

Emergence of Cooperation in N-player games on small world networks

Colm O’Riordan¹, Alan Cunningham¹ and Humphrey Sorensen²

¹NUI, Galway, Ireland

²University College Cork, Ireland.

colm.oriordan@nuigalway.ie

Abstract

The emergence of cooperation in social dilemmas has been addressed in a number of fields. In this paper, we illustrate how robust cooperation can emerge among a population of agents participating in a N-player dilemma when the agents are spatially arranged on a graph exhibiting small world properties. We present a graph structure with a high level of community structure, small diameter and a variance in the node degree distribution. We show that with simple learning rules, robust cooperation emerges. We also show that a population of agents whose interactions are constrained by such a graph can adapt to dramatic environmental changes.

Introduction

Questions regarding cooperation and its emergence, particularly in environments inhabited by self interested individuals, have been addressed in a many domains. They include, among many others, computer science (Chiba and Hiraishi, 1998), biology (Boyd and Richerson, 1988), robotics (Birk, 1999) and social science (Hardin, 1968).

Social dilemma games have been commonly adopted to capture and represent the salient features of interactions in these environments; in particular the conflict between the individually rational actions and the collectively rational group actions and outcomes. The prisoner’s dilemma (and variations) is the most oft studied game. Most previous work has focussed on the case involving two participants. The extended N-player version is less studied but it has been argued by Davis et al. (1976) to have “greater generality and applicability to real life situations”.

In N-player dilemma games defection is the rational choice for all individuals which in turn leads to a sub-optimal outcome for the group. Many researchers have investigated the effect of spatial constraints on agent interactions in both the 2-player and N-player game (Hauert, 2006), (Wu et al., 2005), (Santos and Pacheco, 2005). In these spatially organised games, agents are more likely to interact with a smaller subset of agents than would be expected in simulations where agents are not spatially organised, e.g. randomly organised or round robin type simulations. This

factor has been shown to have a dramatic impact of the likelihood of cooperation emerging.

One form of spatial arrangement or topology that has generated much attention recently is that of a *small world graph* (Watts, 1999). Small world graphs are typified by the fact that most nodes are reachable from all other nodes in a short number of steps. These graphs also tend to have a high clustering coefficient with a high presence of cliques or near-cliques. Another property often associated with small world graphs is that the node degree distribution follows a power law distribution.

One key property that we have explored in previous work is that of *community structure* (O’Riordan and Sorensen, 2008b). This property has also been explored in recent work (Lozano et al., 2006). A graph is said to have a community structure if collections of nodes are joined together in tightly knit groups between which there are only looser connections. This property has been shown to exist in many real-world social networks (Newman and Girvan, 2004).

In our previous work, we have shown that by enforcing a high level of community structure robust cooperation can emerge among agents participating in N-player social dilemma games. The topologies explored in our previous work, however, are quite unrealistic and do not possess the other properties found in many naturally occurring graphs, i.e. small world properties including a variance in node degree.

This paper investigates whether it is possible to build graphs that exhibit the properties of small world graphs which induce the emergence of cooperation. We present two different extensions to our previous representations and illustrate that by constructing the small world graph while maintaining a high level of community structure that cooperation can indeed still emerge.

The following sections discuss some background material, particularly in N-player social dilemmas, graphs with community structure and some of our previous findings. We then discuss the particular graph model and agent interaction models used in this work. The experimental set up is then explained with our two algorithms for creating small world

graphs explained. We present results obtained from simulations with these two different topologies. Finally we present some conclusions and briefly outline some intended future work.

Background

N-player social dilemmas

N-player dilemmas are characterised by having many participants, each of whom may choose to cooperate or defect. These choices are made autonomously without any communication between participants. Any benefit or payoff is received by all participants; any cost is borne by the cooperators only. A well-known example is the *Tragedy of the Commons* (Hardin, 1968). In this dilemma, land (the commons) is freely available for farmers to use for grazing cattle. For any individual farmer, it is advantageous to use this resource rather than their own land. However, if all farmers adopt the same reasoning, the commons will be over-used and soon will be of no use to any of the participants, resulting in an outcome that is sub-optimal for all farmers.

