Intelligence Integration in Distributed Knowledge Management

Dariusz Król
Wroclaw University of Technology, Poland

Ngoc Thanh Nguyen
Wroclaw University of Technology, Poland

Information Science REFERENCE

INFORMATION SCIENCE REFERENCE
Hershey · New York

Director of Editorial Content: Kristin Klinger
Managing Development Editor: Kristin M. Roth
Senior Managing Editor: Jennifer Neidig
Managing Editor: Jamie Snavely
Assistant Managing Editor: Carole Coulson
Copy Editor: Lanette Ehrhardt
Typesetter: Jeff Ash
Cover Design: Lisa Tosheff
Printed at: Yurchak Printing Inc.

Published in the United States of America by
Information Science Reference (an imprint of IGI Global)
701 E. Chocolate Avenue, Suite 200
Hershey PA 17033
Tel: 717-533-8845
Fax: 717-533-8661
E-mail: cust@igi-global.com
Web site: http://www.igi-global.com

and in the United Kingdom by
Information Science Reference (an imprint of IGI Global)
3 Henrietta Street
Covent Garden
London WC2E 8LU
Tel: 44 20 7240 0856
Fax: 44 20 7379 0609
Web site: http://www.eurospanbookstore.com

Library of Congress Cataloging-in-Publication Data

Intelligence integration in distributed knowledge management / Dariusz Krol and Ngoc Thanh Nguyen, editors.

p. cm.

Includes bibliographical references and index.

Summary: "This book covers a broad range of intelligence integration approaches in distributed knowledge systems, from Web-based systems through multi-agent and grid systems, ontology management to fuzzy approaches"--Provided by publisher.

ISBN 978-1-59904-576-4 (hardcover) -- ISBN 978-1-59904-578-8 (ebook)

1. Expert systems (Computer science) 2. Intelligent agents (Computer software) 3. Electronic data processing--Distributed processing. I. Krol, Dariusz. II. Nguyên, Ngoc Thanh.

QA76.76.E95I53475 2009

006.3--dc22

2008016377

British Cataloguing in Publication Data
A Cataloguing in Publication record for this book is available from the British Library.

Chapter XI
Mechanisms to Restrict Exploitation and Improve Societal Performance in Multi-Agent Systems

Sharmila Savarimuthu
University of Otago, New Zealand

Martin Purvis
University of Otago, New Zealand

Maryam Purvis
University of Otago, New Zealand

Mariusz Nowostawski
University of Otago, New Zealand

ABSTRACT

Societies are made of different kinds of agents, some cooperative and uncooperative. Uncooperative agents tend to reduce the overall performance of the society, due to exploitation practices. In the real world, it is not possible to decimate all the uncooperative agents; thus the objective of this research is to design and implement mechanisms that will improve the overall benefit of the society without excluding uncooperative agents. The mechanisms that we have designed include referrals and resource restrictions. A referral scheme is used to identify and distinguish noncooperators and cooperators. Resource restriction mechanisms are used to restrict noncooperators from selfish resource utilization. Experimental results are presented describing how these mechanisms operate.

INTRODUCTION

By nature's design different kinds of people exist in a society. Every society has cooperative and uncooperative members. In the real world, it is not possible to get rid of all uncooperative members in the society. It is an injustice to exclude people with certain behaviour from the society. However, it is possible to restrict the performance of uncooperative members and prevent the cooperative members from being exploited. The uncooperative members take advantage of cooperative members by making suckers out of them and also causing damage to the common good. Special mechanisms need to be designed and deployed to control the behaviour of such particular groups, especially in electronic societies such as P2P file sharing and so forth.

BACKGROUND

In previous work (Purvis, Savarimuthu, Oliveira, & Purvis, 2006) we used simple tags and showed the self-organization of cooperative and uncooperative groups. We have now included the referral mechanism into our approach. In this work, we investigate how effective the referral mechanism is in reducing the performance of uncooperative members and increasing the overall society benefit. Additionally, we also show that resource restriction for uncooperative members improves the society benefit.

This chapter is organised as follows. First, we discuss some concepts which are related to our experiments. Then, we explain our experiments which use tags and referral mechanisms. Next, we explain about the experiments which use resource restriction mechanism. Finally, we present our conclusion and future work.

COOPERATIVE BEHAVIOUR IN MULTI-AGENT SOCIETY

For a society to operate effectively, agents within the society must obey certain social rules and norms. So far, much of the focus in this area has been on work devoted to the identification of malevolent agents, where the goal is to identify a noncooperator and exclude it from the society. However, in the real world, it is not going to be applicable in all situations. Our focus is on situations where society members are behaving in an uncooperative manner, but are not necessarily "evil" and deserving of expulsion. This is the issue of the "Tragedy of the Commons" (Hardin, 1968).

