Help Under Risky Conditions:

Robustness of the Social Attitude and System Performance[*]

Amedeo Cesta
IP-CNR
National Research Council
Viale Marx 15
I-00137 Roma, Italy
amedeo@pscs2.irmkant.rm.cnr.it

Maria Miceli
IP-CNR
National Research Council
Viale Marx 15
I-00137 Roma, Italy
maria@pscs2.irmkant.rm.cnr.it

Paola Rizzo
Center of Cognitive Science
University of Torino
Via Lagrange 3
I-10123 Torino, Italy
paola@di.unito.it

Abstract

This paper deals with the problem of co-existence in the same environment of agents with different attitudes towards each other. In previous works of the same authors, agents endowed with different help-seeking and help-giving attitudes have been defined and simulated in a multi-agent system, and the effect of their behavior on the global performance of the system has been measured. It has been shown that the "social" agents, which both give and ask for help, are the most successful, and may be able to tolerate very well the presence of exploiters. In this paper the robustness of the social strategy is tested by putting help-giving under risky conditions. The following three variables have been manipulated: number of resources available; threshold triggering the help-giving behavior; and ratio between different kinds of exploiters and social agents. The results show that, with a few exceptions, the social strategy is quite robust in several risky conditions. Some considerations about the usefulness of such a strategy for multi-agent systems are also outlined.

Keywords: Cooperation and Coordination, Evaluation of Multi-Agent Systems

Introduction

Multi-Agent Systems (MASs) have recently gained major attention, also due to technological improvements. The possibility of having very different kinds of MASs, from the more flexible to the more rigid, has fostered investigation in multiple directions. This paper presents a series of experiments conducted in a simulated MAS environment in which agents are characterized by rigid interaction attitudes toward each other.

The problem of coordination and co-evolution in the same environment has become central both in Distributed AI (Durfee 1988; Rosenschein & Zlotkin 1994) and robotics (Mataric 1995). Increasing evidence has been provided of the possibility to have an environment where different kinds of artificial agent may interact either for a common goal or for performing independent tasks. The current approaches to the study of such problems may be roughly distinguished along the *game-theoretic perspective*, which mainly focuses on the individual perspectives of the single agents, and on the personal utility each agent pursues in choosing its "social" behavior, e.g. (Rosenschein & Zlotkin 1994); and the *DAI perspective* that instead focuses on the performance of the system as a whole even if different agents may pursue different goals.

This paper adopts the latter point of view. However within the DAI perspective different aims have been pursued with regard to the problem of co-existence. The *coordination based* approach, e.g. (Durfee, Lesser, & Corkill 1987; Conry *et al.* 1991), has concerned the synthesis of reliable protocols for maximizing the "useful" interactions among agents designed as homogeneous. Even if systems are considered as open, all the possible behaviors are a-priori considered and ways to cope with exceptions envisaged. Conversely the *learning-based* approach starts from the reasonable assumption that it is impossible to make up all the possible agents' behaviors as far as the system is open to new agents. As a consequence, agents should be endowed with learning abilities to achieve self-adaptation to different contexts. Several recent works have begun to explore this approach, e.g. (Tan 1993; Sen, Sekaran, & Hale 1994).

We believe that space is left for a different investigation in the same area of problems. In particular we would like to understand how different attitudes (more specifically "interaction" attitudes) may influence the overall performance of the system when no reciprocal adaptation is allowed and each kind of agent behaves according to its own rigid strategy, independent of the strategy shown by others.

A Robotic Example

An example useful to specify the problem addressed is shown in the following. Suppose a number of companies form a joint venture to share a space mission to Mars in order to gather environmental data (e.g., data on the Mars athmosphere). The agreement implies sharing the expenses for a round trip to Mars by hiring a space vehicle the Space Ship. Each com-

---

[*] The authors' names are listed in alphabetical order.