In the N-player dilemma game there are N participants. Each player is confronted with a choice: to either cooperate or defect. We represent the payoff obtained by a strategy which defects given i cooperators as $D(i)$ and the payoff obtained by a cooperative strategy given i cooperators as $C(i)$.

Defection represents a dominant strategy, i.e. for any individual, moving from cooperation to defection is beneficial for that player (they still receive a benefit without the cost):

$$D(i) > D(i-1) \quad 0 < i \leq N-1 \quad (1)$$

$$C(i) > C(i-1) \quad 0 < i \leq N-1 \quad (2)$$

$$D(i) > C(i) \quad 0 < i \leq N-1 \quad (3)$$

However, if all participants adopt this dominant strategy, the resulting scenario is sub-optimal and, from a group point of view, an irrational outcome ensues:

$$C(N) > D(0) \quad (4)$$

If any player changes from defection to cooperation, the performance of the society improves, i.e. a society with $i+1$ cooperators attains a greater payoff than a society with i cooperators:

$$(i+1)C(i+1) + (N-i-1)D(i+1) > (i)C(i) + (N-i)D(i) \quad (5)$$

Small world Graphs

As mentioned in the introduction, small world graphs are a class of graphs or topologies such that nearly all nodes are reachable from all other nodes in a few steps. Watts and Strogatz (Watts, 1999) demonstrated that a regular lattice can be transformed into a small world network by making

a small fraction of the connections random. The algorithm involves taking a regular lattice (ring, grid) and repeatedly removing some edge (a, b) and replacing it with an edge (a, c) . If the node c is selected with probability based on its degree, then the notion of preferential attachment is present which results in a graph with node degree distribution following a power law.

The property of community structure has been reported in several real world networks (Newman and Girvan, 2004) and many algorithms have been proposed to measure the level of community structure present in the graph (Donetti and Munoz, 2004) (Zhang et al., 2007).

Such graphs have been used to constrain agent interactions in social dilemma games in interesting work (Wu et al., 2005), (Santos and Pacheco, 2005) which show that cooperation can be induced in 2-player games. Our work differs by addressing the N-player version which has been shown to be more challenging to induce cooperation in evolutionary settings (Yao and Darwen, 1994). We also show that the maintenance of one key property, that of community structure is of importance.

N-player dilemmas and Community Structure

In previous work, we have created a range of lattices which can be tuned to exhibit different levels of community structure (O’Riordan and Sorensen, 2008b). These graphs do not exhibit node degree distribution according to power laws (in fact, the degree is constant throughout the graph) and they also do not exhibit other small world properties. In the model previously adopted we created graphs with strongly connected clusters of agents who were loosely connected to neighbouring clusters. We varied the degree of community structure by simply varying the ratio of the weights on intra-community edges to intra-community edges. Agents were chosen to interact based on the strength of the edge weights. We allowed agents learn from their immediate neighbours; agents effectively imitated their more successful neighbours. If all immediate neighbours perform similarly, agents were allowed to learn from neighbouring clusters. We showed that cooperation emerged. Our initial model is discussed more in the following section.

Model

Initial Graph Topology

In the simulations described in this paper, agents are located on nodes of a graph. The graph is an undirected weighted graph. The weight associated with any edge between nodes represents the strength of the connection between the two agents located at the nodes. This determines the likelihood of these agents participating together in games.

The graph is static throughout the simulation: no nodes are added or removed and the edge weights remain constant.

We use a regular graph: all nodes have the same degree. In the initial topology, nodes have four neighbours. We use two

different edge weight values in each graph: one (a higher value) associated with the edges within a community and another (a lower value) associated with the edges joining agents in adjacent communities. All weights used in this work are in range [0,1].

The graph is depicted in Fig. 1, where the thicker lines represent intra-community links (larger value as edge weight) and the thinner lines indicate inter-community links between neighbouring communities. The rectangles of thicker lines represent a community; the vertices represent agents.