Tragedy of the Commons

In Hardin's classic paper (Hardin, 1968), "Tragedy of the Commons," he outlines the "tragedy." A common pasture is open to herders, each of which tries to maintain as many cattle as possible on the commons. A herder will reckon that the positive benefits of adding one additional animal will all go to him, alone, whereas the negative effects from overgrazing of that one additional animal will be shared borne by all the herders. Accordingly, self-interested herders may continue adding one more animal to their herds, even if they know that collectively this is destroying the commons. The question is: how are restrict selfish herders to avoid the tragedy?

The Tragedy of the Commons can be related to the "Prisoner's Dilemma" situation (Axelrod, 1984). Two collaborating criminals are imprisoned and questioned separately. Each criminal may cooperate with his fellow criminal by refusing to divulge details of the crime or defect by ratting on his colleague. It is possible to establish a reward structure (see Figure 1) such that:

- If both criminals cooperate they get a reward, R,
- If they both defect, they are punished (punishment, P),
- If one player defects and the other cooperates, then the defector gets high reward (temptation, T) and the other gets a severe punishment (sucker, S)
- And T > R > P > S, and 2R > T + S
  Under these reward conditions, each individual criminal will reason that if the other
  - Cooperates, he does better by defecting, and if the other
  - Defects, he also does better by defecting.

Thus, the Nash equilibrium situation for this game is for both players to defect, even though they would collectively get a higher reward if they were both to cooperate. The Tragedy of the Commons can be likened to a situation in which the individual herder is playing the Prisoner's Dilemma game against the collection of all the other herders: his selfish interests lead him to defect, even though they are all better off if they cooperate.

Another cooperation game that is discussed in the literature is the Stag Hunt game. The metaphor here is two hunters who may cooperate to hunt a stag (high reward, S). If they operate by themselves, they each can only catch a rabbit (lower reward, R). A hunter seeking to hunt a stag without cooperation gets nothing. But there is no sucker's reward here. The reward structure is shown below (see Figure 2).

Figure 1. Payoff matrix for prisoner's dilemma

Criminals 1 / 2	Cooperate	Defect
Cooperate	R/R	S/T
Defect	T/S	P/P

Power Laws in Network Behaviour

A related issue in networked situations is associated with "power law" behaviour (Barabási, 2002; Huberman, 2001; Shirky, 2003b). Globalised economic environments in which previous barriers to resource access have been greatly reduced exhibit power law properties. For example, Shirky (2003b) has observed that Web logs ranked by number of inbound links closely follows a power law distribution. The spread of telecommunications and the resulting easy media access has meant that a great proportion of earnings made in connection with professional musicians, actors, writers, and other entertainers goes to a small group of people. Whereas it was once the case that almost every town had some musicians that people eagerly wanted to hear, now most people get their music listening pleasure from the electronic media, and traditional skills in instrumental music are increasingly unrewarded and disappearing.

In some circumstances, this power-law situation may be a good thing, but there may be times when it is desirable to achieve a more equable distribution. In particular, a more equable distribution of networked links may lead to a more robust and dynamically adaptable society. It may be the case that preferred nodes "hubs" in a network become single points of failure and risk being choked with the amount of traffic that they must handle. As with the Tragedy of the Commons situation, we may find that there are situations in which there are advantages in introducing some sort of regulations in the traffic, just as we find it necessary to regulate and sometimes restrict vehicular traffic in urban environments.

Figure 2. Payoff matrix for Stag Hunt

Hunters 1 / 2	Cooperate	Defect
Cooperate	S/S	0/R
Defect	R/0	R/R

In our view, this is what needs to be done in connection with agent societies. Mechanisms are needed not only to uncover malefactors, but also to guide and sometimes restrict agent behaviour so that a more cooperative environment is fostered.

P2P File Sharing

P2P file sharing, using Napster, Gnutella, Kazaa, or BitTorrent, is widely engaged in, but uncooperative behaviour is frequently observed. BitTorrent is currently particularly popular, and Hales and Patarin's analysis of BitTorrent's workings (Hales & Patarin, 2005) is of interest. With BitTorrent, groups of users "swarm" with an interest in a specific media file coordinate to speed-up the process. A given file is partitioned into pieces, and each peer is responsible for obtaining and sharing with the other peers some of the pieces. Each swarm is managed by a "tracker," which keeps track of the peers interested in a file or group of files. Peers may query the tracker for a random list of other peers in the swarm, and once obtained, the peers can exchange their piece lists so that they may determine which peers may have pieces that they need. Because a peer may not be able to service at once all the peers that need its piece, it only services up to a limited number of other peers, with remaining peers being left out "choked." Presumably, peers will choose to cooperate with those peers which, in turn, have cooperated with it, and so cooperative behaviour is presumed to be induced by an implicit "tit-for-tat" strategy. But Hales determined that it would be easily possible to cheat under these arrangements and be a "free rider" the bane of all P2P file sharing systems; yet such cheating is not observed in connection with BitTorrent, but is observed on other systems. Why?

