The Role of Social Control Systems in Conflict Resolution

Matthias Scheutz and Paul Schermerhorn
Department of Computer Science and Engineering
University of Notre Dame
Notre Dame, IN 46556, USA
{mscheutz,pscherm1}@cse.nd.edu

April 23, 2003

Abstract

This paper describes research with agents employing the ability to evaluate external cues given by other agents to determine their opponents' behavioral dispositions in conflicts. We introduce two agent types, "asocial" agents whose behavior is determined solely by their own dispositions, and "social" agents who also consider others' dispositions when deciding how to act in a conflict. The experiments described here suggest that such "social" mechanisms can be beneficial for agents, indicating a need for further research into the costs and benefits of such mechanisms.

1 Introduction

There are many situations in which animals encounter one another and conflicts ensue. They may be attracted by the same resource (say, a food source) or the same potential mate. One may simply wander into the other's territory. In any case, when such conflicts arise, there must be some way of resolving them. In some cases, a fight ensues immediately to determine which animal obtains the resource. In others, there is some mediating process that may determine the winner. In green anole males, aggressive display is used to resolve territorial conflicts [2]. Behaviors such as head bobbing and dewlap extension provide signals to opponents in aggressive encounters. If neither combatant backs down, a fight will ensue. Male cuttlefish adopt the "Intense Zebra Display" during the initial phases of a conflict [1]. Facial darkness can vary in these displays, with lighter-faced combatants more likely to retreat. If both combatants maintain dark faces, however, fighting ensues. Male Mediterranean field crickets engage in two behaviors, antenna fencing and mandible spreading, that are thought to indicate fighting readiness and fighting ability, respectively [3]. When there is a large difference in the pace of antennal fencing between two combatants, the slower of the two is likely to back down. Similarly, when there is a difference in displayed mandible strength, the weaker is more likely to step down. These examples represent strategies that allow the users to discern information about their opponents, including motivational states, to aid them in making decisions about when to fight.

The focus of this paper will be on the display of characteristics that allow agents to predict other agents' behavioral dispositions (i.e., the agents' behavioral dispositions are correlated with the external display behaviors). An agent's behavioral disposition can be in the range from aggressive to fearful. The more aggressive an agent's disposition for conflicts is, the more likely it is to decide to fight. We use the term "asocial agents" to refer to agents that consider only their own dispositions in making decisions in conflicts. "Social agents," on the other hand, are agents whose own dispositions are mediated by their opponents' dispositions. An "asocial" agent with an aggressive disposition is likely to fight, even when faced by a more aggressive opponent, because it ignores that agents' dispositional cues. An aggressive "social" agent faced with a more aggressive opponent, however, will be less likely to decide to fight than a similarly aggressive "asocial" agent.

Using information contained in opponents' displays of aggression or fearfulness, an agent is able to compare its own level of aggression or fear with that of his opponent and select an appropriate action based on that. Generally speaking, if my opponent is displaying a great deal of aggression, it is unlikely that it will back down, and it may be in my best interest to move along and find another area or resource. Similarly, if I sense that my opponent is fearful of an encounter, it will often prove worthwhile to confront it, since it is likely to retreat. There are many ways in which affect can be displayed, from twitching or perspiring when nervous to baring teeth when aggressive. For agents that can interpret these cues, they can provide useful input for behavior selection. This mechanism is social in the sense that it relies on information provided (willingly or otherwise) by other agents.

To explore the effects of social conflict resolution mechanisms, we implemented agents with the ability to sense the fear and aggression of other agents. Social and asocial versions of these agents were then allowed to interact in the SimWorld artificial life environment, and their performance was evaluated in terms of the average number of survivors over a number of simulation runs. We found that the simple social control architectures we implemented performed well against asocial agents in mixed environments.

This paper begins by providing some background on the simulation environment and the agents that were modified for this project. A detailed description of the social control mechanism is then given. Next will be the results of our tests, followed by conclusions drawn from this research and directions for future work.

