

Group interactions based on agent opinion

Parijat Prosun Kar, Partha Sarathi Dutta and Sandip Sen

Mathematical & Computer Sciences Department

University of Tulsa

Tulsa, OK 74104

{karpa,partha,sandip}@ens.utulsa.edu

Abstract

Cooperative behavior among single agents based on interaction history has been a stimulating field of research among social scientists and multi-agent researchers. The situation becomes complex in case of a group of agents seeking help from another group. The opinion of the members of the helping group about each of the asking group members can be combined to evaluate such a request for help. Exploitative agents would want to be part of groups that receive helps from other groups, but will try to avoid having to help other groups. Such agents, revealing false opinion about the reputation of others, can cause unwarranted penalization of the requesting group. This leads to global performance degradation in terms of reduced inter-group cooperation, and increased cost for the individual agents. We assume randomly formed, short-lived group compositions and study two strategies to use the collective opinion of the members of a helping group and identify situations where truthful agents are able to maintain lower costs even in presence of lying agents. We also study the relative merits of the two strategies in generating contracts in presence of lying agents.

Introduction

As agent-based systems continue to be incorporated into novel practical applications, the likelihood of agents interacting in open environments increase (Bradshaw 1997; CACM March 1999 issue 1999). Interaction strategies to promote cooperation in groups of self-interested agents and thwart malevolent behavior of exploitative agents have been an active area of study for social scientists and multi-agent researchers (Biswas, Sen, & Debnath 2000; Day & Taylor 1998; Zeggelink, de Vos, & Elsas 2000). Most of these studies, however, focus on inter-agent interactions rather than interactions between groups of agents (Sen 1996).

We draw our motivation from a prior work by Sen et al. (Sen, Biswas, & Debnath 2000) that gave insight on the relative performances of different agent behavioral strategies in terms of their effectiveness to generate

savings for the agents. Single agent interactions were considered and two variants of agent behaviors were evaluated, viz., *selfish* and *probabilistic reciprocity*. The selfish agents would ask for help, but never extend help to anyone. *Lying selfish* agents were used who “bad-mouthed” others by revealing false information about their reputation. Reciprocatative agents used a probability based mechanism that was based on past interactions with other agents to decide whether or not to honor a help request. Reciprocatative agents, who based their decision on the opinion of others about the agent asking for help, were also introduced and was found to be more robust in face of the exploitations of the selfish.

In this paper, we extend the inter-agent reciprocity to facilitate interactions among agent groups. We consider a generic situation where tasks are generated and assigned to a group of agents. The group that the task is assigned to checks whether its members are able to do the task or not. A task consists of multiple sub-tasks, each requiring a different expertise to be finished. A group is able to do the task only if it has at least one agent for all the expertise (sub-tasks) in the task. Otherwise, it asks for help from another group for that task. We assume that the group that has been assigned the task has knowledge of other groups with all the expertise needed to do the task. The decision to honor the help request depends on the manner in which the members of the selected group share their opinion about all agents of the asking group. Though the notion of a “contract” is usually associated with monetary transactions, we use the term “contract” to designate the situation when one group helps another to do the task of the latter. The task, in this case, is contracted to the helping group. When a group helps another group, the members of the helping group incur a cost by which they increase their balances with the members of the helped group. The members of the helped group, on the other hand, save a cost and hence, reduce their balances with those of the helping group. The “opinion” that an agent have about another agent is the balance it has with the latter. In our work, an individual agent maintains balance information about other agents based on past interaction history. The individual agent is a part of a group and the group acts

as a unit where help-giving decisions are based on the opinion of the constituent agent members. The motivation behind this conceptualization is to study whether the peculiarities of inter-agent interaction strategies reflect in inter-group interactions where the decision of honoring a help request depends on the shared opinion of the group members.¹

We assume that agents are not statically assigned to groups. Rather, in each time period, new groupings of agents are formed from agents with complementary expertise. So, the same agent can form groups with different agents at different times. An agent, however, accumulates and retains its interaction history over time with other agents in the environment. The characteristics of the individual entities of groups and the way the opinion of different individuals are shared dictate the number of successful contracts and therefore, influence the health of the economy.

