

Original citation:

Franks, H. P. W., Griffiths, Nathan and Jhumka, Arshad (2010) Image scoring in ad-hoc networks : an investigation on realistic settings. In: 8th European Workshop on Multi-Agent Systems, Paris, France, 16-17 Dec 2010

Permanent WRAP url:

<http://wrap.warwick.ac.uk/47426>

Copyright and reuse:


The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

A note on versions:

The version presented in WRAP is the published version or, version of record, and may be cited as it appears here.

For more information, please contact the WRAP Team at: publications@warwick.ac.uk



<http://wrap.warwick.ac.uk/>

Image scoring in ad-hoc networks: An investigation on realistic settings

Henry Franks, Nathan Griffiths, and Arshad Jhumka

Department of Computer Science, University of Warwick, CV4 7AL, UK

Abstract. Encouraging cooperation in distributed Multi-Agent Systems (MAS) remains an open problem. Emergent application domains such as Mobile Ad-hoc Networks (MANETs) are characterised by constraints including sparse connectivity and a lack of direct interaction history. *Image scoring*, a simple model of reputation proposed by Nowak and Sigmund, exhibits low space and time complexity and promotes cooperation through indirect reciprocity, in which an agent can expect cooperation in the future without repeat interactions with the same partners. The low overheads of image scoring make it a promising technique for ad-hoc networking domains. However, the original investigation of Nowak and Sigmund is limited in that it (i) used a simple idealised setting, (ii) did not consider the effects of incomplete information on the mechanism’s efficacy, and (iii) did not consider the impact of the network topology connecting agents. We address these limitations by investigating more realistic values for the number of interactions agents engage in, and show that incomplete information can cause significant errors in decision making. As the proportion of incorrect decisions rises, the efficacy of image scoring falls and selfishness becomes more dominant. We evaluate image scoring on three different connection topologies: (i) completely connected, which closely approximates Nowak and Sigmund’s original setup, (ii) random, with each pair of nodes connected with a constant probability, and (iii) scale-free, which is known to model a number of real world environments including MANETs.

1 Introduction

The emergence of cooperation remains an open problem in the agent community. Although significant progress has been made in a number of domains, few proposed techniques are fully applicable given the challenges of emergent domains such as MANETs. These domains are characterised by unique constraints such as low computational capacity, high agent turnover, lack of centralised authority, sparse connectivity and lack of single ownership, and therefore require low-overhead, distributed solutions. Investigating such mechanisms for promoting cooperation has become a major theme of research in open MAS.

Nowak and Sigmund’s [16] image scoring mechanism biases partner selection towards cooperative individuals while requiring low space and time overheads. Agents maintain a subjectively perceived image score of other agents, based on

observation of the interactions those agents participated in. This image score is then used to aid decision making in future interactions.

Nowak and Sigmund (N&S) found that cooperation does emerge, but is often cyclic as non-cooperative agents invade populations of unconditionally cooperative agents and gain higher payoffs, causing the population to be subsequently dominated by conditionally cooperative agents, who are then superseded by unconditionally cooperative agents. Agents in the setup used by N&S are randomly chosen and paired from the entire population for interactions, with the total number of interactions per round (m) being at most one order of magnitude larger than the number of agents in the population (n).

Our paper is organised as follows. We discuss in Section 2, given our illustrative domains, why the assumptions made by N&S are unrealistic. In Section 3, we discuss our investigation of image scoring on completely connected, random, and scale-free connecting topologies, and with a higher ratio of interactions to agents (m/n). Our view is that many domains have a high ratio of m/n . Taking data routing in MANETs as an example, an interaction might represent the transfer of a packet between nodes, such that m may be several orders of magnitude larger than the number of agents. Given the modification of this ratio, agents observe a smaller proportion of the total interactions. The distance between an agent’s subjective perception of an image score and its absolute value is thus increased, leading to incorrect choices in actions. We show the effect of these errors on emergent cooperation and explore the efficacy of image scoring given common social network topologies in Section 4. Finally, we discuss the implications of our work on the wider problem of incomplete information in open MAS, and identify further work, in Section 5.

2 Background

Several cooperation mechanisms have been proposed for MAS, typically incorporating notions of trust or reputation. Many such mechanisms are effective in constrained domains, but rarely fully address the challenges posed by ad-hoc or mobile networks. In this paper we focus on *image scoring*, which appears suited to such domains, but has not to our knowledge been investigated with these constraints in mind. In this section we introduce the image scoring approach, give an overview of the motivating domain, and discuss how the social connections between individuals can be represented by a network topology.