2 Background

This section begins with a description of the basic agent architecture employed in the experiments. Section 4 describes the extensions made to this architecture for these experiments. After the basic architecture is described, the SimWorld simulation environment is introduced.

2.1 Agent Architectures

The basic agent architecture employs a schema-based approach to information processing and behavior selection. Inputs include sonar, smell, and touch receptors. Sonar is used to detect the presence of other agents, scent receptors detect food sources, and collisions with other agents are detected by touch, as is contact with food. Effectors include motors for moving and turning and a mechanism for eating. The inputs from the sonar and scent receptors are force vectors representing the scaled sums of all vectors between the agent and entities of the appropriate type (agents and food). The weighted sum of these force vectors is mapped onto the motors to produce agent movements. Input from the touch receptors trigger either a reflexive retreat mechanism in the case of an immanent collision or an ingestion mechanism for food when in contact with a food source. The input from the sensors described here determine the activations of the effectors of the agents. They tend to move away from other agents (because collisions are fatal) and toward food sources.

2.2 The Simulation Environment

The simulation environment used for the experiments described below is SimWorld [5], an artificial life environment based on the SimAgent toolkit [7]. SimWorld provides a two-dimensional surface in which agents forage for resources while avoiding hazards. Hazards can include obstacles and other agents. Obstacles are defined as either static or moving, and are specified at the beginning of an experimental run. A collision with an obstacle is fatal to an agent. Resources (e.g., food and water) are typically generated at random intervals in random locations, although it is possible to define regions in which particular resources are found. Food and water sources yield a fixed amount of energy and water when consumed. Agents use resources both for moving and for bodily maintenance. Movement costs agents energy and water as a (typically quadratic) function of their speed. The simulation setup used for the current paper did not include obstacles, and the agents did not require water, simplifying the conflicts to single-resource problems. Section 4.2 describes the experimental setup in more detail.

After agents reach a configurable procreation age, they are able to reproduce when their resources reserves are

C_F	Cost of fighting
C_R	Cost of retreating
FS	Food Source value
P_{GF}	Probability of acquiring food after winning encounter
P_{OFm}	Probability of acquiring food within m steps after retreating
B_F	$P_{GF} \times FS$, utility of winning an encounter
B_R	$P_{OFm} \times FS$, utility of retreating (losing)
$EV_F(n+1)$	$C_F \times n + B_F + C_F$, utility of winning an $n+1$ round game
$EV_R(n+1)$	$C_F \times n + B_R + C_R$, utility of losing an $n+1$ round game

Table 1: Analysis Terminology

high enough. Reproduction is asexual, and there is a recuperation period after birth before an agent can reproduce again. SimWorld also contains a complex mutation mechanism that allows the architecture of the parent to be modified before transmission to the offspring. This allows new architectures and behaviors to emerge during the course of an experiment. For the experiments described here, however, no evolutionary mechanisms were employed.

3 Conflict Resolution Model

We investigated conflicts in a game-theoretic framework [4]. Agents that decide to fight are considered defectors, whereas agents that decide to retreat are considered cooperators. There are two kinds of games: those in which there is a food source present and those in which there is not. In the latter case, the agents are fighting over territory. In each game there are four possible outcomes. Both agents can decide to fight, both can decide to retreat, the first can fight and the second retreat, or the first can retreat while the second fights. Table 1 lays out some of the terms used in this analysis. The costs and benefits are ordered according to the following inequality:

$$C_F < C_R < 0 < B_R < B_F$$

The cost of fighting is much greater than the cost of retreating, but the potential benefits are also higher. The probability of acquiring food after winning an encounter is 0 when agents are competing for territory and 1 when the food source in contest is guaranteed to be present at the end of the fight. If the resource has a chance to escape, or to somehow spoil, the value will be somewhere in between. The probability of acquiring food after retreating depends on the density of food in the environment; when food is plentiful, there is less incentive to fight over it.