Hales's suggested answer to this question (Hales & Patarin, 2005) concerns the way that file metadata is handled with BitTorrent. BitTorrent does not provide a central distribution for metadata; instead the acquisition of metadata is left to the users. To download a file using BitTorrent, one must supply information which can be found in a special .torrent file, but the user must use his or her own devices, such as user-run Web sites, to find this file. This means that the connectivity of this "network" of interested users is not complete: separate, possibly somewhat isolated, groups of users will form and share the metadata. Although this is sometimes thought to be a weakness of BitTorrent, Hales suggests that this may be an advantage, because separate swarms with their individual trackers can be formed for the same file. This can lead to a swarm selection process, whereby higher performing swarms (with more cooperative members) are selected and poorly performing swarms (with more free-riders) are deselected and eventually die off.

The suggested mechanism at work here is that, by means of probably unintended limitations in terms of metadata access, there is an arrangement in BitTorrent that can lead self-interested peers to generate multiple groups and a group-swarm-selection process that ultimately yields more cooperative (and hence higher overall performing) groups. Thus, by having some restrictions in a group, the Tragedy of the Commons can be avoided.

Cooperation Using Tags

Advocates of the Semantic Web envision an IT future in which intelligent agents achieve effective collaboration by employing automated reasoning facilities in connection with rich online ontological information. Others (Doctorow, 2001; McCool, 2005, 2006), have expressed doubts that a realisation of this vision can be practically achieved in the foreseeable future, because the Semantic Web requires too many new tools and constant new data encoding to respond to ever-changing contexts in order to be able to achieve the required take-up and that simpler, lower threshold structures and mechanisms are needed (Shirky, 2003a). One simpler idea is to use simple tags that do not define their semantics, but which are interpreted by application agents for their own particular circumstances (Hales & Patarin, 2005). Since Holland (1993) invented Tags, it has been interest of so many researchers.

Research investigating how cooperation has arisen in biological and social groups, for example, has suggested that simple tagging may provide a better account for the evolution of cooperation than do notions of "tit-for-tat" reciprocity and the "shadow of the future" (Riolo, Cohen, & Axelrod, 2001). In these scenarios, tagging offers a simple mechanism that can facilitate cooperative behaviour on the part of selfish individuals. Individuals just need to like or feel comfortable interacting with other individuals who are readily observed to be like them because they have the appropriate visual tag. This is certainly a natural phenomenon in ordinary human social intercourse.

Feedbacks/Referrals

Information about other members of the society is important and helpful. Before selecting the strategy, getting feedback about the opponent is a rational thing to do. In the model described in Purvis et al. (2006), agents were playing the Prisoner's Dilemma game. They chose a strategy to cooperate or defect by their value of cooperativeness which was assigned randomly when they were created. They did not change their strategy based on their knowledge of the past behaviour. The agents were playing with the nonchanging strategy. The purpose was to simplify the experiments and show the self-organization of cooperative and noncooperative groups achieved by using simple tags.

In this chapter, we have adopted a hybrid approach that uses the concept of tags and recorded history of agent interactions. We change the aspect of nonchanging strategy by allowing the player to ask for feedback about the opponent. By getting the feedback about the opponent, the player can decide whether to cooperate or defect. Here, the feedback about an agent is called the referral. In our approach, there is no lying in the referrals because it is happening within the group. Agents give feedback based on their own observation about the other agent. We do not associate a cost for referrals as we use referrals as a mechanism to improve overall societal benefit. And also, if the referral is positive about the opponent, the agent cooperates. Otherwise, the agent defects to avoid getting a sucker reward.

In the following section, we outline our experiments that take advantage of some of these concepts.