We describe two different agent behaviors in this paper. The notion of a selfish agent that was used in (Sen, Biswas, & Debnath 2000) is modified in this paper. We introduce *lazy liars* who, while in a group of agents assessing a help request from the asking group, lie about the balances they have with the agents in the asking group, thereby reducing the chance of a job being contracted. They are “lazy” because they do not prefer to be in a group that helps another group by taking a contract from the latter. If a contract is accepted, then all agents in the helping group, including the lazy liars, have to perform subtasks that define a contract (details of a contract is given in Section *Simulation Framework*). The *reciprocatives* in this case, are those that reveal truthful opinion about the reputation of agents in the asking group.

We use two different strategies to use the collective opinion of the members of the asked group, viz., *average* and *worst*. We present preliminary results to show that such strategies are able to curb the derogatory tendency of the lazy liars. Our simulation based study reveals that the reciprocative agents, even in groups of a majority of lazy liars, are able to maintain lower costs on an average. This indicates the effectiveness of the strategies in pruning the harmful effects of the liars and is a measure of the stability of individual agent or, local performance. Also, we show that the *worst* strategy is able to generate more contracts than the *average* strategy with a non-trivial percentage of lazy liars in the population. This highlights the relative merits of the two strategies in upgrading global performance.

Simulation framework

We consider an information processing domain where a set of A agents are considered. Each agent has one of k different expertise. A set of T tasks are generated.

¹A parallel to this research is seen in the real world, where organizations pair up and jointly decide to exploit dynamic fleeting opportunities in modern economic markets (Henry *et al.* 2002).

Each task t has m ($m \leq k$) sub-tasks, each of which requires one distinct expertise to be completed. Having an expertise in a task type x implies tasks of type x can be completed incurring less cost compared to tasks of other types. We have used two metrics to compute the task costs, *time* to complete and *quality* of performance. An expert in task type x performs tasks of type x with less time and high quality and performs all other task types with higher time and lower quality. Task cost is defined as the ratio of time to quality, hence, an agent incurs low cost in task types for which it is an expert and high cost for all other task types. m is chosen randomly between 1 and k and then, m distinct expertise are assigned to the m subtasks. The group that the task is assigned to is selected globally by randomly selecting m agents among the population. This forms the asking group G . A group is able to complete an assigned task only if it has at least one agent with the required expertise for every sub-task. Since the m agents are selected randomly, it is not guaranteed that G will always be able to complete the task assigned to it.

If G is not able to complete its own task, it requires help from another group H . We assume that the members of G have perfect knowledge about the expertise of all agents, using which they select m agents where each agent has a distinct expertise and exactly those required to complete the task t . For each member of G the members of H provide their *opinion* which is averaged over all members of H . The opinion that agent i gives about agent j is the balance that the agent i has with j (*balance* _{i,j}).² The opinion that a reciprocative reveals about another agent is the true balance it has with the other agent. The opinion ($OP_{i,j}$) that a lazy liar agent i reveals about another agent j is defined as follows.

$$\begin{aligned} OP_{i,j} &= -\gamma * \text{balance}_{i,j}, \text{ if } \text{balance}_{i,j} < 0 \\ &= \text{balance}_{i,j}, \text{ otherwise} \end{aligned}$$

This set of equations show that the lazy liars express false opinion about those from whom they had earned more help in the past by increasing the balance γ times. Increasing the balance reduces the probability with which the help request of G is honored. The opinions of all agents in H are combined using two strategies that are described in Section *Selection Strategies*.

The decision to agree or refuse the help request of G is probabilistic and is based on the following equation. The probability that group H will help group G to do task t is given by

$$Pr(H, G, t) = \frac{1}{1 + \exp^{\frac{C_h^t + B_{hg} - \beta}{\tau}}},$$

where C_h^t is the cost to perform task t by group H , which is the sum of the costs of the individual agents

²*balance* _{i,j} is negative if i has received more help from j than it has given to j . When i helps j , *balance* _{i,j} is increased by the cost that i incurs, whereas *balance* _{j,i} is reduced by the cost that j would have incurred by doing the task on its own.

in H ; B_{hg} is the net balance that group H has with G , as computed by one of the strategies. β and τ are the only two constant parameters used in the function, where β is used to set the cost a group is ready to incur to help an unknown group with the hope of initiating a long-term cooperative relationship and τ is used to set the shape of the probability curve. This is a sigmoidal probability function where the probability of helping increases as the balance decreases and is more for less costly tasks.