From: Proceedings of the Second International Conference on Multiagent Systems. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

pany owns a number of robots used for similar tasks on Earth. These robots are of different kinds, but share a number of features. All of them carry a special recording tool to store data from the environment, they are unable to directly communicate data to Earth but they should carry the tool back with themselves in the Space Ship to return to Earth. All robots gather uranium (that might be generally viewed as *food* for them) from the environment, and trasform it into energy to survive. If they run out of energy they are unable to return to the Space Ship for the trip back (they *die* on Mars). Robots differ in their ability to move in the environment and in their skills for communicating with other robots. As far as such differences are concerned, we can group robots into four categories:

First generation robots: they are autonomous robots able to move in an uncertain environment, but they have no communication device. So, they are able to gather information, and are able to look for uranium in order to survive, but cannot communicate with any other robot. Let us call those robots *solitary* for their attitude to ignore others.

The following three kinds of robots are "second generation" ones. They are able to communicate to similar robots their state of danger represented by a low level of energy but differ in other abilities.

Second generation static robots have a limitation in their moving abilities. They are able to scan a very narrow surface in the environment, so they may be considered as static. This fact does not undermine the data gathering task, but their ability to supply themselves with energy by finding uranium. They just "eat" uranium in the space they can scan otherwise they send requests for help to other agents for having uranium carried to them. Because of their total dependence on others and absence of reciprocative helping behavior, we call them *parasites*.

Second generation mobile robots freely move in the environment, and therefore are able to gather and trasform uranium. However, when in particular shortage of energy they stop their search and ask for help, like the previous ones. In addition, they never help other needy robots because they are unable to carry uranium. We call those robots *selfish* because they ask for help if in danger, but never give help to any other robot.

Second generation social robots in some sense are the more skilled ones in terms of flexibility of behavior. They are able to autonomously move, can trasform food after searching it, but they are also able to bring food to needy robots, as well as to ask for help if in need. For their helping attitude we call those agents *social*.

The important thing is that robots gather information in the environment and safely survive in order to be able, after a certain time, to fly back to Earth. As a consequence it is important not to fall below a certain level of energy, that would cause the robots to die.

Companies have a certain number of the different kinds of robot and are not interested in studying which are the best robots to send to Mars (suppose they don't have enough of any type, and in particular of the social ones, which can be seen are more likely to survive if left by themselves) but in understanding what happens if they put together the kinds of robot they have, and try to fill the Space Ship as much as possible. In a word, companies are willing to use the existing robots, and study how to maximize their revenues with them. Our experimental setting and the experiments performed are aimed to shed light on different aspect of this problem.

A previous work (Cesta, Miceli, & Rizzo 1996) has provided first data for this kind of experiments showing that in a basic situation the social attitude appears quite robust, in that it allows to tolerate the exploiters without risking dangerous consequences for the social system; in addition a certain degree of self-sufficient attitude (like that shown by the selfish) proves more advantageous, if compared with total dependence on, or exploitation of, others (typical of the parasites). In the present paper, after a description of the experimental environment and the agent architecture, we first summarize and discuss those results, and then present new experiments aimed at testing the robustness of the social strategy by putting help-giving under a variety of risky conditions. A discussion of the different results closes the paper.

A Controlled Artificial Setting

Our simulated environment metaphorically represents a multi-agent system in which some agents, behaving in a common world (say, Mars) with limited shared resources (say, uranium ores), can pursue different and autonomous goals, and find themselves in dependence relations with one another. In the following we shall ignore the possible specific tasks assigned to such agents (like the information gathering task assigned to robots on Mars which is of no concern here) and will limit our description to the general "task" of surviving and to the different interaction attitudes typical of these agents, which, as we shall see, are likely to affect survival itself.

The environment is a two-dimensional grid where some uranium (from now on, food) is randomly located; this world is populated by simple agents that need to look for food and eat it in order to survive, and that can interact with one another.

The simulated environment is interfaced with the MICE testbed (Montgomery & Durfee 1990), a public domain software for simulating two-dimensional worlds in which agents can interact.

Agent Architecture

The agent architecture is quite simply composed of a visual sensor, a goal generator, a planning module, and a set of effectors that are here briefly described.

Sensor. The sensor lets the agent perceive both the pieces of food and the other agents within a limited area. Such area is measured in terms of grid units, and consists in a square of side $2n + 1$ around the agent, where n is a given sensorial range.