Encounters last until at least one agent decides to retreat, unlike the iterated prisoner's dilemma. Thus, encounters may last many rounds, if both agents continually decide to stay and fight. Fight-fight outcomes are most undesirable, because both agents pay the expensive cost of fighting (C_F) and neither agent gains any benefit (either B_F or B_R). Likewise, long encounters are undesirable, because the value of the resource in contest decreases relative to the cost paid to obtain it. Retreat-retreat outcomes are less costly, but are also undesirable, since agents may then be leaving behind some potentially valuable resource. The best outcome is the fight-retreat outcome in which one agent decides to retreat and the other to stay in the first round. Strategies that are more likely to lead to short encounters in which one agent reaps the reward should, therefore, have an advantage over other strategies.

The utility of winning an encounter that lasts $n+1$ rounds ($EV_F(n+1)$) is $C_F \times n + B_F + C_F$. The agent pays the cost of fighting for the first n rounds and in the final round receives the benefit (B_F) and pays one final fight penalty (assessed for whatever action chased its opponent away). This makes explicit the benefit of short encounters; the fewer rounds the encounter lasts, the fewer times the agent has to pay the fight cost. Similarly, the utility of losing an encounter that lasts $n+1$ rounds ($EV_R(n+1)$)

		Player 2	
		Retreat	Fight
Player 1	Retreat	$EV_R(n+1)$ (P2); $EV_R(n+1)$ (P1)	$EV_F(n+1)$ (P2); $EV_R(n+1)$ (P1)
	Fight	$EV_R(n+1)$ (P2); $EV_F(n+1)$ (P1)	$EV_F(n+1)$ (P2); $EV_F(n+1)$ (P1)

Table 2: Payoff Matrix for $n+1$ Round Agent Encounters

is $C_F \times n + B_R + C_R$. The agent pays the cost of fighting for the first n rounds, receives the benefit of retreating, and pays the cost of retreating. Using these values, we obtain the payoff matrix shown in Table 2. Note that this is not a zero-sum game; losers are not required to lose as much as their opponents gain.

Each agent's behavioral dosposition for conflicts can range from aggressive (likely to fight) to fearful (likely to retreat). These dispositions can be mapped onto the range from 0.0 to 1.0. D_S represents an agent's own disposition, while D_O represents an opponent's perceived disposition. In an encounter, the probability that an agent will decide to stay and fight P_F is based on these D values. For asocial agents, $P_F = D_S$ (i.e., the probability that it will choose to fight is just the measure of its behavioral disposition for conflict situations). Social agents determine their probability using the following equation:

$$P_F = \begin{cases} D_S + \left(\frac{(1-D_S)}{D_S}\right) \times (D_S - D_O) & \text{if } D_S > D_O \\ D_S - \left(\frac{D_S}{D_O}\right) \times (D_O - D_S) & \text{if } D_S \leq D_O \end{cases}$$

These equations map the difference in D values for the two agents into the space available to increase or decrease social agents' probabilities of fighting depending on whether the agent's disposition value D_S is higher than its opponent's or not. In this way, social agents are able to make more "informed" decisions in conflicts, and can minimize costly fight-fight and flee-flee outcomes.