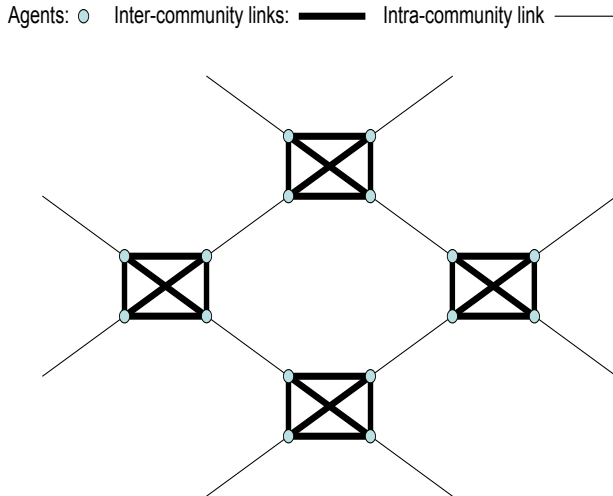


Figure 1: Graph with community structure

Agent Interactions

Interaction Model Agents in this model can have a strategy of either cooperation (C) or defection (D). Agents interact with their neighbours in a N-player prisoner's dilemma. The payoffs received by the agents are calculated according to the formula proposed by Boyd and Richerson (Boyd and Richerson, 1988), i.e. cooperators receive $Bi/N - c$ and defectors receive Bi/N , where B is a constant (in this paper, B is set to 5), i is the number of cooperators involved in the game, N is the number of participants and c is another constant (in this paper, c is set to 3).

Each agent may participate in several games. The algorithm proceeds as follows: for each agent a in the population, agents are selected from the immediate neighbourhood of agent a to participate in the game. Neighbouring agents are chosen to participate with a probability equal to the edge of the weight between the nodes. This means that, for a population with a high community structure, most games involve an agent's local community members. This allows a high degree of insulation from agents in neighbouring communities. An agent's fitness is calculated as the average pay-

off received in the interactions during a generation.

Learning

Agents may change their behaviours by comparing their payoff and that of neighbouring agents. We adopt a simple update rule whereby an agent updates their strategy to those used by more successful strategies. Following each round of games, agents are allowed to learn from their neighbours. Again these neighbours are chosen stochastically; the neighbours are chosen according to the weight of the edge between agent and neighbour.

We incorporate a second update mechanism. The motivation for its inclusion is as follows. Following several iterations of learning from local neighbours, each community is likely to be in a state of equilibrium—either total cooperation or total defection. Agents in these groups are receiving the same reward as their immediate neighbours. However, neighbouring communities may be receiving different payoffs. An agent that is equally fit as its immediate neighbours may look further afield to identify more successful strategies.

In the first update rule, agents consider other agents who are immediate neighbours. Let $s_adj(x)$ denote the immediate neighbours of agents x chosen stochastically according to edge weight. The probability of an agent x updating their strategy to be that of a neighbouring agent y is given by:

$$\frac{w(x, y) \cdot f(y)}{\sum_{z \in s_adj(x)} w(x, z) \cdot f(z)} \quad (6)$$

where $f(y)$ is the fitness of an agent y and $w(x, y)$ is the weight of the edge between x and y .

The second update rule allows agents to look further afield from their own location and consider the strategies and payoffs received by agents in this larger set, i.e. agents update to a strategy y according to:

$$\frac{w(x, y) \cdot f(y)}{\sum_{z \in adj(adj(x))} w(x, z) \cdot f(z)} \quad (7)$$

where again $f(y)$ is the fitness of agent y and now $w(x, z)$ refers to the weight of the path between x and z . We use the product of the edge weights as the path weight. Note that in the second rule, we don't choose the agents in proportion to their edge weight values; we instead consider the complete set of potential in the extended neighbourhood. In this way all agents in a community can be influenced by a neighbouring cooperative community.