Goal Generator. The goal generator chooses a goal to pursue on the basis of the sensorial information and of the agent's internal state: the latter is related to the energy level, which ranges with integer values from 0 to 100. Agents die when their energy goes below 0. The energy has two intermediate thresholds that play a role in the switch of the agent's internal state, as shown in Figure 1.

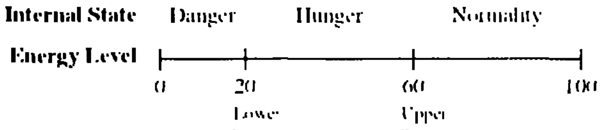

Figure 1: Relationship between Energy Level and Internal States

Planner. The planner selects a plan, composed of a sequence of actions, suitable to pursue the agent's goal, and controls its execution. At present, the planning module limits itself to choose a plan from 3 canned ones indexed by goals (see Table 1).

Effectors. The effectors execute elementary actions: moving (from one grid location to another at a time), taking, giving, and eating (one piece of food at a time), choosing a piece of food among a set of perceived ones, and signalling a needy state to the other agents (by changing one's own appearance). Each action affects the agent's internal state by lowering the energetic level in a specified amount, except for the action of eating which increases it by an amount equal to the food energetic value. Actions have either a procedural implementation or a direct translation in terms of MICE commands.

Agents' Behavior and Interaction Attitudes

As already mentioned, we are interested in exploring the effect of different interaction attitudes upon the agents' survival; such attitudes are generated by embedding different associations between the symbolic internal states (the set *Danger, Hunger, Normal*), and the agent's goals. The relationships among types of agent, internal states, and goals are summarized in Table 1.

Four types of agent are defined:

"Solitary", that just ignore one another, so that there is no interaction among them; their goal is always to individually find food, regardless of their internal state.

"Parasites" that, regardless of the internal state, always stay still, having the only goal of looking for help and eating the food received by other agents; this means that if parasites are not given help they will definitely die.

"Selfish" that, depending on their internal state, either ask for help (when in danger) or autonomously look for food (when hungry or in the normal state). The only social behavior showed by both the selfish and the parasites is help-seeking.

"Social", that may have different goals according to their different internal states; more precisely, in case of danger their goal generator activates the goal of looking for help; when hungry, their goal is to find food; and finally, in case of normal state, i.e., when they reach the upper threshold of 60, if there are any visible help-seekers, the goal of giving help is activated; otherwise, they go on looking for food.

Type of Agent	Internal State	Goal
Solitary	*any*	Find Food
Parasite	*any*	Look for Help
Selfish	Danger	Look for Help
	Hunger, Normal	Find Food
Social	Danger	Look for Help
	Hunger	Find Food
	Normal	Give Help Find Food

Table 1: Relationships among Types of Agent, Internal States, and Goals

Experimental Conditions

The world is a 15 x 15 grid which contains 60 food units and 30 agents, all randomly located. The food units keep constant until the end of the simulation by randomly reappearing on the grid each time one or more agents eat some food[1]. The initial energy of each agent is set to 50; when they move, agents waste 2 units of energy per time unit, while for staying still they waste 1 unit of energy per time unit. The sensorial range is 3 (i.e. agents can perceive things within a 7 x 7 square around themselves). They choose to

---

[1] The reappearance of food units is justified by the fact that our agents (unlike robots on Mars) cannot go beyond the borders of the grid, which is quite small. So, we simulate the agents' search for uranium ores in a large area by letting the food units reappear within the border of the grid.

help the nearest recipient, and to eat the nearest piece of food. Each simulation lasts 500 iterations (an iteration ends when the cycle "sensing - goal generation - planning - executing" has been performed for every agent), and data are collected every 50 iterations (for a total amount of 10 temporal samples). Therefore, The results shown in the following figures concern the percentages of alive agents after 500 time steps, plotted against the food energetic values; each point in the figure represents the mean value across 10 simulations. The x-axis has been put in logarithmic scale to produce clearer plots. All results have been submitted to the analysis of variance, with a level of probability of 5%; all the differences mentioned in the description of the results are significant.