2.1 Image scoring

N&S introduced and extensively investigated *image scoring*, a simple instantiation of reputation modelling indirect reciprocity, in which cooperation emerges without requiring subsequent interactions with the same individuals [16, 17]. This property is key to its suitability in open decentralised systems. In the formulation used by N&S, each round consists of m pairings of agents, such that two agents are randomly selected from the population to act as *donor* and *recipient*

in a pairing, or *interaction*. The donor can cooperate, conferring benefit b on the recipient at personal cost c (where $b > c$), or choose non-cooperation, with no change in payoffs. For the purposes of our discussion, we denote non-cooperation as *defection*, although non-cooperation is not explicitly malicious. Agents maintain an *image score* of others, in the range $[-5, 5]$. A number of agents will *observe* each interaction and adjust their perceived *image score* of the donor, I_{donor} , accordingly. Cooperation increases the perceived image score by 1 and defection decreases it by 1. Each agent has a *strategy*, k , which is an integer in the range $[-5, 6]$. A donor will cooperate with a recipient if its strategy, k , is less than or equal to the recipient's image score, namely if $k \leq I_{recipient}$. Thus, $k = 6$ implies the donor will never cooperate, regardless of another's image score, whereas $k = -5$ implies unconditional cooperation. If an agent has no observations of the potential partner, it assumes an image score of 0. After a round, the most successful strategies are propagated to the next round proportionally to the net benefit gained by agents using them (see Section 3 for more detail). 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 (also termed *selfish* by N&S). We further divide the cooperative strategy space into *unconditionally cooperative* ($-5 \leq k \leq -2$), *conditionally cooperative* ($-2 < k \leq 0$).

N&S observe cooperation, using a population size of $n = 100$ and $m = 125$ interactions per round. With these settings, each agent is chosen for an interaction (as donor or recipient) 2.5 times on average per round. Two variations of observability are considered by N&S: *full observability* where *all* agents observe an interaction, and *partial observability* where a proportion of agents observe. N&S assert that as m increases cooperation becomes more likely with full observability where image scores are universally known. Given partial observability (N&S consider 10 randomly selected observers per interaction), increasing m increases the likelihood that agents have not witnessed all interactions of a given agent, since only a proportion of agents observe each interaction. This leads to a disparity between the perceived image score and its absolute value.

2.2 Motivating Application Domains

MANETs are self-organising wireless networks characterised by lack of central authority and high turnover, with limited computation, connectivity, memory and energy resources. Vehicular Ad-hoc Networks (VANETs) are a subset of these in which the majority of nodes are vehicles, imposing additional constraints of intermittent, sparse connectivity and a dynamic connecting topology (although computational capacity and energy are less restricted). We are concerned with features of these domains (such as complex connecting topologies) but we do not attempt to model a realistic environment. These features limit resources, constrain connectivity, and lead to *incomplete information* about others due to partial observability and sparse connectivity.

Research on cooperation in VANETs has focused on network security and routing [3, 9, 13] and content distribution [22]. Since vehicles have intermittent

contact¹ and few repeat interactions, techniques requiring a history of direct interactions, including many traditional trust and reputation mechanisms, are not appropriate. Trust and reputation have been used to address routing problems in MANETs [2, 11, 19]. Several mechanisms show encouraging results but require historical interaction data that may not always be available. Since image scoring has less reliance on historical data it is potentially more suited to the constraints of VANETs, and so is our focus in this paper.

The connection properties of these domains make it challenging to achieve cooperation. MANETs exhibit scale-free properties in their connections [20], and VANETs exhibit extremes of sparse connectivity (e.g. a free-flowing motorway) and highly clustering (e.g. a traffic jam). Typical applications, such as data sharing, imply that the number of interactions between agents is significantly larger than the population size. N&S do not consider such a setting, and in this paper we investigate the performance of image scoring in this context.

We aim to explore (i) the efficacy of image scoring given assumptions representing features of these domains, (ii) the effect of incomplete information, and (iii) the characteristics of network topology that affect the levels of emergent cooperation achieved through image scoring.

2.3 Topologies

The populations considered by N&S do not have any form of connecting topology, and donor-recipient pairs are randomly selected from the entire population. In realistic application domains, agent populations are situated on some connecting topology that restricts agents' observations and interactions. The setup used by N&S can be viewed as a completely connected topology (where each agent is directly connected to each other agent), with interaction partners and observers chosen from the neighbour-set of an agent. This is an idealised topology that does not accurately reflect the networks that underlie real-world systems.