If H helps G , group members of each of G and H update their balances with the members of the other group. The balances of each agent in H is increased with each of the members in G by $\frac{c_i^t}{|G|}$, where c_i^t is the cost incurred by agent i in H to do the i^{th} subtask of t (t_i) and, $|G|$ is the cardinality of G . So it is assumed that i incurred equal cost for each of the agents in G . Also, since there are exactly that many agents in H as there are subtasks in t , each agent i in H does exactly one subtask (i^{th} subtask) of t in which it is an expert. The policy in which the members of G update their balances with those of H is described in the following.

```

cost=0
For each subtask s in t do
begin
  x=number of agents in G with expertise
                               required for s
  if (x >= 1)
    cost=cost + expert's cost / x
  else
    cost=cost + non expert's cost / |G|
end

```

$|G|$ is the number of agents in the group G . In the above procedure we assume that for a subtask of t for which there are one or more experts in G , that subtask is equally shared by the corresponding experts. For a subtask in t for which there are no experts, it is divided among all members of G . The balances of the members of G with those of H are decreased by $\frac{cost}{|H|}$, assuming the cost saving was due to all the agents in the helping group.

Selection Strategies

We have designed two simple strategies to combine the individual opinions of the members of H on those of G . We describe these strategies in the following. The opinions of the members of H are the balances that they have with the agents in G . The overall opinion of the group H about an agent $j \in G$ is given by,

$$O_j = \frac{\sum_{i \in H} balance_{i,j}}{|H|}$$

Average Strategy: In this strategy the opinion of H about G is computed as the average of the overall opinions O_j about each member j of G , as given above. Hence,

$$B_{hg} = \frac{\sum_{j \in G} O_j}{|G|}$$

Worst Strategy: In this strategy H tries to punish any selfish agent hiding in G and makes the decision on the basis of the worst (maximum) balance that any member of H . Hence,

$$B_{hg} = \frac{\max_{j \in G} O_j}{|G|}$$

Results

In this section we report preliminary results to study the global performance of the system in terms of the total contracts made, as well as individual agent performance in terms of their average savings. We use the strategies described in Section *Selection Strategies* and measure the total number of contracts made. We also study the average balances of reciprocative and lazy liar agents. In the simulations, the population of lazy liars is increased from 10 to 90% in steps of 10%. We have used different values of agents and total tasks generated and the results are averaged over 10 runs.

Total number of contracts

In this set of experiments we measure the global performance of the system by comparing the total number of contracts made by using the two strategies for combining agent opinion. We used 100 agents and two values of tasks, 1000 and 2000. The percentage of lazy liars was varied from 10 to 90% in steps of 10%. The total number of contracts made for each value of lazy liars were averaged over 10 runs. We used $\gamma = 1$, which implies that the lazy liars are not imposing extra blame on others. The results are shown in Figure 1.

In Figure 1 (top) we plot the average number of contracts with increasing percentage of lazy liar agents with 1000 tasks and in Figure 1 (bot) the results with 2000 tasks are plotted. From both the plots we see that the worst strategy leads to a poorer performance than the average strategy for all values of lazy liar percentages. The worst strategy considers the largest balance that all agents in the helping group have with any agent in the asking group. When the lazy liar percentage is small, they have a reduced chance of damaging the reputation of the members of the asking group by falsely revealing balances higher than the true values. Hence the penalization of the asking groups, on an average, is less severe. As lazy liars increase in the population, the effect of their false opinions causes a havoc on the help asking groups. Higher the balance the members of the asking group have with those of the helping group, the severity of penalization by the worst strategy magnifies rapidly. With an increased number of lazy liars, the opinion that the asking group gets is mostly high and the worst strategy dictates that the probability of honoring the help request be reduced rapidly. This leads to a reduced number of contracts. The severity of penalization of the average strategy is less than the worst strategy. Hence, the number of contracts are, on an average, higher when the average strategy is used.

We notice that the difference between the plots in Figure 1 (top) and (bot) is in the percentage contracts

made. For the same percentage of lazy liar agents and the same strategy, the percentage contracts made with 2000 tasks is less than with 1000 tasks. We verified this trend by running similar experiments with 3000 tasks. We conjecture that, since the agent groups interact for longer times when there are larger number of tasks, the lazy liars affect other agents more severely by revealing false opinion over more extended time. This leads to more frequent rejection of help requests of a group when the number of tasks is higher.