Questions and Results of Previous Experiments

Once identified those different characters, our previous experiments were aimed at answering a number of questions, among which the following are relevant to our present purposes:

> *What is the effect of the helping behavior on the overall performance of a multi-agent system, in terms of the agents' survival?* Would the social agents, that help each other when in need, outperform the solitary?
>
> *How does the coexistence of exploiting agents (that ask for help but never give help) and social agents affect the performance of both the entire system and of the social agents themselves?*

Solitary vs. Social: The Positive Effect of Mutual Help

The answer given to the first question is that *help giving and seeking increases the probability of survival of the entire system.*

Each simulation has been performed with a homogeneous group of agents (either 100% of solitary or 100% of social agents). The comparison between solitary and social agents shows that the social condition increases (in a statistically significant way)[2] the percentages of survived agents. Help-giving increases the probability of survival of the needy agents, and since each agent can become needy in given circumstances, help-giving, in accordance with the theory of reciprocal altruism (Trivers 1971), "helps" the helper itself, and ultimately the whole social system.

Social with Exploiters: Robustness of the Social Strategy

The answer given to the second question is that *the social agents are able to tolerate exploiting help-seekers without great danger both for the entire system and for*

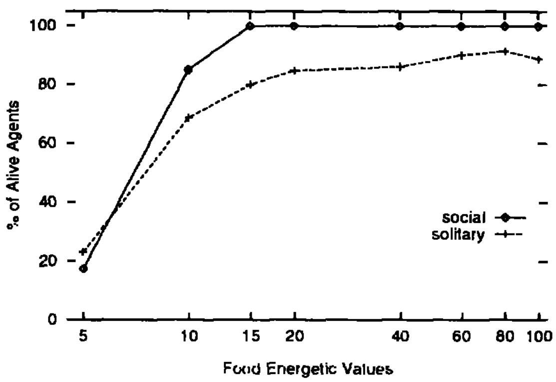

Figure 2: Comparison between Solitary and Social

themselves. As far as different kinds of exploiters are concerned (parasites vs. selfish), some interesting differences emerge. The simulations have been done by putting 15 social agents with either 15 parasites or 15 selfish in the same environment (with the same parameters as in the previous experiments). Results of both series of simulations are sintesized in Figure 3.

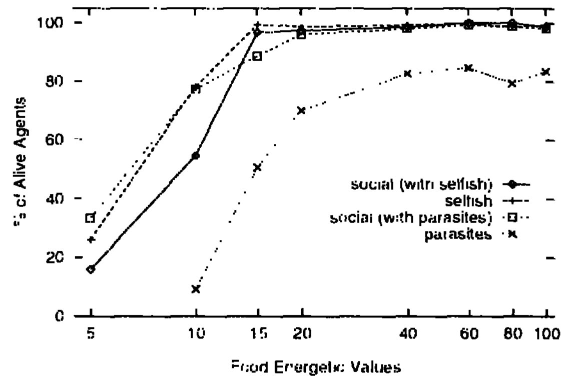

Figure 3: Social with Parasites vs. Social with Selfish

As one can see, the system as a whole performs pretty well, both in the case of social with parasites and in the case of social with selfish.[3] Comparing the destiny of the group of social agents in the two conditions with that of the social agents by themselves (see Figure 2), we can see that: (a) their interaction with the parasites does not affect their survival; (b) when the social agents come to interact with the selfish, both groups perform at their best (100% of survived agents of each type) for each energetic value of the food but

---

[2] For a more detailed account of these and the following results, see (Cesta, Miceli, & Rizzo 1996).

[3] In the graphs, "parasites" and "selfish", unlike otherwise specified, are always considered in interaction with social agents. When the same graph shows social agents in interaction with either kind of exploiters, "social (with parasites)" and "social (with selfish)" indicate social agents that interact with either parasites or selfish, respectively.

the two lowest ones, when the selfish significantly outperform the social agents. In fact, when the food energetic value is very low, the energy the social agents waste while helping is not easily restored by the food they receive from others or get autonomously. Thus, under such conditions, the interaction with the selfish is more dangerous for the social agents than the interaction with the parasites. The selfish, in fact, by looking for food when not in danger, actively compete with the social agents over the available resources, while the parasites can do nothing but staying still and waiting for help all the time.