We investigate image scoring using three connection topologies: a completely connected topology, a random network that has topological assumptions closer to realistic settings, and a scale-free topology which is known to model many real-world networks, including ad-hoc networks [20].

In random networks each pair of nodes is connected with a probability p [8]. Completely connected topologies are idealised in assuming that an agent is connected to every other agent, whereas a random topology allows investigation of situations where agents are only connected to a small proportion of the rest of the population, as is the case in many real world environments.

Scale-free networks are a subset of random networks in which the node degree d for each node is given by the probability distribution $P(d) \approx d^{-\gamma}$, where γ is some arbitrary constant². Many real-world networks have scale-free properties,

¹ Ott and Kutscher [18] estimate 50-60% packet loss for vehicles passing and communicating with a roadside access point.

² In networking literature, k is typically used to denote node degree. We use d to prevent confusion with use of k for agent strategy.

including the world wide web, the internet, social network graphs, collaboration networks of researchers, cellular networks and phone call pattern networks [1]. These properties make scale-free networks suitable for use when investigating many domains of multi-agent systems, including ad-hoc networks [20].

3 Experimental Setup

We reproduced the original setup used by N&S as follows: each agent is associated with a strategy k , chosen uniformly at random in the range $[-5, 6]$. Each agent maintains a matrix of image scores calculated from observing other agents interact and its own image score, which allows us to calculate the number of misclassified interactions (see below). Image scores are initialised at 0 and constrained to the range $[-5, 5]$. A number of pairs, m , of agents are randomly chosen from a population of n agents each round, with one agent 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, with $b > c$ (we adopt the values of $b = 1, c = 0.1$ used by N&S). An agent assumes an image score of 0 if it has no data on the recipient. If the donor acts altruistically and donates (cooperates), then its image score increases by one, and the observers of that interaction increase their perception of the donor's image score by one (the recipient's image score remains the same). If the donor does not cooperate, then its image score falls by one, and observers of the choice not to cooperate decrease their perceived image score of the donor by one.

N&S consider both *complete* and *partial* observability of interactions. In the partial observability settings, N&S allow 10 agents, chosen at random, to observe each interaction. To allow investigation of the effects of changing the number of observers, we define an observability parameter, o , in the range $[0, 1]$, as the probability of observing an interaction of a neighbour. Observers are selected at random from the union of the neighbour-set of the donor and recipient for each interaction. If N_i denotes the set of neighbours for a given agent i , then the average number of observers of an interaction is given by $o \times |N_{donor} \cup N_{recipient}|$. Thus, given $n = 100$, an observability of $o = 0.1$ is equivalent (given a completely connected topology) to the original setup of N&S. Observability, in the static (i.e. time invariant) connection topologies investigated in this paper, can be viewed as a simple abstract model of constraints such as hardware or communications failure, or the intermittent connectivity characteristic of VANETs.

After m interactions have been performed, offspring are generated in proportion to an agent's final payoff. N&S do not give more precise details and so in our simulations we use the following mechanism. If agent a_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 will produce $\frac{n \times f_i}{F}$ offspring. The strategy of the offspring is an exact copy of the parent strategy, 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 (termed *discriminators* by N&S), who only help other cooperative individuals. These agents are then superseded by unconditionally cooperative agents (also termed *altruists* by N&S), who are subsequently invaded by selfish agents (termed *defectors* by N&S).

N&S used parameters of $n = \{20, 50, 100\}$ and $m = \{125, 200, 300, 500, 1000\}$, which are unrealistic in many instantiations of open MAS. Consider a typical data routing problem in MANETs: a cooperative interaction represents successful packet transmission between two agents. The number of interactions is thus likely to be large, such that m is much greater than n (i.e. the ratio of m/n is significantly greater than 10, the maximum ratio considered by N&S). To investigate the effects of scaling m and the ratio m/n , we simulated $m = 1000, 5000, 10000, 20000, 50000$ for $n = 100$ (i.e. a maximum ratio of $m/n = 500$). It would be useful to also investigate the effects of scaling n for various values of m/n , but this is beyond the scope of this paper, and will be considered in future work.