Average balance

In this set of experiments we have studied the variation in the average balance of a reciprocative and a lazy liar agent with the two strategies for combining agent opinion. We have used 100 agents and two values of tasks, 1000 and 2000. A γ value of 1 is used. Figure 2 shows the results for the two strategies with 1000 tasks for different values of lazy liar percentage in the population.

In both Figures 2 (top) and 2 (bot), we see that with an increase in the lazy liar percentage, the average balance of both lazy liar (*LL*) and the reciprocative agents increases. Thus, they are able to save less cost with more lazy liars in the population, which corresponds to a degraded individual performance. For the same percentage of *LL* agents, the worst strategy leads to poorer performance of both agent types, i.e., generates higher balances. This is due to the characteristic of the worst strategy that penalizes heavily an asking group if the latter has even one member about whom the helping group generates a positive balance. With increase in the number of the lying agents, this penalization is boosted by the false opinion they reveal, leading to lesser number of contracts and hence, degraded individual performance. It is observed, however, that the balance of the reciprocative agents remain better (more negative) than the *LL* agents all throughout. This indicates that the reciprocatives are able to maintain greater savings than the *LL* agents even in the presence of extensive revelation of false reputation by the *LL* agents. The average strategy is more beneficial in this context because the reciprocatives earn better balances for the same percentage of *LL* agents than the worst strategy.

We conducted a similar set of experiments to study the average balances of *LL* and reciprocative agents with 2000 tasks. The results were similar and hence, we have excluded the figures.

We were interested to investigate whether there exists situations where the *LL* agents perform better than the reciprocatives. Since the *LL* agents reveal false reputations about others, we used a constant penalizing factor of $\gamma = 40$ in the next set of experiments. We used 5000 tasks and the results are shown in Figure 3.

From Figures 3 (top) and (bot) we observe that the balances of reciprocative and *LL* agents using the average strategies are better than the worst strategies. This is similar to the results we obtained in the previous experiments with $\gamma = 1$. We notice, however, that the re-

ciprocatives perform better than the *LL* agents by earning more balances when the percentage of *LL* agents is low (30%). When the *LL* agent population grows above that percentage, the tendency to undermine the reputation of other agents by the *LL* agents becomes too overpowering for the reciprocatives to maintain a better performance. This is because of using a more powerful bias of $\gamma = 40$ which is equivalent of expressing false opinion about other agents more severely. Thus, the truthful reciprocatives suffer and fail to accumulate as much help as it could when γ was equal to 1. The reciprocatives continue to be outperformed by the *LL* agents from thereon.

Results from the above set of experiments give us insight about different situations when the reciprocatives are viable and can earn better balances than the lazy liars. For the same strategy, an increase in the viciousness of the lazy liars outperform the truthful reciprocatives.

Conclusions and future work

We have investigated the impact of two decision procedures for task exchange between groups, where these decisions are based on the opinions of group members about the individuals in the other group. In particular, we are interested in the number of task exchanges and the relative performance of reciprocative and lazy liars as the percentage of the latter is varied in the population.

We have seen that reciprocative agents are performing better than the lazy liars even at a high proportion of the lazy liars in the environment. The reciprocative agents performed better in spite of the negative opinion reported by the lazy liars. The only exception is when the percentage of lazy liar agents is high and their negative opinion of other agents is largely exaggerated. In the current work, groups are formed randomly. Instead we can allow agents to form groups based on their knowledge of other agents. In this framework all the agents have knowledge about all other agents, their behavior, performance and expertise. Agents can then choose to form teams with truthful agents. This will in turn further restrict the exploitative tendencies of lazy liars.

References

- 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.
- 1999. Communications of the ACM, march 1999, volume 42, number 3. Special Issue on Multiagent Systems on the Net and Agents in E-commerce.
- Day, T., and Taylor, P. D. 1998. The evolution of temporal patterns of selfishness, altruism, and group cohesion. *The American Naturalist* 152(1):102-113.

Henry, R. A.; Kmet, J.; Desrosiers, E.; and Landa, A. 2002. Examining the impact of interpersonal cohesiveness on group accuracy interventions: The importance of matching versus buffering. *Journal of Artificial Societies and Social Simulation* 87(1):25-43.

