Emergence of stable coalitions via negotiation of task exchanges

Partha Sarathi Dutta and Sandip Sen
Mathematical & Computer Sciences Department
University of Tulsa
partha@euler.mcs.utulsa.edu, sandip@kolkata.mcs.utulsa.edu

Abstract

Autonomous agents interacting in an open world can be considered to be primarily driven by self interests. Previous work in this area has prescribed a strategy of reciprocal behavior for promoting and sustaining cooperation among self-interested agents. In an extension to that work the importance of using other's "opinion" by the reciprocative agents to effectively curb the exploitative tendencies of selfish agents has also been evaluated. However, one of the restrictions so far has been the metric of "cost" evaluation. So far we have considered only the time of completion of a task to be the measure of the task cost. In this paper, we expand the cost metric by including both time of delivery and quality of work. This allows us to model more realistic domains. In contrast to our previous work, we also allow for agents with varying expertise for different job types. This necessitates the incorporation of the novel aspect of learning about other's capabilities within the reciprocity framework. In this paper, we evaluate the hypotheses that self-interested agents with complementary expertise can learn to recognize cooperation possibilities and develop stable, mutually beneficial coalitions that is resistant to exploitation by malevolent agents.

Introduction

Agent-based systems have been an active area of research for the past decade both in the academia and industry. Agent-based systems are an important aspect of real world applications like electronic commerce, recommender systems and personal assistants. Agents deployed in these applications often interact in an open environment with other agents or humans (Bradshaw 1997; CACM July 1994 issue 1994; CACM March 1999 issue 1999; Huhns & Singh 1997). The interactions involve cooperation, collaboration or competition for resources to achieve the specified goals of these agents. With increase in the complexity of agent interactions, the behavioral characteristics of agents acting in a group should be studied more deeply and suitable interaction strategies developed that optimize system performance.

We have been interested in agent strategies for interactions with other agents that can promote cooperation in groups of self-interested agents. Our approach is different from other researchers who have designed effective social laws that can be imposed on agents (Shoham & Tennenholtz 1992). We assume that typical real-world environments abound in *cooperation possibilities*: situations where one agent can help another agent by sharing work such that the helping cost of the helper is less than the cost saved by the helped agent. As agent system designers we can define rules of interactions to increase the likelihood of cooperation possibilities. We prescribe reciprocative behavior as a stable sustainable strategy that creates cooperative groups in the society. This behavior not only sustains group formation in a homogeneous society of self-interested agents, but also helps to ward off exploitative tendencies of selfish agents in the society (Biswas, Sen, & Debnath 2000). This strategy of reciprocal behavior becomes more stable if the helping agent takes into account the opinion of all other agents before extending any favor (Biswas, Debnath, & Sen 2000).

A restriction of the previous work on reciprocity was the incorporation of only a single cost metric, time, used by the agents. In real-life scenarios multiple objectives like time, quality, dependability, etc. will be involved when an agent evaluates the benefit of interacting with another agent. As a first step to such a scenario, we expand on the set of cost metrics by including both time and quality in an agent's evaluation scheme. The measures of *time* and *quality* of a work need clarification. The *time* refers to the absolute time units required for completing a particular task, and the *quality* refers to the required quality of performance to complete a particular task. Naturally, these values will depend on the types of the tasks.

A second restriction in the previous work was the explicit assumption that all agents have the same capabilities. This was done deliberately to focus on the help-giving behaviors of the agents. Having established a basic competence of probabilistic reciprocity based agents in recognizing mutually cooperative relationships and effectively neutralizing exploitative behavior, we now turn our attention to the interesting aspect of variance in agent expertise. We assume that different agents have different skill sets which make them more effective in accomplishing some tasks versus others. We require agents to learn the capabilities of themselves and others through performance and interaction.

The goal of this work is to evaluate whether self-interested agents can learn to recognize agents with complementary expertise and develop a self-sustaining relationship through exchange of help. This can be described as an augmentation of the basic probabilistic reciprocity model with the concept of *selecting a partner*. Additionally, such help exchange inclinations must be resistant to exploitation by malevolent agents who do not reciprocate help-giving actions. Our hypothesis is that when combined with an appropriate learning scheme, probabilistic reciprocity based strategies will enable the development of stable, mutually beneficial coalitions of self-interested agents with complementary skill sets.