Along with the characterization of the strategy space identified above (Section 2.1), 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 termed *incorrect cooperation* if an agent cooperates when it should have defected. An *incorrect defection* is an interaction in which an agent defects (i.e. does not donate to the recipient) when it should have cooperated. The number of misclassified interactions is the sum of the incorrect cooperations and incorrect defections. Incorrect defections are detrimental to the society since they cause the donor’s image score to fall, leading to fewer agents donating to the donor in the future, thus undermining the levels of indirect reciprocity in the society. Incorrect cooperations are undesirable since they allow selfish agents to gain higher payoff, and thus become more likely to have their strategies propagated. Note that the absolute value of an agent’s image score that is maintained (to allow calculation of misclassified interactions) includes any incorrect cooperations or defections that that agent has made — it is the result of an agent’s actual actions rather than how they should have acted given complete information.

4 Results and Discussion

Our investigation focused on two main metrics: the strategy distribution for the population and the numbers of misclassified interactions. We conducted four types of simulation: (i) a replication of the original setup described by N&S without a connecting topology, (ii) using a completely connected topology, (iii) using a random topology and (iv) using a scale-free topology. For each type, we investigated the effects of varying both m and topological parameters such as p in a random topology or the number of edges for a scale-free topology. The

results given are averaged over 10 runs for each parameter configuration, giving a standard deviation ranging from 1–14%. We used $t = 10000$ generations of evolution and, unless otherwise stated, $o = 0.1$.

4.1 Effect of the Number of Interactions

First we replicated the original setup of N&S, using $n = 100$ agents, $o = 0.1$ (i.e. an average of 10 agents observing each interaction), $b = 1$, $c = 0.1$ and $\mu = 0.001$, and varied the number of interactions per round. Figure 1(a) shows the proportion of the population within each strategy category averaged over all generations, while varying m . Due to the cyclic nature of strategies identified by N&S, analysing results at an arbitrarily chosen generation (e.g. the final generation of $t = 10000$) is unlikely to provide a representative view of the simulation. For $m = 1000$ we observe a highly cooperative society in which only 1.07% of agents have a strategy $k > 0$ (selfish). Out of all the interactions performed, 0.6% were misclassified (i.e. around 6 interactions were misclassified per round). As m increases, we see two effects: (i) the population becomes more selfish and (ii) the proportion of misclassified interactions rises. Figure 1(b) shows the percentage of interactions over the entire simulation that were misclassified for varying values of m . At $m = 50000$, the selfish proportion of the population has risen to 31.4%. The proportion of misclassified interactions has risen to 3.5%, or an average of 1751 per round. Figure 2(a) shows the average strategy, sampled every 100 generations, for a representative run with $m = 1000$. The two horizontal lines show the delineation of our strategy classes, with everything above the top line classified as selfish, everything inside the lines being conditionally cooperative, and everything below the bottom line corresponding to unconditionally cooperative. We see the cycles between strategy classes described by N&S, but with the population remaining mostly cooperative.

It is interesting to note that the proportion of misclassified interactions falls very slightly from $m = 20000$ to $m = 50000$, while the selfish proportion of agents still rises. As later results show, the resultant strategy distribution of the population is important in determining the number of misclassified interactions. Image scoring, through indirect reciprocity, induces a feedback effect in which cooperative actions cause more cooperative actions, and vice-versa for defection. In a cooperative society, the donor is more likely to cooperate and the recipient is more likely to have cooperated in the past. Thus it is more likely that that choice to be cooperative is correct, even if the donor has not observed many interactions of the recipient. The same is true, vice-versa, for defecting societies. When a society is mixed between cooperators and defectors however, uncertainty about a recipient’s strategy is higher and the level of incorrect choices rises. We witness this effect from $m = 20000$ to $m = 50000$, where the population becomes much more selfish (20% for $m = 20000$ to 31.4% for $m = 50000$) but the proportion of misclassified interactions is more stable (with a slight drop from 3.75% to 3.5%).

The ratio of m/n is clearly important in determining the level of cooperation emergent in the population. An interesting effect emerges when we look at the misclassification of cooperative and defective interactions separately. This




Fig. 1. Results showing, for differing values of m , $n = 100$, $o = 0.1$, $\mu = 0.001$, $b = 1$, $c = 0.1$, (a) the proportion of population displaying each strategy type and (b) the percentage of all interactions that were misclassified (y-axis shows 0–10% for clarity).