EXPERIMENTS USING TAGS AND REFERRALS

For our experiments, an artificial agent simulation environment has been set up with a society of 100 agents divided into 5 subgroups of 20 agents each. In each subgroup, each agent played 10 games with 19 other agents in its group. Among these 10 games, we call the first five games as first half and the next five games as second half. In the first half, agents play using their value of cooperativeness assigned to them. Every agent keeps the history of the first half which can be referred for the second half, so they know who cooperated or defected with them in the past five games. But they know nothing about how much the other agents might have cooperated with each other. In the second half, each agent asks for referral about the opponent to its best five cooperators of the first half. Among five of them, if at least three of them say that the opponent is a cooperator, the agent will cooperate; otherwise it defects. So each agent plays 190 (19 *10) games in a round. Then, the subgroup monitor will conduct a survey among the subgroup members to determine which is the most cooperative and least cooperative member of each subgroup.

When a player plays with every other member of its subgroup, it computes a cooperation score for each player it played with. The performance of an individual agent is measured by its individual score (not the cooperation score) and is denoted as P.

The score of an individual member is different from its Degree of Cooperation denoted as DC, as voted by its fellow members. For instance, in the Prisoner's Dilemma game, an uncooperative group member may make suckers out of its fellow members and achieve a high score (P), even though being considered least cooperative by its fellow group members (low value of DC). At the end of each round, the voting is performed. Voting is the process of ranking of subgroup members based on their degree of cooperativeness (DC) for that round.

The agents' vote is based on their individual playing experience with other agents. The votes are tallied by the subgroup monitor. Thus, after surveying each subgroup member, the monitor will know its most cooperative (highest DC) and least cooperative (lowest DC) member for that round. The monitor uses this information to kick out the least cooperative member and promote the most cooperative member to other subgroups.

In addition, for each round, the five subgroups are themselves ranked in terms of their Overall Performance (OP), which is the sum of the individual scores of all of its members in the games as given by the formula below.

$$OP = \sum_{p=1}^{p=20} P$$

where P , the score of a player (performance) is given by

$$P = \sum_{g=1}^{g=10} S_g$$

where Sg is the score per game.

To determine movement between subgroups, the procedure given below is followed:

- The highest ranked subgroup in terms of performance (P) kicks out its least cooperative member.
- The 2nd, 3rd, and 4th subgroups in terms of performance (P) also kick out their least cooperative members, but also promote their most cooperative members for movement to a new group.
- The lowest ranked subgroup in terms of performance (P) promotes its most cooperative member for movement to a new group.

There are, thus, eight agents that have been placed into a separated pool for moving to another group: four promoted from the 2nd, 3rd, 4th, and 5th ranked subgroups and four kicked out of the 1st, 2nd, 3rd, and 4th ranked subgroups. Now tagging is employed. Commonly in Hales's (Hales, 2003a, 2003b, 2004a, 2004b; Hales & Edmonds, 2004) and Riolo's(Riolo, 1997; Riolo et al., 2001) work, all the agents are tagged and the tags serve the purpose of showing the identity of the agent and also specifying which group the agent belongs to. Here, our purpose of using tags is just to represent the status of an agent which is currently in the pool. The status could be high (promoted) or low (kicked out). The four promoted agents are given blue tags, signifying promotion, and the four kicked out agents are given red tags, signifying demotion. The monitor agent chooses players with blue tags in preference to players with red tags without knowing the performance scores of the players in the pool. The monitor agent takes players with blue tags if they are still available when it comes to its turn to choose.

- At this stage, the highest performance ranked subgroup gets one agent among the pool members in order to replace the member that has been kicked out (1st group gets 1 blue tagged agent).
- Then, the 2nd, 3rd, and 4th ranked subgroups get two agents to replace the two agents that they have lost (2nd group gets 2 blue tagged

agents, 3rd group gets a blue tagged and a red tagged agent and the 4th group gets 2 red tagged agents).

- Finally, the lowest ranked subgroup winds up with the remaining agent of the pool (5th group gets a red tagged agent).

With the newly created subgroups, another round of play is initiated. As mentioned before, individual game-playing agents are programmed to cooperate or defect with a tendency determined by a constant cooperation threshold parameter value of cooperativeness, CT that was randomly initialised to have a value between 0 and 1. Then, when each game was played, a random number was selected, and if it was less than CT, the player defected on that occasion, while if it was equal to or above CT, the player cooperated. At the outset of the game, each of the five subgroups populated with a random collection of 20 players having various tendencies to cooperate or defect within the group. Remember, only in the first half of every game, all the agents use the assigned constant cooperation threshold value to select the strategy. In the second half, they use referral scheme to select the strategy. The goal of the experiments is to show how well the performance of noncooperators can be restricted and society benefit can be improved. An agent is considered to be cooperative if DC is at least over a threshold value CT. In our case we set CT = 0.5. That corresponds to having a degree of cooperation (DC) such that: $DC \geq 0.5$. For instance, if an agent has DC=0.4, it cooperates 4 times out of 10 times. Because its DC value is less than CT, it is considered to be uncooperative.

