2011 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology

Norm Establishment via Metanorms in Network Topologies

Samhar Mahmoud, Jeroen Keppens, Michael Luck
Department of Informatics
King's College London
London WC2R 2LS, UK.
samhar.mahmoud@kcl.ac.uk

Nathan Griffiths
Department of Computer Science
University of Warwick
Coventry CV4 7AL, UK.

***Abstract*—In order to established a norm in a society of agents, metanorms have previously been proposed as a means of ensuring not that norms are complied with, but that they are enforced. Yet while experimental results have shown that metanorms are effective in fully-connected environments such as that used by Axelrod, there has been limited consideration of such metanorm models with different but more realistic topological configurations. In this paper, therefore, we consider the use of metanorms in supporting norm establishment in lattices and small world networks. Our results suggest that norm establishment is achievable in lattices and small worlds.**

I. INTRODUCTION

In peer-to-peer systems, agents share resources (hardware, software or information) with others, but if there is no cost to access files nor any limit on the number of files accessible, then there is no incentive to respond to requests and, more generally, to establish cooperation in the system. Yet cooperation is needed: when self-interested autonomous agents must exchange information without any central control, non-compliance (due to selfish interests) can compromise the entire system. The use of *norms* to provide a means of ensuring cooperative behaviour has been proposed by many but, as shown by Axelrod [1], norms alone may not lead to the desired outcomes. In consequence, *metanorms* have been proposed as a means of ensuring not that norms are complied with, but that they are enforced. While experiments have shown that metanorms are effective in fully-connected environments as used by Axelrod, there has been limited consideration of metanorms with different but more realistic topological configurations, which fundamentally change the mechanisms required to establish cooperation.

Some work has already been undertaken on examining the impact of different topologies on norm establishment. For example, Savarimuthu et al. [2] consider the *ultimatum game* in the context of providing advice to agents on whether to change their norms in order to enhance performance for random and scale-free networks. Delgado et al. [3] study norm emergence in coordination games in scale-free networks, and Sen et al. [4] examine rings and scale-free networks in a related context. Additionally, Villatoro et al. [5] explore norm emergence with memory-based agents in lattices and scale-free networks. While these efforts provide valuable and useful results, the context of application has been limited, with only two agents involved in each encounter, rather than a larger population of agents. This simplifies the problem when compared with those in which the actions of multiple interacting agents can impact on norm establishment. In particular, Axelrod's seminal model [1] has provided the foundation for several investigations into norm emergence, yet offers a very general framework comprising the use of norms and metanorms in populations of agents where the overall behaviour determines whether a norm is established. In this paper we extend Axelrod's model to address the context of different topological configurations.

The paper begins with an outline of Axelrod's metanorms game, adjusted to suit the purposes of this paper, and augmented with a learning mechanism. Section III then considers the problems that arise from the use of different topologies, and Sections IV and V describe in detail the impact of applying the model in lattices and small worlds.

II. THE METANORMS GAME

Inspired by Axelrod [1], our simulation focusses only on the essential features of the problem. Agents play a game iteratively; in each iteration, they make a number of binary decisions. First, each agent decides whether to comply with a norm or to defect. Defection brings a reward for the defecting agent, and a penalty to all other agents, but each defector risks being observed by the other agents and punished as a result. These other agents thus decide whether to punish agents that were observed defecting, with a low penalty for the punisher and a high penalty for the punished agent. Agents that do not punish those observed defecting risk being observed themselves, and potentially incur metapunishment. Thus, finally, each agent decides whether to metapunish agents observed to spare defecting agents. Again, metapunishment comes at a high penalty for the punished agent and a low penalty for the punisher.

The behaviour of agents in each round of the game is random, but governed by the probability of being seen, boldness, and vengefulness. Each round, agents have a fixed number of opportunities to defect, each of which has a randomly selected probability of a defection being seen. Bold-

978-0-7695-4513-4/11 $26.00 © 2011 IEEE
DOI 10.1109/WI-IAT.2011.169

ness (B) determines the probability that an agent defects, such that if an agent's boldness exceeds the probability of a defection being seen then the agent defects. Vengefulness (V) is the probability that an agent punishes or metapunishes another agent. Thus boldness and vengefulness of an agent comprise its strategy. After several rounds, each agent's rewards and penalties are tallied, and successful and unsuccessful strategies identified. By comparing themselves to other agents on this basis, the strategies of poorly performing agents are revised such that features of successful strategies are more likely to be retained than those of unsuccessful ones. We need not be concerned with the details of this in this paper, beyond the fact that boldness and vengefulness are simply revised upward or downward as appropriate, in line with a specified learning rate. If most agents employ a strategy of low boldness and high vengefulness, it can be argued that the norm has become *established*, because strategies that lead to defection or to sparing defecting agents are unlikely and lead to high penalties.