Fig. 2. (a) Average strategy over time, sampled every 100 generations, for $m = 1000$, for representative run from Figure 1, and (b) interactions misclassified as proportion of interaction type for simulation runs in Figure 1, that is, the proportion of cooperations that are incorrect is the number of interactions in which the donor donated incorrectly divided by the total number in which the donor donated (as opposed to the total number of interactions as denominator).

is illustrated in Figure 2(b), which shows the percentage of actions which were incorrect for all actions of that type (cooperation or defection), for varying values of m . At $m = 1000$, only 0.24% cooperative interactions are misclassified, meaning that of all the interactions in which the donor donated, only 0.24% were incorrect given the absolute image score of the recipient. However, 26.4% of defections were misclassified, meaning that in 26.4% of the interactions in which the donor did not donate it should have done so. Given a society in which approximately 90% of the agents are cooperative, this makes sense, since any interactions of the recipient which the donor did not observe are likely to be cooperative, and thus increase the actual image score of the recipient. The donor's perception of the image score is likely to be lower than its absolute value, and the agent chooses defection. In a society that is mostly selfish, this effect is reduced. Indeed we see this as m increases to 50,000, and the society becomes markedly more selfish, the proportion of defections that are misclassified rises (to a peak of 62%) and then falls again (for the reasons discussed above).

These results show two important relationships: the proportion of (i) misclassified interactions due to incomplete information and (ii) selfish strategies both increase as the ratio m/n increases, leading to a decrease in cooperation emerging in the society.

4.2 Effects of Network Topology

The above results are for an environment without a social network topology connecting agents. We also investigated the effects of incomplete information given three different topologies: completely connected, which closely approximates the setup above, random, and scale-free. Instead of pairs of agents being chosen randomly, the donor is now selected at random from the population and the recipient is chosen at random from the donor's neighbour set.

Figure 3(a) shows the strategy distributions obtained while varying m in a completely connected topology. Figure 3(b) shows the corresponding percentage of misclassified interactions, with the maximum proportion reaching 6.87% at $m = 20000$. The results are almost identical to Figure 1, supporting our claim that a completely connected topology closely approximates the original setup used by N&S. The proportions of misclassified interactions follow the same trend, with a slight drop as the diversity of the population strategy distribution falls at $m = 50000$. We found no statistically significant difference between the results for no topology and a completely connected topology using a two-tailed t-test.

Completely connected topologies are, however, highly unrealistic. Random topologies represent a useful middle ground between a completely connected topology and scale-free. In a random topology, each pair of agents is connected with a probability p . Given 100 agents, a completely connected graph will involve $\frac{n(n-1)}{2} = 4950$ edges. A connection probability $p = 0.1$, with 100 agents, will mean each agent is connected to, on average, 4.95 other agents³ (assuming

³ 4950 potential edges with connection probability $p = 0.1$ gives 495 edges for the entire population. Given 100 agents, this equals 4.95 edges per agent on average.

an undirected graph, as used in this investigation). By varying p , we can explore the effects of more sparse or more connected topologies. Given that sparse connectivity is a major feature of many open MAS, especially VANETs, it is useful to evaluate the effect this can have on emergent cooperation.

Figure 4(a) shows the effects of varying p , using $n = 100, m = 1000, o = 0.1$ and $t = 10000$. Observability is kept constant, and thus with $p = 0.1$ and $o = 0.1$, it is likely that only one agent will observe an interaction on average: assuming the donor and recipient neighbour sets are distinct, the total number of neighbours is around 10 (on average)⁴. While this assumption may not always hold (e.g. when the neighbour sets of the donor and recipient overlap), it provides a useful upper bound. Since fewer and fewer agents will observe an interaction as p falls, it becomes less beneficial to cooperate and selfishness becomes dominant in the society. At $p = 0.01$, selfishness is highly dominant: with so few edges, not enough agents will observe a cooperation to make it beneficial in the long run. As p rises from 0.001 to 0.1, cooperation becomes more viable, since enough agents will observe a cooperation to allow indirect reciprocity to have an effect. However, above $p = 0.1$, we see selfishness start to rise again. As p rises past this point, a smaller proportion of an agent's neighbours observe each interaction, and thus the effects of incomplete information become more pronounced. Figure 4(b) shows this effect, with a maximum percentage of misclassified interactions of 2.09% at $p = 0.5$. At $p = 0.01$, where most of the population is selfish, there are almost no misclassified interactions. As p rises past 0.1 a higher proportion are misclassified. It is evident that there is a minimum number of observers necessary for cooperation to emerge.