When the Prisoner's Dilemma game was played over 100 rounds, initially the groups were approximately having equal number of cooperators and defectors in every group. But after about 15 rounds, we could clearly distinguish different groups:

- Two groups had almost all cooperators;
- Two other groups had mostly defectors; and
- A middle group that had about half cooperators and half defectors.

Thus, these groups self organise themselves. This kind of emergence falls on the category of Type 3 according to Fromm (2005). In his chapter, Formm has explained comprehensive classification of the major types and forms of emergence in Multi-agent Systems. According to his classification, the result we got shows multiple emergence because it is influenced by several factors like tagging, referrals and assigned value of cooperativeness. This is shown in Figure 3.

With the separation of groups (based on their behaviour) we compared the scores (see Table 1).

Our general observation when using the referral scheme is that the groups who have more cooperators in the population gain a higher score than what it scored in the first half. And the uncooperative group (full of defectors) scores less than what it scored in the first half. In such groups, by

Figure 3. Multiple emergence influenced by tagging, referrals and assigned value of cooperativeness

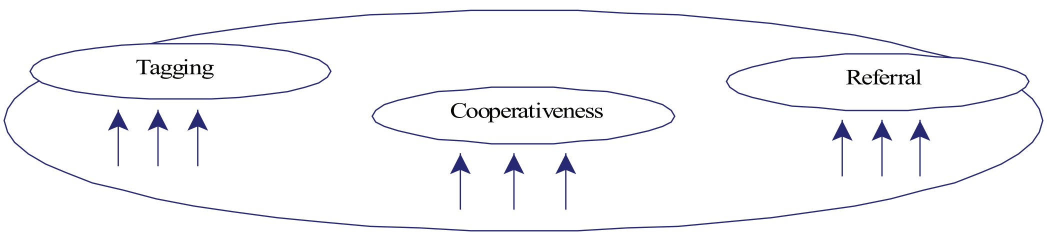

Table 1. Different groups and their performance in PD game

Groups (ranked by OP)	1	2	3	4	5
C : D	18:2	16:4	10:10	4:16	0:20
P in Firsthalf	3397	3293	2961	2688	2247
P in Secondhalf	3531	3374	3147	2636	2174

OP–Overall Performance C-Cooperators
D-Defectors P-Performance

using referrals, every agent comes to know that the opponent is a defector, so they all defect each other and get a low score. Because of this, their group score also goes low. The performance of a group which has an equal number of cooperators and defectors is determined by its overall tendency to cooperate/defect. If the tendency to cooperate is higher than defection, the group increases its score from its first half score. This is also the same for the groups who have a very small difference in the number of cooperators and defectors. For instance, a group with 8 defectors and 12 cooperators can get a lesser score in the second half, if the groups overall tendency to defect is more than to cooperate. For the first three groups, the score increased in the second half. But for the last two groups, the score decreased because they are mostly defectors (see Figure 4).

Also, we observed that the sum of scores of all groups in the second half is always higher than that of the first half's (see Figure 5). This shows that the referral scheme works well to improve the scores which is good for the overall society.

As the second experiment, the Stag Hunt game was played (S = 9, R = 7),over 100 rounds in the same manner as the previous PD game experiment. This also showed similar results as PD game in score increase while using referral (similar to Figure 5).

With the separation of groups, we compared the scores for Stag Hunt (see Table 2). Here, all the groups showed a score increase in the second

Figure 4. Performance of groups for PD game

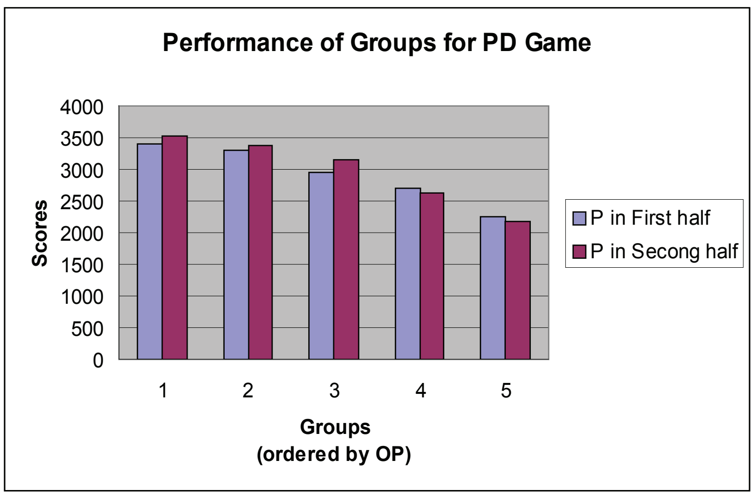

Table 2. Different groups and their performance in Stag Hunt game

Groups (ranked by OP)	1	2	3	4	5
C : D	18:2	15:5	10:10	1:19	3:17
P in First half	13777	12126	11156	11022	10665
P in Second half	15172	12225	11504	11322	11398