Small World version of Graph

In order to create a graph topology more reflective of naturally occurring graphs, the basic graph topology must be changed. This is achieved by adopting the approach proposed by Watts (1999).

Our first approach involves taking our existing graph structure and re-attaching edges i.e. the following procedure is repeated: edge (a, b) is randomly selected from the set of edges present and replaced with the edge (a, c) where node c is selected in proportion to its degree. The new edge will have a weight equal to the deleted one. This approach, while introducing the small world property and the desired node degree distribution seriously damages the community structure. We hypothesise that this should negatively impact on the emergence of cooperation.

Our second approach begins with another regular graph structure; we place the communities of agents on a ring (depicted in Fig. 2). We again re-attach edges in the graph, but with the following constraint; only inter-community edges are deleted and re-attached. Thus, we choose an edge (a, b) randomly such that both a and b are on the circumference of the ring; this edge is deleted and re-attached as (a, c) such that c is selected randomly from those nodes positioned on the circumference. Again, the new edge will have a weight equal to the deleted one. This approach maintains the community structure in the graph while introducing the desired small world graph properties.

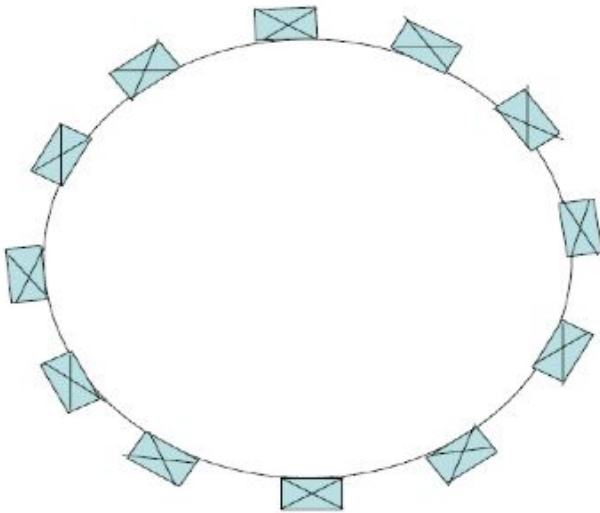


Figure 2: Ring structure with communities

The following tables present some data to illustrate some of the properties of the resulting graphs for the two different algorithms.

Experiment Setup

A population of 800 agents is used. Strategies are assigned to agents randomly. We allow simulations to run for 200 generations.

Following each generation, the first learning rule is applied. Following every four generations (sufficient for community to reach an equilibrium), the second learning rule

Re-attachment Rate	Dev. in Node Degree	Av. Diameter
0%	0	15.59
1%	0.261	12.14
5%	0.620	8.52
10%	0.835	7.49

Table 1: Properties of resulting small world graph using first algorithm (initial lattice structure, all edges considered for re-attachment) for different levels of re-attachment

Re-attachment Rate	Dev. in Node Degree	Av. Diameter
0%	0.5	132.9
1%	0.512	41.17
5%	0.532	5.655
10%	0.589	2.136

Table 2: Properties of resulting small world graph using second algorithm (initial ring structure, inter-community edges considered for re-attachment) for different levels of re-attachment

is applied. This is not necessary in most cases; we merely choose to let the local interactions stabilise prior to applying the second rule. This eases analysis in some cases where fluctuations can occur if community structure levels are not sufficiently high. We include another plot in a later section in this paper where we use the ring graph with re-attachment and modify the rates of application of learning rules such that they both occur every generation. The outcome is similar.

In all the simulations we initially enforce a high level of community structure; the intra-community links are held constant with a value of one. The value of the inter-community are set to 0.1. The lower the value, the more insulated clusters are and hence should promote cooperation.

We vary the level of re-attachment and measure the resulting levels of cooperation.