Preliminary Conclusions and Open Questions

The previous results point to the well-known advantage of mutual help. Social agents, through their mutual support, are able to compensate for the exploitation they suffer from the non-reciprocating agents. The helping behavior is a powerful strategy for increasing the probability of survival of the helper itself. And not only the social agents do not risk dangerous consequences in terms of their own survival, but the entire social system benefits from their strategy.

However, a basic objection could be raised against the robustness of the social strategy. In the experiments presented so far, the parameters chosen make help-giving relatively "safe". For instance, the number of resources is quite large (60 food units for 30 agents), and help is given when the agent has a considerable amount of energy — at least 60/100. It is plausible to assume that help given under more risky conditions would be more disturbing for the social agents, and one might wonder how much the robustness of the social strategy depends just on certain experimental conditions (i.e., it is an experimental "artifact"). To dispel this doubt, one should verify the robustness of the social strategy under more risky conditions for help-giving.

The Social Strategy Under Risk: Experimental Conditions and Results

We put the social strategy's robustness to the test by modifying three experimental variables.

Number of resources. If resources are cut down, help-giving should be more risky. In fact, the probabilities to eat are reduced; hence, the helpers' probabilities to reach high levels of internal energy should decrease; in addition, the requests for help should increase[4]. As a consequence of both facts, help-giving should cause the social agents to often risk reaching low energy levels, hence not being able to tolerate exploitation.

Threshold for help-giving. If that threshold is lower, help-giving should be more risky. As already observed, in our previous experiments help-giving could not occur until the social agent had reached the threshold of the "normal" state, i.e., when its energy level was 60/100 or more. If the threshold is lowered, the social agents would help more frequently and at lower energy levels, and this would make it more difficult for them to restore and maintain high levels of energy after each helping act, and therefore to tolerate exploitation.

Ratio between social agents and exploiters. If the exploiters are in the majority, help-giving should be more risky, in that the minority of social agents should have some difficulty in facing the exploiters' requests for help, and risk succumbing to exploitation.

Reduced Resources

In these experiments, 15 social agents have been put with either 15 parasites or 15 selfish, with the same parameters as in the previous experiments, except for the number of resources: 30 food units (rather than 60) for 30 agents. However, we focus here on the interaction between social agents and selfish (Figure 4), because, as expected, this interaction put the social agents to a serious test.

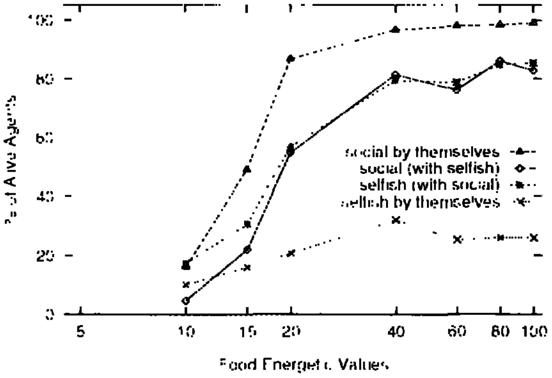

Figure 4: Social and Selfish with reduced resources

As shown in Figure 4, by halving the resources we find that even 30 social agents by themselves, if compared with the social agents by themselves of our first experiment (Figure 2), perform of course a little worse. In fact they show to suffer from the scarcity of resources when the food energetic value is low (from 5 to 20); however, they are able to already stabilize around 100% of survival when the food energetic value is 40. When interacting with the selfish, social agents are disturbed by their exploitation more than in the original experiment (see Figure 3). However, they are able to stabilize around 80% of survival when the food energetic value is 40. In addition, in spite of their exploiting

---

[4] This is limited to the selfish, because the parasites, as we know, ask for help all the time, independent of their energy level.

From: Proceedings of the Second International Conference on Multiagent Systems. Copyright © 1996, AAAI (www.aaai.org). All rights reserved.

help-seeking, the selfish are not able to outperform the social agents, except when the food energetic value is very low, showing the same trend as in the original experiment (see Figure 3). Finally, comparing the curve of the selfish by themselves with the curve of those receiving help, we can see that the latter actually benefit from the help they receive, showing a far better performance. Thus, we can conclude that, though more at risk in the present condition, the social strategy still proves quite effective both in dealing with exploiters and in favoring a better performance for the entire system.