In the random topology, we require a minimum number of interactions (as noted by N&S) for cooperation to prevail, and this number is dependent on how many agents observe each interaction. Figure 6(a) shows the effects of varying m for a small value of $p = 0.1$. Each agent is connected to far fewer neighbours than in the completely connected topology, and we see that at $m = 1000$, which is enough for cooperation in completely connected topologies, selfish agents prevail. At $m = 5000$, we see the more familiar highly cooperative strategy distribution. Once again, as m rises after establishing cooperation, we see selfishness representing a slightly higher proportion of the strategies in the population. It is much less prevalent than in the completely connected topology, however, and the use of a random topology appears to aid the establishment of cooperation.

Comparing the random topology to the completely connected topology, we find that the level of selfishness is different with statistical significance. Calculated t-values for a two-tailed t-test range from 0.031 at $m = 1000$ to 6.37×10^{-20} at $m = 50000$. The statistical significance becomes increasingly strong as m rises, showing that random topologies significantly affect the operation of image scoring. The levels of cooperators are not consistently significantly different. We witness the same trend, that the proportions become more significant as m increases. At $m = 50000$, the levels of unconditional and conditional cooperators

⁴ Given 4.95 average neighbours per agent as above.




Fig. 3. (a) Strategy classifications and (b) levels of misclassified interactions, using a completely connected topology, with all other settings as Figure 1.

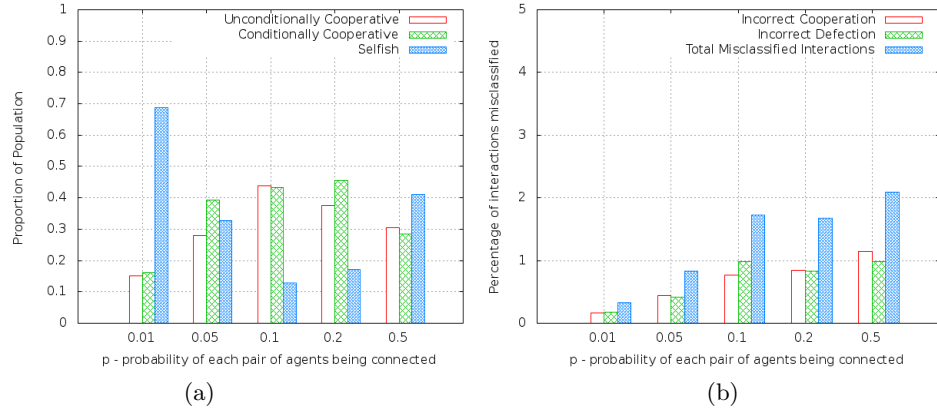


Fig. 4. (a) Strategy classifications and (b) levels of misclassified interactions, using a random topology, varying p , with $m = 1000$ and all other settings as Figure 1.

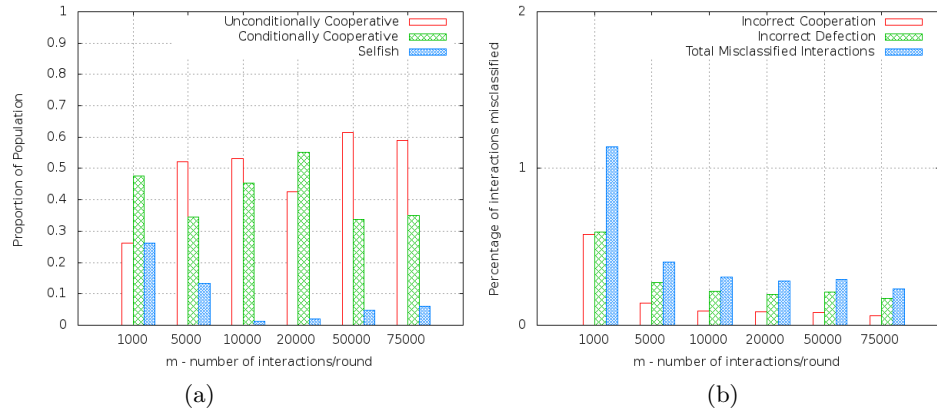


Fig. 5. (a) Strategy classifications and (b) levels of misclassified interactions, for scale-free topology, varying m , using 1000 edges in total, and with all other settings as Figure 1.

are significantly different from completely connected topologies, with t-values of 0.003 and 0.009 respectively.