III. IMPOSING TOPOLOGIES ON METANORMS

Despite its success, Axelrod's model omits consideration of some important aspects. In particular, real-world domains, such as peer-to-peer and wireless sensor networks, are not fully connected, with agents tending to interact with a small subset of others on a regular basis, yet it is only through such interactions that defection can be observed and punishment administered. Other network topologies must therefore be considered, reflecting different potential configurations, in which agents are connected only to a subset of other agents, their *neighbours*. This constraint on connectivity between agents implies some adjustments to Axelrod's model.

First, Axelrod assumes that an agent's defection penalises *all* other agents in the population. The introduction of a topology enables us to restrict the penalty to only those agents with which the defector interacts. Second, Axelrod assumes agents to be able to observe the entire population. By introducing a topology, we employ a more realistic model in which an agent can only observe (with a certain probability) those agents with which it interacts. Third, Axelrod does not require observation of misbehaviour for punishment. However, by introducing constraints on observation and rendering the model more realistic, an agent should only punish a defector if the agent can observe the defector. In addition, an agent should only metapunish an agent that fails to punish a defector if it can observe both the defector *and* the agent that fails to punish the defector. Finally, in order to enhance an agent's individual performance, it compares itself to others in the population before deciding whether to modify strategy. However, since agents can only observe their neighbours, these are the only agents they are able to learn from. In what follows, we consider these modifications to the basic model in the context of lattices and small worlds.

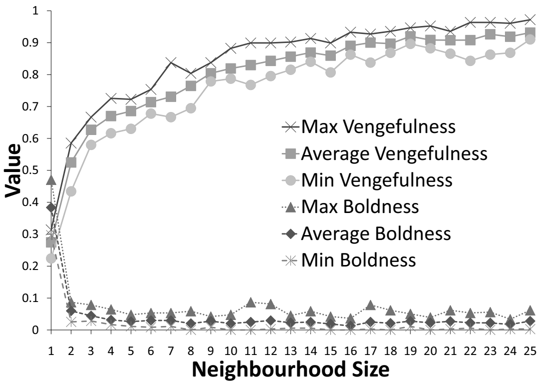

Figure 1. Lattice: impact of neighbourhood size on B and V

IV. METANORMS IN LATTICES

In a (one-dimensional) lattice with neighbourhood size n, agents are situated on a ring, with each agent connected to its neighbours n or fewer hops (lattice spacings) away, so that each agent is connect to exactly $2n$ other agents. Thus, in a lattice topology with $n = 1$, each agent has two neighbours and the network forms a ring. In a lattice topology with $n = 3$, each agent is connected to 6 neighbours.

A. Neighbourhood Size

It is clear that, depending on neighbourhood size (NS), lattices may be more or less connected. Those with larger neighbourhood sizes are more similar to Axelrod's fully connected model; our hypothesis is that an increased NS enables punishment and metapunishment to become more effective in reducing boldness and increasing vengefulness. To investigate this, we ran several experiments.

In our first set of experiments, we used 51 agents (so we have an even number, plus 1, to account for the $2n$ neighbours plus our original agent), and varied the NS between the least connected lattice (the ring topology) and the most connected lattice ($n = 25$). Each experiment involved 10 separate runs, with each run comprising 1,000 timesteps, for a particular NS. The results of all runs were averaged, as shown in Figure 1, with NS plotted against B and V.

For the least connected lattice (NS of 1), no norm is established, as runs ended in both relatively low boldness and relatively low vengefulness: though agents rarely defect, they also rarely punish defection. This is an unstable situation in which defecting could be a rewarding behaviour for agents as it is relatively unlikely to be penalised. However, by increasing NS to 3, boldness drops almost to 0, with agents not defecting. While vengefulness increases, it still does not correspond to norm emergence, since agents might still not punish a defection without being metapunished.