Lowered Threshold for Help-Giving

In these experiments, 15 social agents have been put with either 15 parasites or 15 selfish, with the same parameters as in the original experiments, except for the threshold of the social agent's energy for giving help, that has been lowered from 60 to 40 in a first series of simulations, and to 20 in a second series of simulations. Interestingly enough, here, as shown in Figure 5 (threshold at 40), we find that the parasites are quite disturbing for the social agents, much more than the selfish.

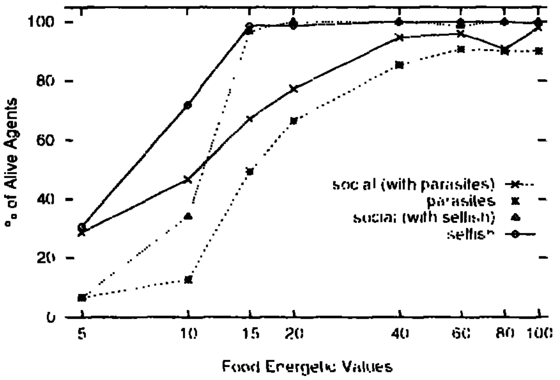

Figure 5: Social with Selfish vs. Social with Parasites with lowered threshold (40) for help-giving

In fact, when helping the parasites, the social agents are able to reach a respectable performance only when the food energetic value is around 40. On the contrary, when interacting with the selfish, the social agents not only "tolerate" the selfish, but they are able to reach the top of 100% of survival starting from a food energetic value of 15. The same holds for the selfish that, as usual, slightly outperform the social agents when the food energetic value is very low, and then stabilize around 100% of survival.

Even more striking results are obtained when the threshold for help-giving is lowered to 20. In this case, as far as the interaction between social agents and parasites is concerned, the former clearly show they are *not* able to tolerate exploitation. In fact, both parasites and social agents die before 500 iterations. For that reason, we are not able to summarize the results of the simulations in a graph. On the contrary, as shown in Figure 6, the interaction with the selfish shows the same trend as before.

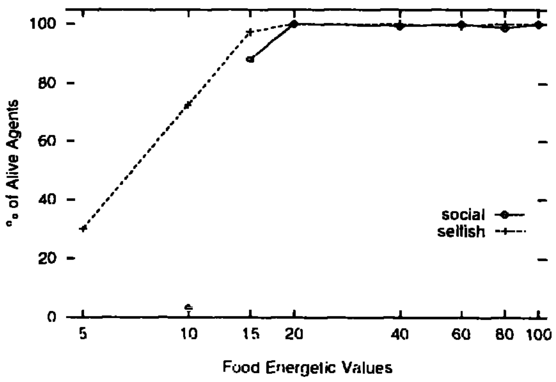

Figure 6: Social and Selfish with lowered threshold (20) for help-giving

The conclusion we can draw from these results is that such a risky condition implies very different consequences depending on the kind of exploiters involved. While the interaction with the selfish does not impair the robustness of the social strategy, as well as their usefulness for the selfish, the interaction with the parasites is very disturbing. Such a difference, however, makes sense if we consider the typical behavior of the parasites. As we know, they ask for help all the time, irrespective of their internal state (while the selfish ask for help just when in danger). When the threshold for giving help is quite high (as in the original experiments), the social agents are able to "ignore" their requests for most of the time, and to increase their own energy without suffering from their presence. But now, with a lowered threshold, the social agents should help the parasites almost continuously, even when they are endowed with scarce resources of their own, with ruinous consequences for themselves, and no special advantage for the entire system. Conversely, the lowered threshold does not constitute a serious risk with the selfish, thanks to the latter's self-sufficient attitude.

Social Agents in the Minority

In these experiments, 5 social agents have been put with either 25 parasites or 25 selfish, with the same parameters as in the original experiments.