Reciprocity

Altruistic behavior has been the topic of research for mathematical biologists and economists. The objective of such research has been to verify the rationality of adopting such behavior, to analyze the success of altruistic individuals in an environment of self-interested agents and to study the evolution of genetic patterns supporting such behavior among communities (Badhwar 1993; de Vos & Zeggelink 1994; Goranson & Berkowitz 1966; Krebs 1970; Macy & Skvoretz 1998; Schmitz 1993; Trivers 1972). The concept of reciprocal altruism has been evaluated to explain the emergence of stable cooperative groups among a society of interacting agents (Zeggelink, de Vos, & Elsas 2000). Majority work in the field of game theory have evolved around the idealized problem called Prisoner's dilemma (Rapoport 1989). Robert Axelrod has shown using a simple deterministic reciprocal strategy how stable cooperative behavior can be achieved in self-interested agents (Axelrod 1984).

Even though tit-for-tat strategy explains the stability of reciprocal behavior in the repeated Prisoner's dilemma tournament framework, the applicability of this analysis in more realistic multiagent environments is limited by unrealistic assumptions like symmetrical interactions, repetition of identical scenarios, lack of a measure of work, etc. Hence, such naive schemes cannot be used as a robust solution approach for encouraging the emergence of stable, mutually cooperative, groups in a real-life, open agent society. We have proposed a *probabilistic reciprocative* approach where the probability of extending help to another agent depend on the interaction history with that agent (Sen 1996). Our probability-based reciprocal behavior is designed to enable an agent make decisions about agreeing to or refusing a request for help from another agent. The probability that agent k, having task l, will help agent i to do task j is given by

$$Pr(i,k,j,l) = \frac{1}{1+\exp^{\frac{C_{ij}^{k}+OP-\beta * C_{avg}^{k}-B_{ki}}{\tau}}},$$

where C_{avg}^{k} is the average cost of tasks performed by agent k; B_{ki} is the net balance that agent k has with i, due to the previous helps given to and received from agent i; OP is the balance that agent i has with other agents excluding agent k; C_{ij}^{k} is the cost of agent k to do the task j of agent i; β and τ are constants where β is used to set the cost an agent is ready to incur to help an unknown agent with the hope of initiating a long-term cooperative relationship and τ is used to the shape of the probability curve. This is a sigmoidal probability function where the probability of helping increases as the balance increases and is more for less costly tasks. We include the C_{avg} term because while calculating the probability of helping, relative cost should be more important than absolute cost.

In this paper we aim at extending the basic probabilistic reciprocity framework. We evaluate the importance of "quality of performance" as another cost metric and the applicability of learning other agents' capabilities in developing stable, cooperative coalitions among self-interested agents.

Problem domain

We evaluate our hypothesis using simulations in a job completion problem domain. In this domain each of N agents are each assigned T jobs. There are K job types and each agent has expertise in one of these K job types. An agent who is an "expert" in a particular job type can do the job in less time and with higher quality than other job types. The jobs can be finished at any one of F different machines. Agents are assigned their jobs when the simulation starts. At this point every agent asks for help from some *selected* candidates among all other agents. Initially the probability of selecting an agent for asking help is the same for all agents. This is because the asking agent has no prior model of the other agents' behaviors. However, as agents learn over interactions, some agents become more likely to be selected by an agent when the latter is seeking help. The simulation ends when every agent has finished all of their assigned jobs.

Selecting partners

We have incorporated the concept of agents learning to select another agent for asking help. Agents have estimates of their own abilities to do the different job types. Estimates are of two types: *time estimate*, which reflects the possible time of completion of the job, and *quality estimate*, which reflects the possible performance level of an agent to do that job. Agents also keep estimates of other agents' abilities. Agent estimates are initialized to zero. To obtain estimates about their own abilities, agents must themselves finish a job of a certain type if it has never encountered another job of that type before. Every job requires a certain time and quality of performance to be finished. These task-specific values are used by the agents to measure their performance. Estimates about another agent are updated using the time and quality of performance values that the other agent requires to complete the particular type of job.