Sen, S.; Biswas, A.; and Debnath, S. 2000. Believing others: Pros and cons. In *Proceedings of the Fourth International Conference on Multiagent Systems*, 279-285. Los Alamitos, CA: IEEE Computer Society.

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.

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-help ing-partners'. *Journal of Artificial Societies and Social Simulation* 3(3):212-15.


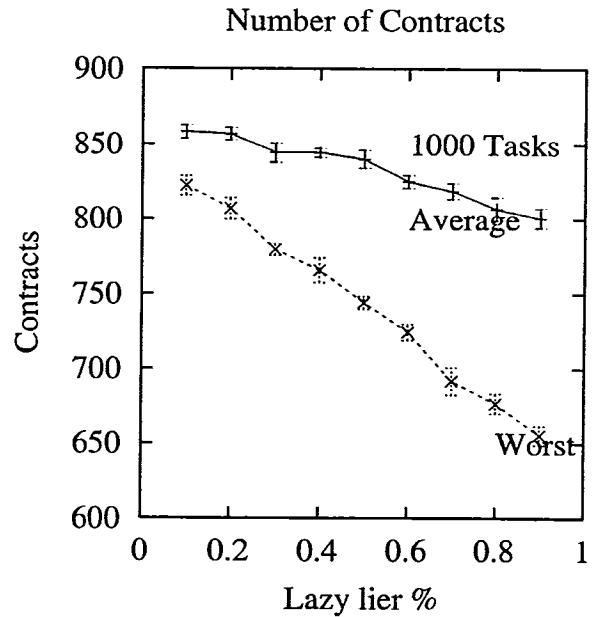


Figure 1: Average number of contracts with 1000 tasks (top) and 2000 tasks (bot).


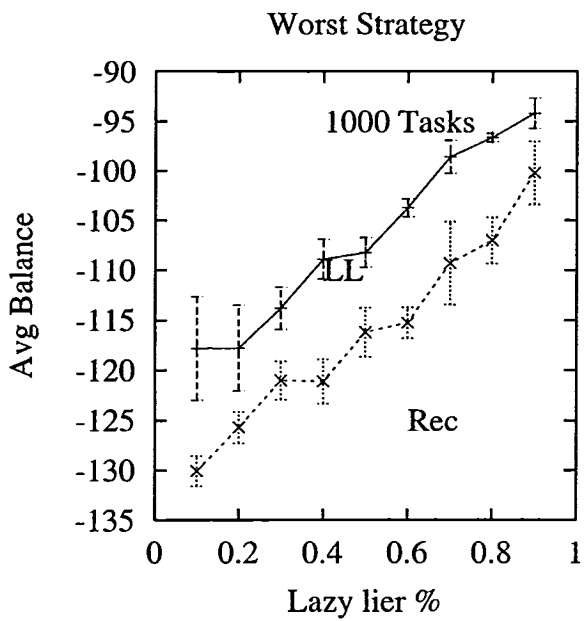

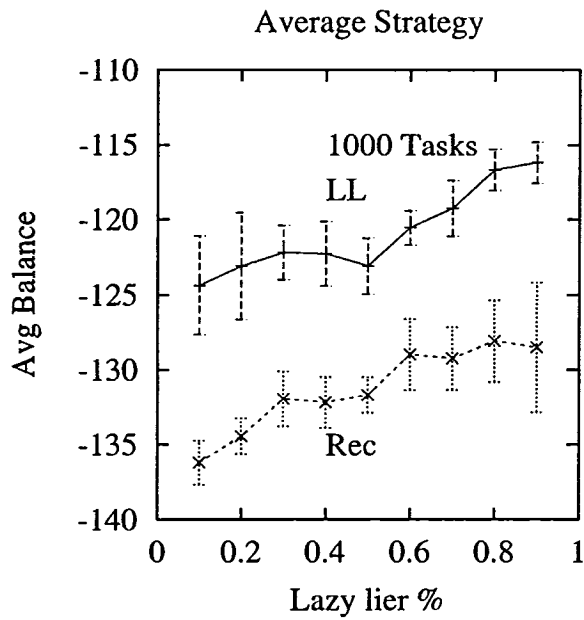


Figure 2: Average balance with average (top) and worst (bot) strategies ($\gamma = 1$).

Figure 3: Average balance with average (top) and worst (bot) strategies ($\gamma = 40$).