Here, as shown in Figure 7, while 5 social agents still tolerate 25 parasites quite well, in that their performance is not impaired by that majority of exploiters, they are less able to make the parasites survive under such conditions. The help the parasites receive permits them just not to die immediately (as it would happen if they were left by themselves), but their percentages of survival are very low for each energetic value of the food. The conclusion we can draw is that under such condition the social strategy is still robust, but its usefulness for the entire system is impaired by their being in such a little number.

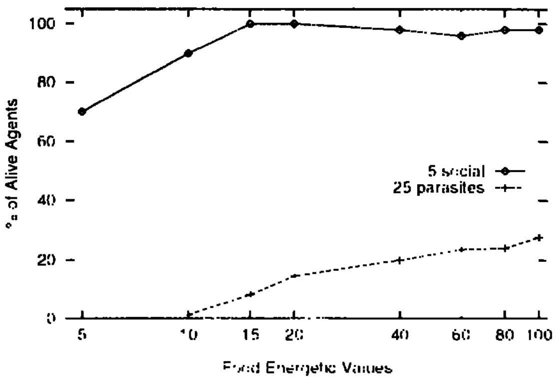

Figure 7: 5 Social with 25 Parasites

As for the interaction between 5 social agents and 25 selfish, Figure 8 shows that the social agents in the minority, if compared with 15 social interacting with 15 selfish (Figure 3), are not greatly disturbed by the selfish (the disturbance is significantly greater than in the original experiment just in one case, when the food energetic value is 15). In fact, starting from the energetic value of the food set at 20, they perform pretty well. In addition, it should be stressed that such a minority of social agents allows also the selfish to perform pretty well, much better than they would do if left by themselves (see the curve of the selfish by themselves in Figure 8). Therefore, we can conclude that under such condition not only the social strategy is still robust, but it is also useful for the whole system.

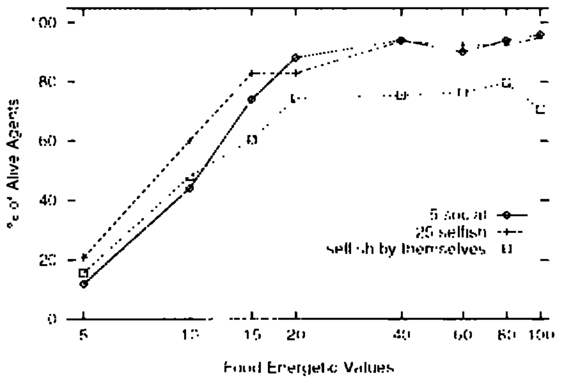

Figure 8: 5 Social with 25 Selfish

Discussion

In this work, our concern has been to see how coexistence in the same environment of agents with different interaction attitudes affects the performance both of the system, and of the various subgroups of agents. In particular, we wanted to test the robustness of the social strategy.

Unlike other works, e.g. (Axelrod 1984; Glance & Huberman 1993; Lindgren & Nordahl 1994; Novak, May, & Sigmund 1995), that look at the interaction among various "strategies" in conditions where agents are sensible to each other's behavior, and adapt their own strategy to those shown by their interactants, our agents' behavior is fixed: each kind of agent behaves according to its built in attitudes, independent of the rewards or punishments it may receive from others. In addition, we are not concerned with how such kinds of agent (and their interaction attitudes) evolve across generations. In a word, we want to see "what happens" in very rigid and simple social systems, before exploring more complex and flexible ones.

From our results we can draw the following general conclusions. First of all, one can notice that the robustness of the social attitude is globally verified, particularly when the social agents interact with the selfish, which in the original experiments appeared as the most dangerous.

In addition, some interesting and counterintuitive results have emerged as well, such as the greater dangerousness of the parasites in particular cases, namely when the threshold for giving help is lowered. This indicates the crucial role played by the specific kinds of risky conditions. In fact, what under certain risky conditions is a less dangerous interaction may turn into a more dangerous one under other risky conditions. Moreover, a generalization one can draw from such results is that it is less the presence of exploiters which can put a social network in danger than the absence of "filters" against exploitation. In fact, the threshold for giving help can be considered as a sort of filter. Going back to our robot example, companies preoccupied with the risk implied by making their social robots interact with exploiters might reach the conclusion that what really matters is to endow their social robots with such a simple filter, rather than avoiding their coexistence with exploiting robots.