A neighbourhood size (NS) as small as 2 is enough to maintain boldness near 0, indicating that agents only defect

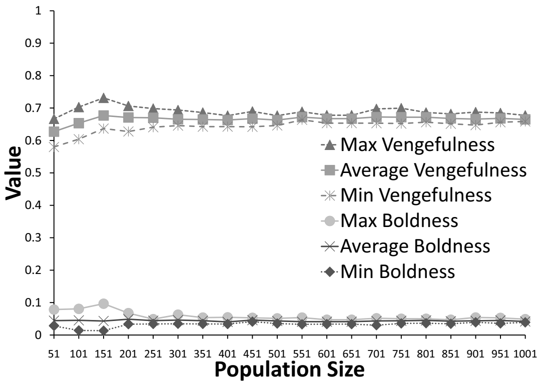

Figure 2. Lattice: impact of population size on B and V (n=3)

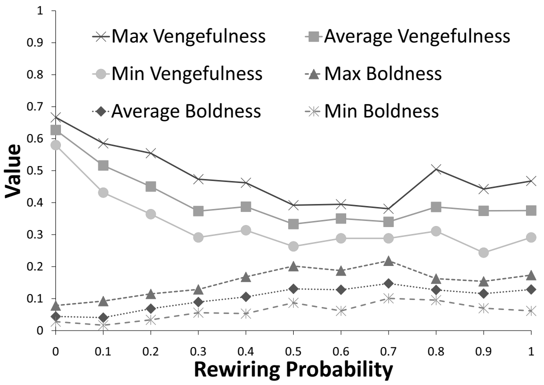

Figure 3. Small world: impact of rewiring on B and V (n=3)

when they *explore* as a result of adopting random strategies (introduced for comparability with Axelrod's model). Conversely, increasing NS has a major impact on vengefulness, until NS reaches around 15 (at which point an agent is connected to more than half the population) when it brings only very minor change. This is because, in poorly connected environments, agents not punishing defection can more easily escape metapunishment than in more connected environments. Increasing NS thus brings a corresponding effect on agents (in terms of boldness and vengefulness). Only the most poorly connected lattices have moderate levels of boldness, with vengefulness increasing monotonically over a longer period before stabilising at norm establishment. Connections between agents give rise to this behaviour, with an increase in connections providing more opportunities for agents to respond to defectors appropriately.

B. Population Size

If we increase population size (PS) while keeping NS static, we decrease the relative number of connections among the overall population. This suggests that convergence to norm establishment should decrease, in line with the above results. In the second set of experiments, therefore, NS was fixed and PS varied between 51 and 1,001 agents. However, the results shown in Figure 2 for NS of 3 (though other values gave similar results) are not as expected, and suggest that increasing PS has no effect on the rate of norm emergence, as all runs for all sizes of population end almost with the same level of boldness and vengefulness.

These results suggest that norm emergence in a community of agents that interact in a lattice is not affected by total PS but by NS. By increasing the number of neighbours, norm establishment becomes more likely, irrespective of PS. In other words, the likelihood of norm establishment is governed by the total amount of punishment that could potentially be brought upon a defector or an agent failing to punish a defector, which may be termed the *potential peer pressure* of a lattice. This is because such lattices comprise multiple overlapping localities in which agents are highly connected: via punishments, agents in these localities impose a strong influence on their neighbours, and increasing PS simply increases the number of such overlapping regions.

V. METANORMS IN SMALL WORLDS

While lattices are regular structures, as opposed to random structures, Watts and Strogatz noted that many biological, technological and social networks lie somewhere between the two: neither completely regular nor completely random [6]. They instead proposed *small world networks* as a variant of lattices in which agents are connected to others n or fewer hops away, but with some of connections replaced by connections to other randomly selected nodes in the network, in line with a specific rewiring probability (RP).

Thus, while lattices essentially create overlapping localities of well connected agents (since agents are connected to $2n$ agents surrounding them), small worlds break these connections. Though the number of connections does not change, the locality effect does, since there may no longer be localities of well connected agents, but instead agents with some connections to their local neighbours, and some connections to others elsewhere in the network. As these local regions break down, the strong influence of an agent's local neighbours, causing compliance with norms, should also break down because of the more sparse connections.

To verify this hypothesis, we investigated the impact of the rewiring probability by running the model with different values, in populations of 51 agents, for different neighbourhood sizes of 3 and 5. The results with neighbourhood size (NS) of 3 are shown in Figure 3, indicating that increasing the RP decreases the final average vengefulness in the population. With a NS of 5 the results are similar (not shown).

This is because, as a result of rewiring, agents no longer affect just their locality, but now affect agents that are much further away, consequently requiring establishment of the norm in multiple localities. For example, in the case of NS of 3, it is clear that not only is the norm not established, but

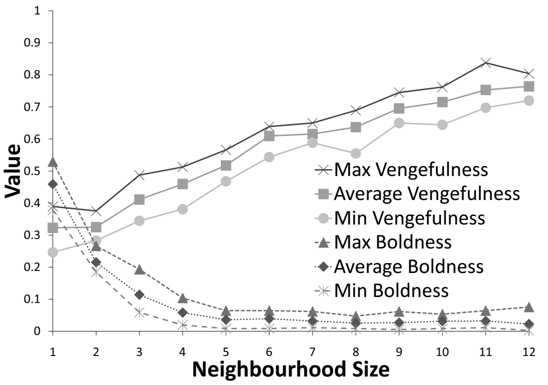

Figure 4. Small world: impact of neighbourhood size on B and V (RP=0.4)