These estimates are used by the agents to compute the cost of a particular task delivery. When asking for help, agents compute the cost $C1$, incurred by itself to do that task. The estimated cost $C2$ of the prospective helping agent for that work is also computed. Help is obtained only if $C2 < C1$. One important issue in determining the estimated cost $C2$ is that, if the task type of the asking agent's present task is different from that of the task the helping agent is executing, then an additional cost factor is added to $C2$. This takes into account the cost incurred by the helping agent to change the setup of a machine for a new kind of job.

Once an agent does a particular task of a certain type, it updates its self-estimate about that task type. If an agent is helped by another agent, it updates its estimate of the helping agent's ability to do that task type. This ensures a reinforcement learning of one's own abilities and those of others for every job type. Given sufficient interaction possibilities between agents, they learn the behavioral characteristics of other agents. This information is used to select an agent for asking help. Naturally, as a consequence of learning others' abilities, an agent i who specializes in job type T_1, when given a job of type T_2, will decide in favor of asking help from an agent j specializing in job type T_2, rather than agents who specialize in other job types. Thus, given sufficient interaction possibilities, agent coalitions develop where self-interested agents have complementary skill sets.

Agent types

The behavior of reciprocative agents is determined by the level of available information to the agents about other agents. We mentioned that agent A will receive help from agent B only if the cost incurred by A to do that task is more than incurred by B for helping A. However, in deciding whether to help or not to help, the reciprocative agents can ask *some* other agents for their *opinion* about the agent asking help. The agents from whom a help-giving agent asks for opinion are only those with whom it has a favorable balance, that is, from those it has received more help than given away. This is a reasonable policy, to believe in those who have *earned trust* by their interactions. The opinion received by the helping agent is the cumulative balance that the agent (who is seeking help) has with the agents from whom opinion is obtained. This reduces if the help-seeking agent has received more help than given away to those giving opinion, and vice versa. An agent who has a "good" record of help giving behavior is more likely to be helped. Reciprocative agents who decide to help or not based on such collaborative opinion are called *earned trust based reciprocative agents.*

Selfish agents by definition do not extend help under any circumstances. This basic premise holds in our selfish agents too. However, they are characterized by the additional property of *lying* when they are approached for opinion by a reciprocative agent. They lie about the balance they hold at any time with the agent who is currently waiting to receive help. More formally, if the agent asking help is A, that deciding whether to help or not is B and asking a selfish C for opinion, then C says "A has a positive balance of X units with me", where X is actually the positive balance that C has with A, because A helped C on previous interactions. Thus the selfish agents add more damage, and we call such agents as *individual lying selfish agents.*

Experimental results

In the first set of experiments we monitored the average savings obtained by the reciprocative agents. *Savings* earned by an agent is the cost it saves by giving up the job to another agent. Higher the savings, better is the agent performance and vice versa. *Savings* can be a measure of time or resource saved by giving up a job to another agent. Our results show that with the increase in the number of per-agent tasks, the average savings of reciprocative agents increase, given the same percentage distribution of selfish and reciprocative agents in the population. However, given the same value of per-agent task level, the average savings of reciprocative agents decrease as the percentage of selfish agents increases in the population (see Figure 1). Increase in average savings with the increase in per-agent jobs is due to more interactions between reciprocative agents that allowed them to identify helpful agents with complementary expertise. The above experiments were conducted by varying the selfish percentage in the population from 50 to 80 in steps of 10. For each such setting, the per-agent task was varied from 20 to 100 in steps of 20.