Fig. 6. Strategy results for (a) random topology, varying m , $p = 0.1$, and (b) scale-free topology, varying total number of edges, $m = 1000$, and with all other settings as Figure 1.

Finally, we simulated image scoring using agents connected by a scale-free topology, as generated by Eppstein’s generating algorithm [7]. The algorithm is parameterised by r , the number of iterations, and the total numbers of edges in the network. As r tends to infinity, the node degree distribution tends towards a power-law. In all simulations, we use $r = 100000$.

Figure 6(b) shows the results of varying the number of edges in the population. We see a similar effect to random topologies at very low numbers, with selfishness being dominant. As the number of edges rises the selfish proportion of the population continues to fall, which differs from random topologies in which the selfish proportion rises after some threshold value of p is reached. Figure 5(a) shows similar effects for varying m (up to 75,000). We still require a certain number of interactions to allow cooperation to emerge, but once it has emerged the level of selfishness remains very low, and the levels of unconditional cooperation in the population are far higher than for the other topologies investigated. Scale-free topologies appear to aid image scoring’s robustness to selfish action. Figure 5(b) strengthens this view: the levels of misclassified interactions are very low and fall as m rises, with a maximum total of 1.13% at $m = 1000$. This is partially due to the fact that the population is very highly cooperative, and as before, it is harder for an agent to choose incorrectly. However, we find no statistically significant difference in the strategy distributions between random and scale-free topologies.

Scale-free topologies are known to have beneficial properties regarding information propagation and robustness to untargeted malicious action. For ex-

ample, Delgado [6] used a model of social convention emergence to show that complex (i.e. scale-free and/or small-world) networks are more efficient than regular graphs with the same average node degrees, and that scale-free networks are as efficient at spreading information as fully-connected graphs. Barabasi and Albert also noted the remarkable fault-tolerance of scale-free networks [1]. The robustness of scale-free networks is partially derived from their clustering: there are highly-internally connected groups with relatively few links to the rest of the population. In the context of our investigation, we hypothesise that this grouping effect allows image scoring to act with a much smaller average connectivity, since there will be many such groups in which agents are highly visible to other agents within that group. As discussed previously, visibility of agents is important for the efficacy of image scoring. We use *visibility* to denote the combination of observability and topological connectivity, since both influence how many agents might observe an interaction. Sen [20] demonstrated the existence of scale-free topological structure in mobile ad-hoc networks, and many other real-world networks are known to be scale-free [1]. The robustness of image scoring on scale-free networks is thus highly important, as it demonstrates the broad applicability of the technique.

We learn three important lessons from these results:

1. *The level of incorrect interaction choices is dependent on the probability of having witnessed a recipient's interactions.* This probability is based on a number of factors, including the degree of a node, the observability in the population, and the number of interactions.
2. *Incomplete information has an observable effect on levels of emergent cooperation.* N&S note in their original paper that when moving from their initial model, equivalent to observability = 1, to observability of 0.1 a larger number of interactions are needed to establish cooperation. Our results corroborate this and establish that higher levels of incomplete information (whether caused by low node degree or high numbers of interactions) lead to more selfish societies.
3. *The levels of cooperation are highly dependent on the underlying topological structure of the social network.* Random graphs, and to a lesser extent, scale-free networks, significantly reduce the detrimental effects of incomplete information and aid the emergence of high levels of cooperation.

5 Summary and Further Work

Image scoring, as a model of the emergence of cooperation via indirect reciprocity, shows promising results and is highly applicable to modern open MAS domains such as MANETs and VANETs. We have explored the application of image scoring in settings more approximate to the realities of these domains than the original formulation of N&S, by scaling the number of interactions each agent engages in and introducing an underlying connection network. Changing from a completely connected to a random topology shows statistically significant differences in the levels of selfishness emergent in the population. While scale-free

networks are not consistently significantly different from random networks in terms of the levels of selfishness and cooperation, there do appear to be some benefits exhibited which require further investigation.

We have also shown that raising the ratio of m/n significantly increases the level of incomplete information in the society, with up to 62% of defection actions (in the completely connected topology scenario) being taken incorrectly. However, our investigation is limited with respect to the number of agents: given our illustrative domains, the number of agents in the population is likely to be far higher. We aim to investigate the effects of group size in future work, with a specific exploration of the relationship between m , n , and the ratio of m/n to levels of selfishness and incomplete information in the society.