Another interesting fact is the relation between the robustness of the social strategy and its usefulness for the entire system. Sometimes the two dimensions are directly related: if the strategy is robust, it can save the entire system; if it is not, the destiny of the social agents affects that of the exploiters. However it is not always like this. Consider for instance the case where 5 social agents interact with 25 parasites: here, the social agents are able to survive quite well, but unable to "save" the parasites. Once again, the threshold for giving help (which in this case is high as usual) allows the social (which are anyway in strong minority) to be robust, saving them from a disruptive exploitation, but on the other hand it does not allow parasites to survive at an acceptable level.

A final consideration concerns the comparison between the two types of exploiters. In addition to being dangerous in different ways depending on the kind of risky conditions, it must be stressed that in any case a selfish agent has been proved to be more "robust" than a parasite: a moderate autonomy allows to get a great benefit from exploitation, while total dependency on others puts the depending agent at the mercy of circumstances.

Our work is based on an abstract view of MASs. We have placed our agents in an environment where their only task is survival. The latter can be considered a good metaphor for a task which needs no social coordination and must be performed autonomously by each agent, but at the same time it can be affected by the other agents' cooperative attitude and produces global effects on the system. Our findings could in principle be replicated in systems with such characteristics. However, this is yet to be tested: we would like to see whether the behavioral attitudes of our agents would produce similar effects on the performance of some specific task in a MAS. In particular, we would like to see whether the social attitude is still robust in different environmental conditions.

Acknowlegments

This research is partially supported by Esprit III BRWG project No.8319 "A Common Formal Model of Cooperating Intelligent Agents (ModelAge)", and by IP-CNR — Division of Artificial Intelligence, Cognitive Modeling and Interaction.

References

Axelrod, R. 1984. *The Evolution of Cooperation.* Penguin: London.

Cesta, A., Miceli, M., and Rizzo, P. 1996. Effects of different interaction attitudes on a multi-agent system performance. In de Velde, W. V., and Perram, J. W., eds., *Agents Breaking Away (Proceedings of Maamaw '96, Eindhoven, The Netherland, January 22-25, 1996).* LNAI 1038, Springer: Berlin. 128–138.

Conry, S. E., Kuwabara, K., Lesser, V. R., and Meyer, R. A. 1991. Multistage negotiation for distributed constraint satisfaction. *IEEE Transactions on Systems, Man, and Cybernetics* 21:1462–1477.

Durfee, E. H., Lesser, V. R., and Corkill, D. D. 1987. Coherent cooperation among communicating problem solvers. *IEEE Transactions on Computers* 36:1275–1291.

Durfee, E. H. 1988. *Coordination of Distributed Problem Solvers.* Kluwer: Dordrecht.

Glance, N. S., and Huberman, B. A. 1993. The outbreak of co-operation. *Journal of Mathematical Sociology* 17:281–302.

Lindgren, K., and Nordahl, M. G. 1994. Cooperation and community structure in artificial ecosystems. *Artificial Life* 1:15–37.

Mataric, M. J. 1995. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems* 16(2-4):321–331.

Montgomery, T. A., and Durfee, E. H. 1990. Using mice to study intelligent dynamic coordination. In *Proceedings of the IEEE Conference on Tools for Artificial Intelligence.* IEEE Computer Society Press.

Novak, M. A., May, R. M., and Sigmund, K. 1995. The arithmetic of mutual help. *Scientific American* 272:50–55.

Rosenschein, J. S., and Zlotkin, G. 1994. *Rules of Encounters. Designing Conventions for Automated Negotiation among Computers.* MIT Press: Cambridge, MA.

Sen, S., Sekaran, M., and Hale, J. 1994. Learning to coordinate without ssharing information. In *Proceedings of AAAI-94.* 426–431.

Tan, M. 1993. Multi-agent reinforcement learning: Independent vs. cooperative agents. In *Proceedings of the Tenth International Conference on Machine Learning.* 330–337.

Trivers, R. L. 1971. The evolution of reciprocal altruism. *Quarterly Review of Biology* 46:35–57.