In the second set of experiments we studied the variations in average savings earned by both selfish and reciprocative agents after each agent had completed all of its jobs. Figures 2 and 3 show the results obtained by varying the percentage of selfish agents in the population under two different values of "tasks per agent", 20 and 40 respectively. The percentage of selfish agents in the population was varied from 50 to 80% in steps of 10%. In both the figures we find reciprocative agents outperform selfish agents, because the average savings earned by the reciprocative agents are more than those of the selfish agents for all selfish percentages in the population. In both cases the savings earned by reciprocatives decrease with increase in the percentage of selfish population. For the same percentage of selfish population, however, the savings earned by the reciprocatives are higher when the per-agent tasks are higher (40 in this case). From these observations, we can conclude that, reciprocative agents learn to select partners and this enables them to enter into mutually beneficial coalitions. Savings earned, which reflects how good the behavior of the agent has been during its interactions

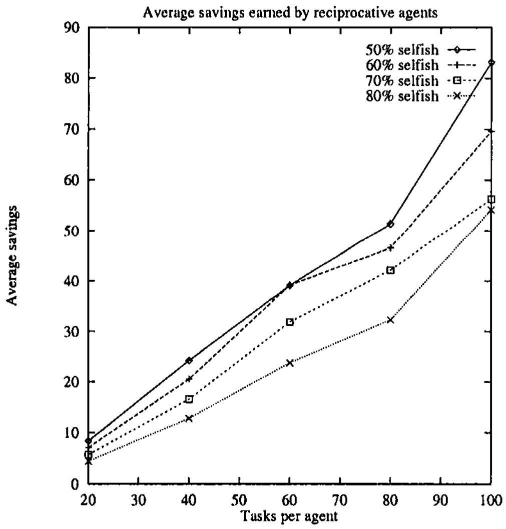

Figure 1: Variation of average savings with per-agent task.

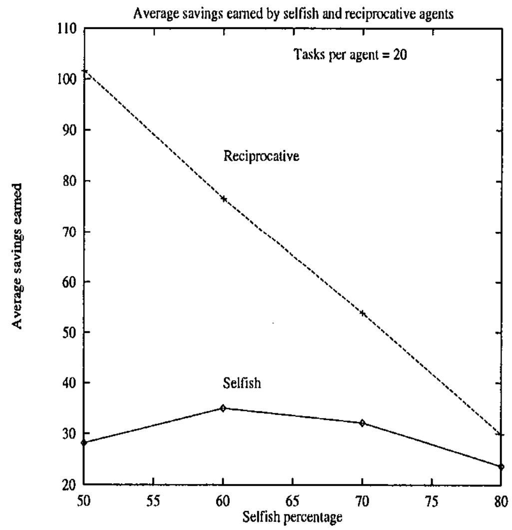

Figure 2: Average savings of agents with 20 tasks per agent.

Groups	1,1	2,2	3,3	1,2	1,3	2,3
Group Size	20	28	12	60	38	47

Table 1: Size of stable coalition between different agent types.

with other agents (more a reciprocative agent helps another reciprocative agent, more is the probability that it will receive help from the one it helped), is therefore, higher in reciprocatives than selfish agents. Lower values of savings earned by the selfish agents indicate that the reciprocative agents have effectively shunned the exploitative selfish agents. The effectiveness, however, decreases with increase in the number of the selfish agents. This is due to the lesser probability of interaction with another reciprocative agent, than with another selfish agent. With increase in the task level, however, the savings of reciprocative agents for the same number of selfishes in the population are more because a higher number of per agent tasks makes it necessary for the agents to stay "on the field" for a longer time. This has a similar effect to what we obtained with a higher number of reciprocative agents compared to a lower value - the average savings earned by the reciprocatives are more.

In the final set of experiments, we studied the nature of coalitions developing among reciprocative agents. In the current context, viable coalition formation depends both on identifying potentially beneficial partnerships and then nurturing them through mutually beneficial exchanges of help. The complementarity of expertise is a useful indicator of the potential of a mutually beneficial partnership. We monitor the development of agent groups where all agents are inclined to help each other. Our results show that groups of agents with complementary expertise have larger sizes. These experiments were run using 100 tasks per agent and an equal distribution of selfish and reciprocative agents in the population.