OP–Overall Performance C-Cooperators
D-Defectors P-Performance

Figure 5. Comparison of scores for first half (without referral) and second half (with referral) is shown for 10 runs

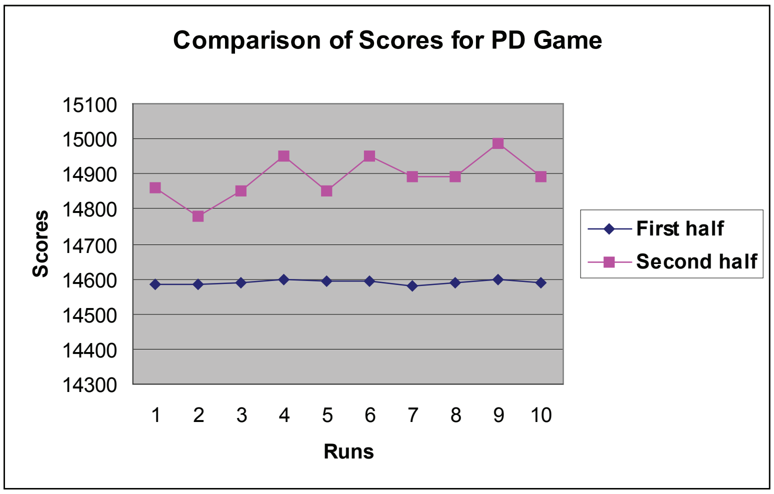

half. It is normal that the first 3 groups show score increase. But the reason why the last two groups also showed score increase is that there is no sucker reward in the reward structure of the Stag Hunt game. One cannot gain at the others' expense. So, in the defectors' groups, when both select to defect they both get a better score than being one of them defected (the score for Defect (7) + Cooperate (0) = 7 and the score for Defect (7) + Defect (7) = 14).When both defect it improves the group score as well. The performance of the groups in the first half and second half is shown (see Figure 6).

RESOURCE RESTRICTION

For agent societies to operate effectively, all agents within the society must play by rules. When some of them do not abide the rules and exploit the common resource, they must be restricted to do so. A related issue is that of Internet access given to students for academic purposes. When some of them violate the rule by downloading music/videos, the speed goes down and everyone is affected by this. To deal with this problem, the Internet access for these exploiters (students) must be limited. Our experimental results describe this.

Figure 6. Performance of groups for Stag Hunt game

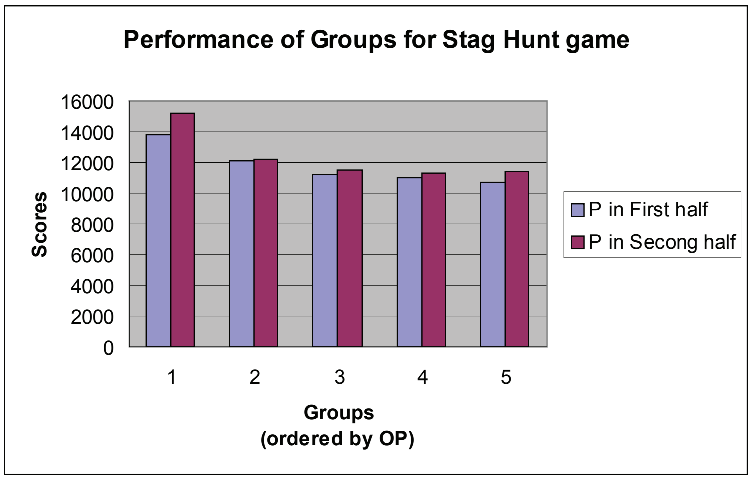

Resource Restriction for Groups

We tried this experiment in the same environment as the previous PD game experiment. There are 100 agents divided into 5 groups each having 20 agents. We made the agents in each group play five games per round (just like the first half of the previous PD game experiment, but no history is stored here). They played 100 rounds. They played according to how they were randomly assigned at the beginning.

After few rounds (about 15 rounds) the groups started separating. Out of five groups, they separated two groups full of mostly cooperators, two groups full of mostly defectors and a middle group having half of both. Now the OP is calculated for each group and the groups were ranked based on that. And also the overall average (sum of the average score of all five groups) is calculated. We call this average *the average without resource restriction*.