We plan to further investigate the effects of scale-free topologies on image scoring, specifically investigating the effect of clustering on levels of cooperation and the ability of selfish agents to invade groups of cooperators. Within groups of highly clustered agents in scale-free networks there will be high visibility of interactions, and these clusters may therefore be very robust to selfish action. Scale-free topologies display high robustness against non-targeted malicious action, and we hypothesise that this will benefit image scoring.

A key effect to note from our results is that of visibility of interactions (as defined by node degree and our observability parameter) on the levels of cooperation observed. Many open MAS domains are characterised by sparse topologies and our results appear to show the efficacy of image scoring is reduced in such settings. To mitigate this effect we plan to investigate the use of gossiping algorithms (which exhibit low overheads and are highly robust) to spread results of interaction observations, thus removing the need for direct observation. Sommerfeld et al. [21] have reported on this effect in human subjects with promising results. Gossiping has been applied successfully within the specific topological challenges of VANETs [3, 5] and MANETs [4], and also within the domain of reputation mechanisms [2, 15].

References

1. R. Albert and A.L. Barabási. Statistical mechanics of complex networks. *Reviews of modern physics*, 74(1):47–97, January 2002.
2. 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.
3. 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.
4. S. Buchegger. Coping with Misbehavior in Mobile Ad-hoc Networks. PhD Thesis, *Ecole Polytechnique Federale de Lausanne*. 2005.
5. 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.
6. J. Delgado. Emergence of social conventions in complex networks. *Artificial Intelligence*, 141(1-2):171–185, October 2002.

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. P. Erdős and A. Rényi. On the evolution of random graphs. *Publication of the Mathematical Institute of the Hungarian Academy of Sciences*, 5:17–61, 1960.
9. E. Fonseca and A. Festag. A Survey of Existing Approaches for Secure Ad Hoc Routing and Their Applicability to VANETS. Technical Report NLE-PR-2006-19, NEC Network Laboratories, 2006.
10. N. Griffiths. Tags and image scoring for robust cooperation. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems*, pages 575–582. International Foundation for Autonomous Agents and Multiagent Systems, 2008.
11. N. Griffiths, A. Jhumka, A. Dawson, and R. Myers. A Simple Trust model for On-Demand Routing in Mobile Ad-Hoc Networks In *Proceedings of the 2nd International Symposium on Intelligent Distributed Computing (IDC 2008)*, pages 105–114, 2008.
12. D. Hales and B. Edmonds. Applying a Socially Inspired Technique (Tags) to Improve Cooperation in P2P Networks. *IEEE Transactions on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 35(3):385–395, May 2005.
13. C. Harsch, A. Festag, and P. Papadimitratos. Secure Position-Based Routing for VANETs. *Proceedings of 66th Vehicular Technology Conference*, pages 26–30, September 2007.
14. M. Milinski, D. Semmann, T. C. Bakker, and H. J. Krambeck. Cooperation through indirect reciprocity: image scoring or standing strategy? *Proceedings. Biological sciences / The Royal Society*, 268(1484):2495–501, December 2001.
15. J. Mundinger and J. Le Boudec. Reputation in self-organized communication systems and beyond. *Proceedings from the 2006 workshop on Interdisciplinary systems approach in performance evaluation and design of computer & communications systems*, page 3, 2006.
16. M. Nowak and K. Sigmund. The dynamics of indirect reciprocity. *Journal of theoretical biology*, 194(4):561–74, October 1998.
17. M. Nowak and K. Sigmund. Evolution of indirect reciprocity. *Nature*, 437(7063):1291–8, 2005.
18. J. Ott and D. Kutscher. Drive-thru Internet: IEEE 802.11b for Automobile users. In *Proceedings of the Twenty-third Annual Joint Conference of the IEEE Computer and Communications Societies*, page 373. March 2004.
19. A. A. Pirzada and C. McDonald. Trust Establishment In Pure Ad-hoc Networks. *Wireless Personal Communications*, 37(1-2):139–168, 2006.
20. Q. Sen. Scale-free topology structure in ad hoc networks. In *11th IEEE International Conference on Communication Technology, 2008. ICCT 2008*, pages 21–24. IEEE, November 2008.
21. 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, October 2007.
22. O. Trullols-Cruces, J. Morillo-Pozo, J. M. Barcelo, and J. Garcia-Vidal. A Cooperative Vehicular Network Framework. In *International Conference on Communications*, pages 1–6. IEEE ICC, 2009.