Results

Emergence of Cooperation

For the first graph structure (regular lattice) with different levels of edge re-attachment, we see that the levels of cooperation is dependent on the degree of re-attachment present (see Fig. 3). For a regular graph with high community structure and no other small world properties, we see that the population quickly converges to cooperation. Introducing 1% re-attachment reduces the diameter of the graph and increases the node degree deviation but also damages the level of community structure. We see that the levels of cooperation reached fall to roughly 700 cooperators in the population. As the level of re-attachment increases, the ef-

fect becomes even more pronounced with a big decrease in the number of cooperators for re-attachment level of 5% and a large collapse in the number of cooperators for re-attachment levels of 10%.

It is worth commenting on the nature of the fluctuations in the separate runs. Consider, as an example, the line indicating re-attachment levels of 5% where the levels of cooperation fluctuate considerably. This is due to the effect of the two learning rules and the frequency with which they are applied. Following a few generations, each community converges to total cooperation or total defection. Following every fourth generation, the second rule is applied with leads to an immediate increase in the number of cooperators as members of non-cooperating clusters imitate more successful clusters. These new cooperators are in most cases interacting with non-cooperators and hence are exploited by their immediate neighbours. These immediate neighbours are then imitated leading to emergence of defection in these clusters.

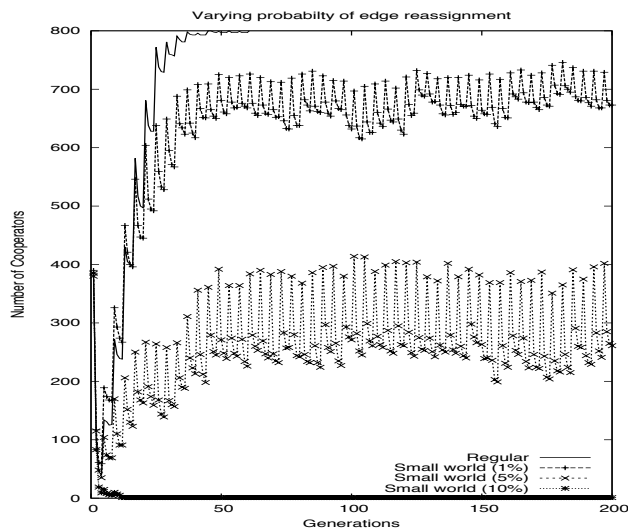


Figure 3: Levels of Cooperation present in population placed on small world graph created from original lattice

Fig. 4 shows the levels of cooperation attained given a graph with small world properties that also has initially a high level of community structure created by re-attaching inter-community edges only.

We see that for levels of re-attachment up to 10%, cooperation still emerges. These results illustrate that we can have small world properties (e.g. small diameter) and still maintain community structure and hence maintain high levels of cooperation.

An interesting point to note is that cooperation reaches the maximum possible level most quickly for re-attachment levels of 5%. This is due to the reduction in the diameter which causes cooperation to spread more quickly as non-

cooperative clusters are more likely to be close to cooperative clusters. However, increasing the level of re-attachment further slows down the spread of cooperation. This is because, despite the potential gain caused by the decrease in diameter, the increased probability of having a number of nodes with a high degree which can be influenced more readily by non-cooperating strategies.




Figure 4: Levels of Cooperation present in population placed on small world graph created from ring; only inter-community links re-attached.

Robustness

In many scenarios that we may wish to model, it is possible for uncertainty or noise to exist—agents may perform their acts incorrectly or imperfectly, their acts may be misinterpreted, agents may learn or imitate others and change their behaviour accordingly, and agents may exit or join the group thereby changing the environment or others. Alternatively, the environment may change dramatically thereby requiring agents to explore and learn new suitable behaviours.

In previous work O’Riordan and Sorensen (2008a), we showed that these graph structures allowed populations to be robust to noise and to dramatically changing environments for a population of generalised *tit-for-tat* strategies. In this experiment we explore again if the population of agents can survive and track dramatic environmental change. We introduced 1% noise to ensure some exploration of the strategy space. At every generation, each agent has a 1% probability of changing strategy. We also introduce dramatic environmental change during the simulation. This involves reversing the payoffs of the game which causes ‘cooperation’ now to be viewed as individually rational and collectively suboptimal and renders ‘defection’ the new socially beneficial and cooperative act.

