002.
8. Y. Fernandess and D. Malkhi. On collaborative content distribution using multi-message gossip. *Journal of Parallel and Distributed Computing*, 67:1232–1239, 2007.
9. A. Frieze and G. Grimmett. Shortest-path problem for graphs with random arc-lengths. *Discrete Applied Mathematics*, 10(1):57–77, 1985.
10. M. Gjoka, M. Kurant, C. Butts, and A. Markopoulou. Walking in facebook: A case study of unbiased sampling of osns. In *Proceedings of the 29th Conference on Information Communications*, pages 2498–2506, 2010.
11. R. Glinton, P. Scerri, and K. Sycara. Exploiting scale invariant dynamics for efficient information propagation in large teams In *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems*, pages 21–30, 2010.
12. T. D. Huynh, N. R. Jennings, and N. R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.

13. D. Kempe, A. Dobra, and J. Gehrke. Gossip-based computation of aggregate information. In *Proceedings of the 44th Annual IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2003.
14. J. Kleinberg. Navigation in a small world. *Nature*, 406(3):845, 2000.
15. M. Kurant, A. Markopoulou, and P. Thiran. On the bias of BFS (Breadth First Search). *Proceedings of the 22nd International Teletraffic Congress*, 2010.
16. J. Mundinger and J-Y. Le Boudec. Reputation in self-organized communication systems and beyond. In *Proceedings from the 2006 workshop on Interdisciplinary systems approach in performance evaluation and design of computer & communications sytems*, page 3, 2006.
17. M. Nowak and K. Sigmund. The dynamics of indirect reciprocity. *Journal of Theoretical Biology*, 194(4):561–74, 1998.
18. M. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437(7063):1291–8, 2005.
19. A. A. Pirzada and C. Mcdonald. Trust Establishment In Pure Ad-hoc Networks. *Wireless Personal Communications*, 37(1-2):139–168, 2006.
20. S. Ramchurn, N. Jennings, Carles Sierra, and Lluis Godo. Devising a trust model for multi-agent interactions using confidence and reputation. *Applied Artificial Intelligence*, 18(9-10):833–852, 2004.
21. J. Sabater, M. Paolucci, and R. Conte. Repage: Reputation and image among limited autonomous partners. *Journal of Artificial Societies and Social Simulation*, 9(2):3, 2006.
22. R. D. Sommerfeld, H.-J. Krambeck, and M. Milinski. Multiple gossip statements and their effect on reputation and trustworthiness. *Proceedings of the The Royal Society B*, 275(1650):2529–36, 2008.
23. R. D. Sommerfeld, H.-J. Krambeck, D. Semmann, and M. Milinski. Gossip as an alternative for direct observation in games of indirect reciprocity. *Proceedings of the National Academy of Sciences of the United States of America*, 104(44):17435–40, 2007.
24. A. Sridharan and J. Nastos, 'Statistical Behavior of Embeddedness and Communities of Overlapping Cliques in Online Social Networks', *Distribution*, in *Proceedings of the 30th Internal Conference on Communication Technology*, pp546–550, (2011).
25. H.P. Young. The economics of convention. *The Journal of Economic Perspectives*, 10(2):105–122, 1996.
26. R. Zhou and K. Hwang. Gossip-based Reputation Aggregation for Unstructured Peer-to-Peer Networks. In *Proceedings of the International Parallel and Distributed Processing Symposium*, pages 1–10, 2007.