Agents are differentiated on the basis of the types of job they are most proficient in. *Groups* are defined as coalitions of agents of such differing expertise. To determine a group, say between reciprocative agents who are proficient in job-types $T1$ and $T2$, we compute the number of agents specializing in job-type $T2$ who have helped an agent specializing in type $T1$. This is calculated for all agents specializing in type $T1$. This calculation is again computed with the roles (helper and helped) of agent types $T1$ and $T2$ reversed. The sum of the two results gives the *group size* of agents with expertise in types $T1$ and $T2$. The values of such group sizes are given in Table 1. The numbers within braces are the job-types in which agents have expertise. Thus, group $\{x, y\}$ is the group of agents who specialize in job types x and y.

These results show that the savings of groups of reciprocative agents with complementary expertise are larger than groups of agents with similar expertise. These results indicate that reciprocative agents tend to form stable and beneficial groups with other reciprocative agents who have complementary expertise.

Our simulations have shown that in a heterogeneous group of agents with varying expertise for different job types, reciprocative agents can learn to detect agents with complementary expertise and form mutually beneficial and stable groups. An extension of this work can be the inclusion of dynamically changing behavior of

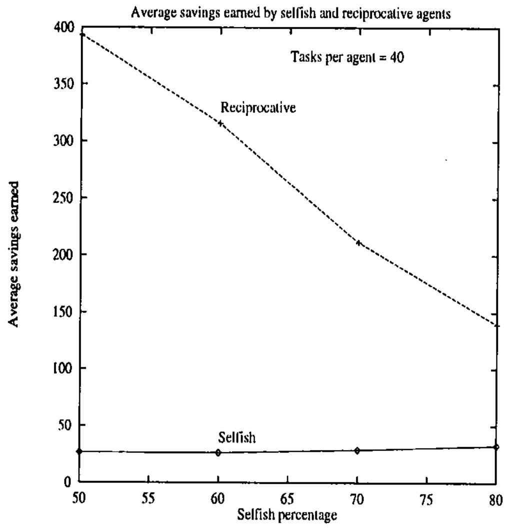

Figure 3: Average savings of agents with 40 tasks per agent.

agents, that might evolve one of the strategies as the stable strategy.

Related work

Using agent coalitions is an interesting and much cultivated methodology to solve complex, distributed tasks. Assigning groups of agents to do a task or multiple tasks has the advantage of individual agent expertise being used to complete a part of the global problem.

However, work in the area of coalition formation in agent societies has focussed on cooperative agents (Shehory & Kraus 1998). A related work on coalitions in agent societies takes self-interested agents into account (Lerman & Shehory 2000). But it does not consider the possible heterogeneity in performance of a task between different agents. Our work is different from these because it takes into consideration self-interested agents that have different performance levels for different task types. We also have a learning parameter in our agents which is used to identify other's qualities. This, in turn, favors those with complementary expertise when asking for help. Mutual understanding of capabilities evolve cooperative groups in the society of self-interested agents. Learning of cooperative behavior has also been addressed in (Denzinger & Krodt 2000). However, there the online learning uses an a priori off-line learning module which generates situation-action pairs using GA. To determine the fitness of the evolved situation-action pairs during off-line learning an agent-model is required. This is obtained by observing other's behavior, which implies learning once more, or the model can be stored in an agent's data area. Our work uses only online learning and agents do not have to store a priori models of other agents. However, in an attempt to learn behaviors of all other agents, our approach can become time consuming in situations where number of agents increase exponentially. A concept of using "congregations" to determine "who" to interact with in an agent society is addressed in (Brooks, Durfee, & Armstrong 2000). However, determining the right congregation again suffers from the time issue when number of agents increase exponentially. Besides, congregations are not true coalitions.

Conclusions and future directions

Adoption of reciprocal behavior in a society of self-interested agents generates incentives that promote cooperative behavior. In this paper we hypothesized that in an agent society agents differ in their expertise, incorporation of a learning component will help generate cooperative groups of agents with complementary expertise. Our simulations confirmed this hypotheses. These groups are mutually beneficial to the agents because they save more time and get their jobs finished with better quality of performance by giving it to another agent who has the required expertise. Such cooperative group formation is stable since the exploitative agents get shunned effectively.

We are still running further set of experiments and plan to present a more extensive set of results detailing the gain in quality, completion time, etc. for groups with complementary expertise.