In our simulation, we count and plot the number of agents choosing the socially beneficial action. We see that following the change in environment at generation 350, the population recovers to high levels of cooperation (Fig. 5).

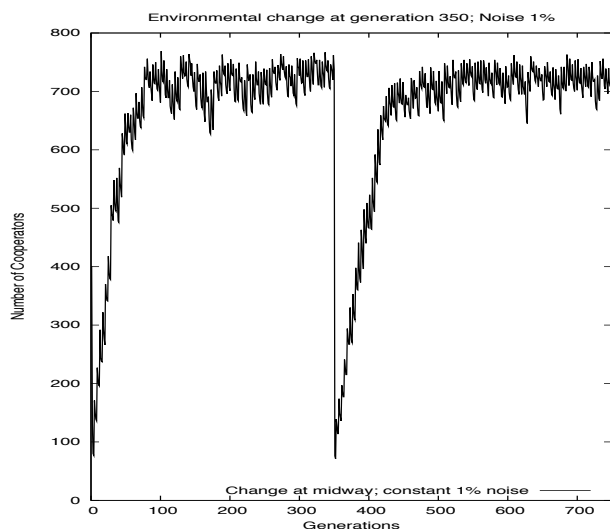


Figure 5: Levels of cooperation present on ring with re-attachment probability 5%. Noise is set to 1%. Dramatic environment reversal occurs at generation 350

Conclusions

In this paper, we wished to explore if cooperation can emerge among self-interested agents participating in N-player social dilemmas where the agents are placed on a small world network exhibiting community structure. Our previous work illustrated the emergence of cooperation given a community structure on regular lattices. In this paper, we showed that by converting the graph to a small world network by re-attaching edges in such a manner that damaged the community structure, cooperation collapses and defection emerges as the norm. We also showed that by converting a regular graph to a small world graph while taking care to preserve the community structure, cooperation can emerge as the norm. The speed of the emergence can also be improved by having small world properties (e.g. reduced diameter).

We have shown that the notion of community structure is a key feature in the emergence of cooperation. In order for cooperative clusters to survive, the agents must be able to protect themselves from non-cooperative agents by insulating themselves and playing mainly among themselves. However, communities or clusters cannot be totally isolated; they must have some link to other communities so as to provide an opportunity to learn more beneficial strategies if possible. A balance must be struck between the risk of exploitation and the potential to learn a better strategy if one exists.

Discussion

In the experiments in this paper, we utilise two learning rules—one involving an agent’s immediate neighbours, the other involving an extended neighbourhood. We allow the agents to learn from the immediate neighbours first and then upon reaching an equilibrium we allow them to learn from the neighbouring communities. We achieve this by allowing the first learning rule every generation and the second learning rule every four generations.

The motivations were primarily to allow local communities reach an equilibrium prior to learning from others as otherwise, in some cases, this causes fluctuations in levels of cooperation and convergence is never reached. This occurs when there are insufficient levels of community structure. In these cases, a local community may be heading towards to defection and then learn from neighbours and heads towards cooperation again etc. We wished to ease the complexity of the interactions for these specific cases.

However, it should be noted for the results presented for the main graph of interest (the ring transformed into a small world graph while maintaining the level of community structure), the rates of application of the learning rules does not dramatically interfere with the results. The same trends are noticed (Fig. 6) where we apply both learning rules every generation. The agents reached the state of total cooperation much more quickly due to the application of the second learning rule every generation.

Future Work

There are several directions for future work. One future direction would be to explore the generalizability of these results. We wish to explore different updating schemes and other social dilemma games to explore if the same effects are detected. We also wish to explore these graph structures in uncertain environments. Another track which we will pursue is to investigate under which conditions these graphs can emerge based on agent interactions. In this paper, agent interactions are constrained by the graph properties. It would be interesting to show that such graphs can emerge based on interactions between agents.