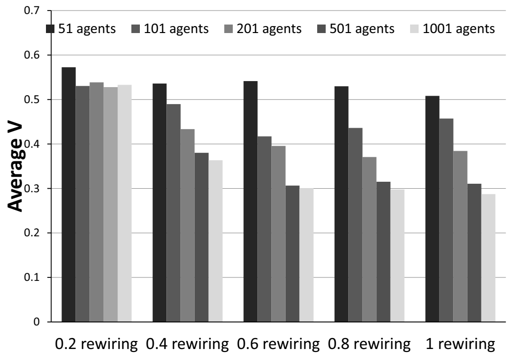

Figure 5. Small world: impact of rewiring and population size on V

as the RP rises above small values, the trend moves further away from establishment, since the connections of agents are increasingly rewired, giving a locality effect similar to lattices with NS of 2 (Section IV-A). In addition, rewiring to other agents further away brings the need to establish the norm in all those localities to which an agent is connected, making it much more difficult. It is clear that the RP of small worlds does not impact the level of defection in the population since, independently, boldness remains very low, indicating that agents are very unlikely to defect.

A. Neighbourhood Size and Rewiring

While increasing NS causes an increase in vengefulness in lattices, in seeking to understand this in small worlds, we repeated the lattice experiments in this new context, for different values of the RP. Results for RP of 0.4 are shown in Figure 4 (with results for other values of the RP being similar in trend), again showing that NS increases vengefulness. However, in comparison to lattices, vengefulness in small worlds is lower for the same NS. This is because agents must now respond to defections in different regions of the network, where there is less influence on behaviour, thus potentially incurring greater enforcement costs.

B. Population Size and Rewiring

Population has been shown to have no effect on norm establishment in lattices due to the *potential peer pressure* arising from the multiple overlapping localities. However, since these concentrated local regions of connected agents are weakened in small worlds, we repeated the previous experiments to determine the effect with RPs of 0.2, 0.4, 0.6, 0.8 and 1.0, and NS of 5. The results indicate that boldness is not affected by the changes of the PS as it is always close to 0 (not shown), but vengefulness decreases as the RP increases. More specifically, when the RP is 0.2, increasing PS has little effect, as shown in Figure 5. However, for the other RP values, increasing PS decreases vengefulness. Again, this is due to rewiring breaking down the strong locality effect, and this is magnified with increasing PS, since there is a greater opportunity for connections to other localities, causing a greater cost for agents seeking to bring about norm establishment in all localities at once.

VI. CONCLUSIONS

In this paper we have investigated mechanisms that encourage norms to emerge in communities of self-interested agents, without interference of a central or outside authority, under the realistic constraint that agents can only influence one another if they regularly interact. Based on Axelrod's seminal work, our model's substantial novel extension examines the impact of different types of topologies of interaction on norm emergence. Our results show that in circumstances in which *each* agent regularly interacts with a small number of other agents, as in lattices and small worlds, Axelrod's mechanisms to encourage norm emergence remain effective. Moreover, we have demonstrated that, given fixed penalties, the effectiveness of Axelrod's approach only depends on the number of neighbours of each agent, *not* on the total population size. Thus, topology must be considered: in the case of a lattice or a small world, Axelrod's proposed approach will be effective for sufficiently large neighbourhood sizes.

REFERENCES

[1] R. Axelrod, "An evolutionary approach to norms," *American Political Science Review*, vol. 80, no. 4, pp. 1095–1111, 1986.

[2] B. T. R. Savarimuthu, S. Cranefield, M. Purvis, and M. Purvis, "Norm emergence in agent societies formed by dynamically changing networks," in *Proceedings of Web Intelligence-Intelligent Agent Technology*, 2007, pp. 464–470.

[3] J. Delgado, J. M. Pujol, and R. Sangüesa, "Emergence of coordination in scale-free networks," *Web Intelligence and Agent Systems*, vol. 1, pp. 131–138, 2003.

[4] O. Sen and S. Sen, "Effects of social network topology and options on norm emergence," in *Proceedings of COIN'09*, 2010, pp. 211–222.

[5] D. Villatoro, S. Sen, and J. Sabater-Mir, "Topology and memory effect on convention emergence," in *Proceedings of Web Intelligence-Intelligent Agent Technology*, 2009, pp. 233–240.

[6] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks." *Nature*, vol. 393, pp. 440–442, 1998.