One assumption in this work has been that agents have fixed behaviors. A more realistic scenario would be for an agent to have the freedom of changing its behavior when it deems appropriate. Such behavior adoption leads to an evolutionary process with a dynamically changing composition of agent group behaviors. We plan to investigate whether such evolving agent behavior might create a society of agents with the same behavior, which could be the optimal behavior to adopt.

Acknowledgments This work has been supported in part by an NSF CAREER award IIS-9702672.

References

Axelrod, R. 1984. *The Evolution of Cooperation.* Basic Books.

Badhwar, N. 1993. Altruism versus self-interest: sometimes a false dichotomy. *Social Philosophy & Policy* 10(1):90–117.

Biswas, A.; Debnath, S.; and Sen, S. 2000. Believing others: Pros and cons. In *Proceedings of the Fourth International Conference on Multiagent Systems*, 279–285.

Biswas, A.; Sen, S.; and Debnath, S. 2000. Limiting deception in groups of social agents. *Applied Artificial Intelligence (Special Issue on Deception, Fraud, and Trust in Agent Societies)* 14(8):785–797.

Bradshaw, J. M. 1997. *Software Agents.* Menlo Park, CA: AAAI Press/The MIT Press.

Brooks, C.; Durfee, E.; and Armstrong, A. 2000. An introduction to congregating in multi-agent systems. In *Proceedings of the Fourth International Conference on Multi-Agent Systems*, 79–86.

1994. Communications of the ACM, july 1994, volume 37, number 7. Special Issue on Intelligent Agents.

1999. Communications of the ACM, march 1999, volume 42, number 3. Special Issue on Multiagent Systems on the Net and Agents in E-commerce.

de Vos, H., and Zeggelink, E. 1994. The emergence of reciprocal altruism and group-living: an object-oriented simulation model of human social evolution. *Social Science Information* 33(3):493–517.

Denzinger, J., and Krodt, M. 2000. Evolutionary online learning of cooperative behavior with situation action pairs. In *Proceedigns of the Fourth International Conference on Multi-Agent Systems*, 103–110.

Goranson, R., and Berkowitz, L. 1966. Reciprocity and responsibility reactions to prior help. *Journal of Personality and Soial Psychology* 3(2):227–232.

Huhns, M. N., and Singh, M. P. 1997. *Readings in Agents*. San Francisco, CA: Morgan Kaufmann.

Krebs, D. 1970. Altruism – an examination of the concept and a review of the literature. *Psychological Bulletin* 73(4):258–302.

Lerman, K., and Shehory, O. 2000. Coalition formation for large-scale electronic markets. In *Proceedings of the Fourth International Conference on Multi-Agent Systems*, 167–174.

Macy, M. W., and Skvoretz, J. 1998. The evolution of trust and cooperation between strangers: a computational model. *American Sociological Review* 63:638–660.

Rapoport, A. 1989. Prisoner's dilemma. In Eatwell, J.; Milgate, M.; and Newman, P., eds., *The New Palgrave: Game Theory*. London: Macmillan. 199–204.

Schmitz, D. 1993. Reasons for altruism. *Social Philosophy & Policy* 10(1):52–68.

Sen, S. 1996. Reciprocity: a foundational principle for promoting cooperative behavior among self-interested agents. In *Proceedings of the Second International Conference on Multiagent Systems*, 315–321. Menlo Park, CA: AAAI Press.

Shehory, O., and Kraus, S. 1998. Methods for task allocation via agent coalition formation. *Artificial Intelligence* 101(1-2):165–200.

Shoham, Y., and Tennenholtz, M. 1992. On the synthesis of useful social laws for artificial agent societies (preliminary report). In *Proceedings of the National Conference on Artificial Intelligence*, 276–281.

Trivers, R. 1972. The evolution of reciprocal altruism. *Quarterly Review of Biology* 46:35–57.

Zeggelink, E. P.; de Vos, H.; and Elsas, D. 2000. Reciprocal altruism and group formation: The degree of segmentation of reciprocal altruists who prefer 'old-helping-partners'. *Journal of Artificial Societies and Social Simulation* 3(3):212–15.