References

- Birk, A. (1999). Evolution of continuous degrees of cooperation in an n-player iterated prisoner’s dilemma. *Technical Report, Vrije Universiteit, Brussel, AI-Laboratory, 1999.*
- Boyd, R. and Richerson, P. (1988). The Evolution of Reciprocity in Sizable Groups. *Journal of Theoretical Biology*, 132:337–356.
- Chiba, K. and Hiraishi, K. (1998). Iterated continuous prisoner’s dilemma game and its usefulness in analyzing multi-agent systems. *IEEE International Conference on Systems, Man, and Cybernetics*, 1:644–649.

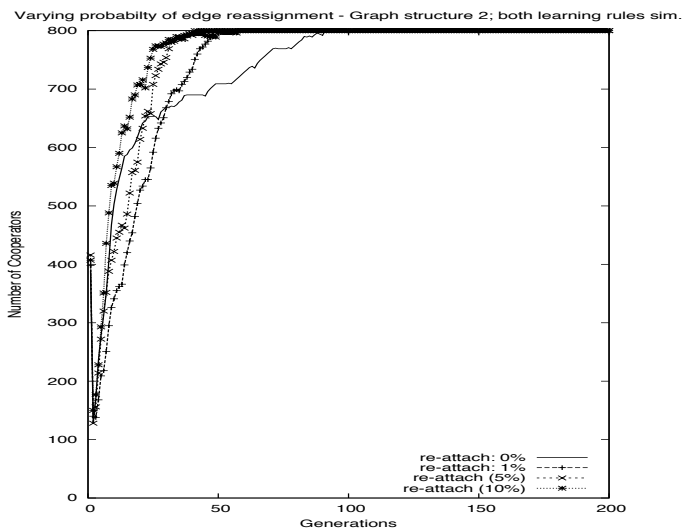


Figure 6: Small world ring; both learning rules applied each generation

Davis, J., Laughlin, P., and Komorita, S. (1976). The social psychology of small groups: Cooperative and Mixed-Motive Interaction. *Annual review of Psychology*, 27:501–541.

Donetti, L. and Munoz, M. (2004). Detecting network communities: a new systematic and powerful algorithm. *Journal of Statistical Mechanics*, P10012.

Hardin, G. (1968). The tragedy of the commons. *Science*, 162(3859):1243–1248.

Hauert, C. (2006). Spatial Effects in Social Dilemmas. *Journal of Theoretical Biology*, 240(4):627–636.

Lozano, S., Arenas, A., and Sanchez, A. (2006). Mesoscopic structure conditions the emergence of cooperation on social networks. *arXiv:physics/0612124v1*.

Newman, M. E. J. and Girvan, M. (2004). Finding and evaluating community structure in networks. *Physics Review E*, 69.

O’Riordan, C. and Sorensen, H. (2008a). Stable cooperation in changing environments (short paper). *7th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*.

O’Riordan, C. and Sorensen, H. (2008b). Stable cooperation in the n-player prisoner’s dilemma: the importance of community structure. *Lecture Notes in Artificial Intelligence LNAI 4865, Adaptive Agents and Multi-Agent Systems III*.

Santos, F. and Pacheco, J. (2005). Scale-free Networks Provide a Unifying Framework for the Emergence of Cooperation. *Physical Review Letters*, 95(9).

Watts, D. J. (1999). *Small Worlds: The Dynamics of Networks Between Order and Randomness*. Princeton University Press.

Wu, Z., Xu, X., Chen, Y., and Wang, Y. (2005). Spatial prisoner’s dilemma game with volunteering in newman-watts small-world networks. *Physical Review E*, 71:37103.

Yao, X. and Darwen, P. J. (1994). An experimental study of N-person iterated prisoner’s dilemma games. *Informatica*, 18:435–450.

Zhang, S., Wang, R., and Zhang, X. (2007). Identification of overlapping community structure in complex networks using fuzzy c-means clustering. *Physica A*, 374:483–490.