The number of games each agent plays is treated as the resource. We want to limit the resource (number of games) to the groups who are not performing up to the standards. The limitation for using the resource varies to groups according to their performance. The 1st ranked group is allowed to play five games as usual, because it is the best performer. The 2nd ranked group is restricted to play four games, the 3rd ranked group to play three games, the 4th ranked group to play two games and the 5th ranked group to play just one game, which is the worst performer among all. Now, again, the overall average is calculated. We call this *the average with resource restriction.*

When we compared the performance (see Figure 7), the total score with resource restriction is less than without resource restriction, because of less number of games. But the overall average is higher. This shows that resource restriction for exploiters is good for society benefit.

Resource Restriction for Individual

In another experiment using the same experimental set up as the previous one, we applied resource restriction for individuals who are uncooperative. Unlike the previous experiment, here the past history is stored. A player can decide whether to play or not with a particular opponent, based on past history. The player calculates how many times a particular opponent cooperated and defected. If the opponent cooperated more times, then the player plays with the opponent; otherwise the player opts not to play the game.

Figure 7. Comparison of performance

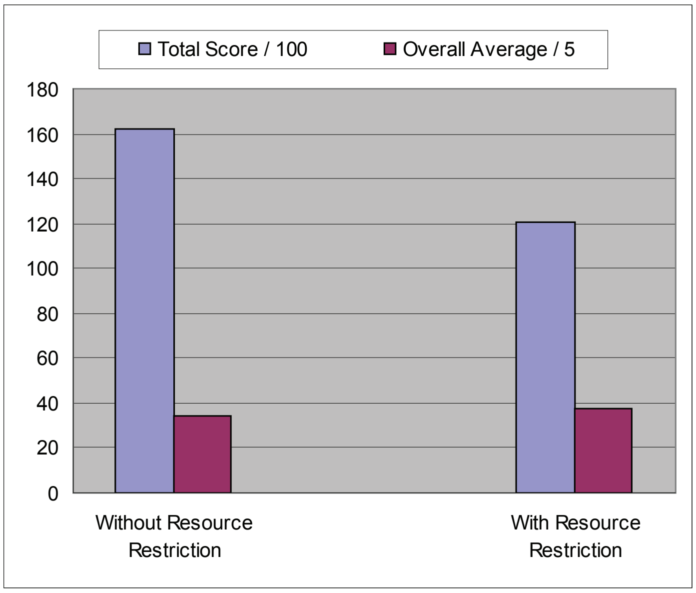

Figure 8. Self-organization of groups based on individual resource restriction is shown for 25 runs

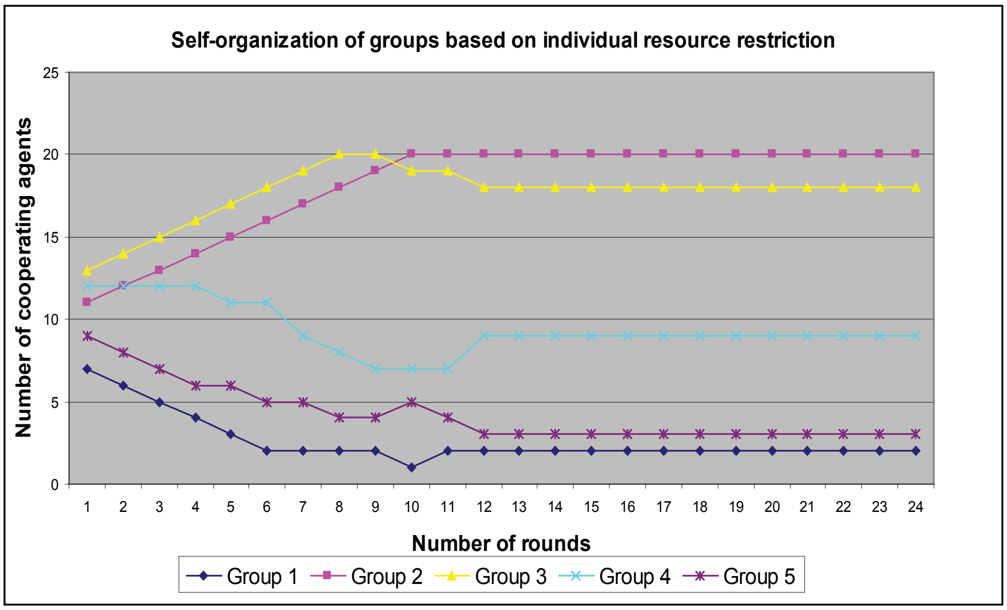

By this mechanism, the groups with a greater number of cooperators score more and the groups with a greater number of uncooperators score less.

In the very first experiment of this chapter, we observed the convergence around 15 rounds. In this experiment, the convergence takes place faster. The groups started separating around 7-10 rounds (see Figure 8).

The emergence of self-organization when using individual resource restriction leads to faster convergence than the mechanism that used referral because this mechanism uses personal history of all the past interactions with other agents, while the referral mechanism used the history of only five games.

CONCLUSION AND FUTURE WORK

In this work, we have described mechanisms that help to resist exploitation and improve overall societal performance in a multi-agent society that has cooperative and uncooperative agents. Our first mechanism uses tags and referrals. We have demonstrated that our proposed mechanism helps in improving the society benefit when the agents play the PD game as well as the Stag hunt game. We have also proposed resource restriction mechanisms and have demonstrated that the group level resource restriction and the individual level resource restriction can be used to improve the overall societal performance.

In the future, we are planning to use the concepts of tags and referrals and propose mechanisms for real life electronic societies (similar to orkut), so that cooperative members can be encouraged and uncooperative members can be restricted from social interactions.

REFERENCES

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.

Barabási, A. L. (2002). *Linked: The new science of networks*. PA: Perseus Books Group.

Doctorow, C. (2001). *Metacrap: Putting the torch to seven straw-men of the meta-utopia* [Electronic Version]. Version 1.3. Retrieved April 3, 2008, from http://www.well.com/~doctorow/metacrap.htm

Fromm, J. (2005). *Types and forms of emergence* [Electronic Version]. Adaptation and self-organizing systems. Retrieved April 3, 2008, from http://www.citebase.org/abstract?id=oai:arXiv.org:nlin/0506028

Hales, D. (2003a). *Evolving specialisation, altruism and group-level optimisation using tags* (Vol. 2581/2003, Lecture notes in computer science). Berlin/Heidelberg: Springer-Verlag.

Hales, D. (2003b, September 16-19). Understanding tag systems by comparing tag models. In *Paper presented at the to-Model Workshop (M2M2) at the Second European Social Simulation AssociationConference (ESSA'04)*, Valladolid, Spain.

Hales, D. (2004a). Change your tags fast!--a necessary condition for cooperation?. In *Paper presented at the The Multi-agent-based Simulation (MABS 2004)*, (pp. 89-98).

Hales, D. (2004b). Self-organising, open and cooperative P2P Societies--from tags to networks. In *Paper presented at the Engineering Self-Organising Applications (ESOA 2004)*, (pp. 123-137).

Hales, D., & Edmonds, B. (2004). Can tags build working systems?—From MABS to ESOA. *Lecture notes in computer science* (pp. 186-194).

Hales, D., & Patarin, S. (2005). *How to cheat BitTorrent and why nobody does. UBLCS* (Tech. Rep. No. UBLCS-2005-12).

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

Holland, J. H. (1993). *The effect of labels (tags) on social interactions* (Vol. SFI Working Paper 93-10-064). NM: Santa Fe Institute.

Huberman, B. A. (2001). *The laws of the Web: Patterns in the ecology of information*. Cambridge: MIT Press.

McCool, R. (2005). Rethinking the Semantic Web, part 1. *IEEE Internet Computing, 9*(6), 86-88.

McCool, R. (2006). Rethinking the Semantic Web, part 2. *IEEE Internet Computing, 10*(11), 93-96.

Purvis, M. K., Savarimuthu, S., Oliveira, M. D., & Purvis, M. A. (2006). Mechanisms for cooperative behaviour in agent institutions. In *Paper presented at the Intelligent Agent Technology (IAT 2006)*, Hong Kong.

Riolo, R. L. (1997). *The effects of tag-mediated selection of partners in evolving populations playing the iterated prisoner's dilemma* (Paper No. 97-02-016). NM: Santa Fe Institute.

Riolo, R. L., Cohen, M. D., & Axelrod, R. (2001). Cooperation without reciprocity. *Nature, 414*, 441-443.

Shirky, C. (2003a). *Permanet, nearlynet, and wireless data* [Electronic Version]. Economics & Culture, Media & Community, Open Source. Retrieved April 3, 2008, from http://www.shirky.com/writings/permanet.html

Shirky, C. (2003b). *Power laws, Web logs, and inequality* [Electronic Version]. Economics & Culture, Media & Community, Open Source, Version 1.1. Retrieved April 3, 2008, from http://shirky.com/writings/powerlaw_weblog.html