Given the fact that fight-fight outcomes are so expensive, we decided to map out the probability-spaces for three types of encounters: asocial-asocial encounters, social-social encounters, and mixed asocial-social encounters. We compared the probability of fight-fight outcomes for each combination. Figure 1 compares the fight-fight probability spaces (i.e., the probabilities for all combinations of D-values) for asocial-asocial and social-social encounters. There is a small area in which social-social encounters are more likely to lead to expensive fight-fight outcomes, however, for most combinations asocial-asocial encounters are more likely to lead to fight-fight outcomes. Figure 2 compares the fight-fight probability spaces for asocial-asocial and mixed asocial-social encounters. When the asocial participant in the mixed encounter has a higher D-value, asocial-asocial encounters were more likely to lead to fight-fight outcomes than mixed encounters, whereas mixed encounters had a higher probability in cases in which the social agent had a higher D-value. However, the differences in the latter cases is significantly less than the differences in the former case. Comparing mixed and social-social probability spaces, we find a similar situation (Figure 3). Social-social encounters have a higher probability of fight-fight outcomes than mixed encounters when the social member of the mixed encounter has a higher D-value than the asocial member, while mixed encounters have a higher probability in the reverse case. Once again, however, the difference in the former cases is less than that in the latter cases, so the net result is that mixed encounters are more likely overall to lead to fight-fight outcomes. Overall, this produces a progression in which asocial-asocial encounters tend to lead to more undesirable outcomes than mixed encounters, which in turn have a higher tendency to lead to undesirable outcomes than social-social encounters.

Retreat-retreat outcomes are less expensive than fight-fight outcomes, but they are still expensive compared to fight-retreat outcomes, since neither agent has a chance to obtain the higher benefit of winning an encounter (B_F). We again mapped out the probability spaces, this time for retreat-retreat outcomes of the three varieties of encoun-

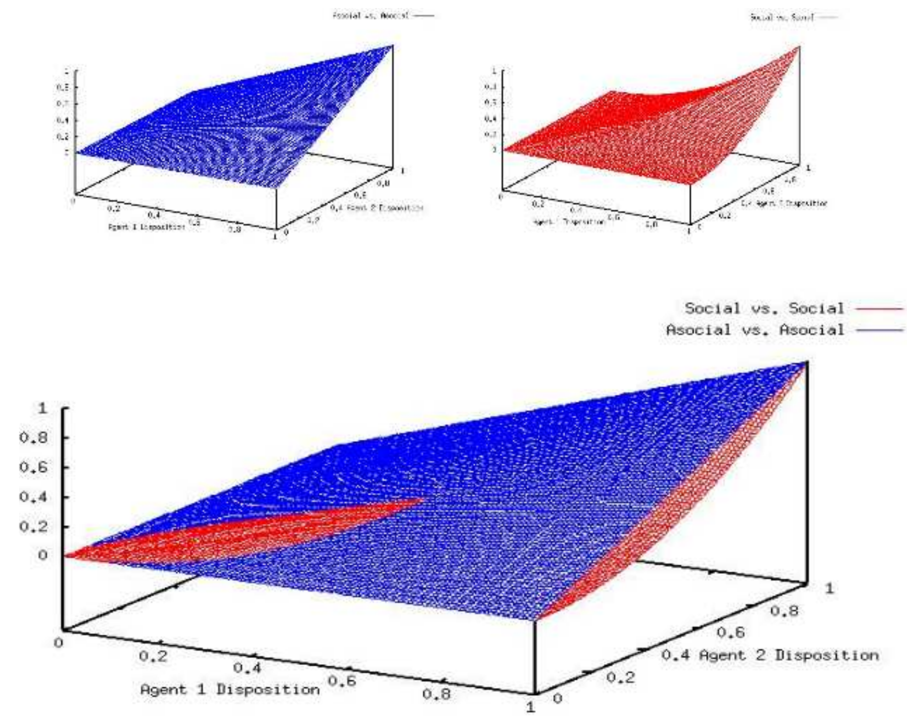

Figure 1: Homogeneous Asocial Fight-Fight vs. Homogeneous Social Fight-Fight Probabilities

ters. In Figure 4 we see that there is only a small area in which social-social encounters are more likely to lead to retreat-retreat outcomes than asocial-asocial encounters. Figure 5 compares retret-retreat probability spaces for asocial-asocial and mixed encounters. This is another example in which each variety of encounter has higher probabilities in half of the cases, but the difference between asocial-asocial and mixed when the social agent has a higher D-value is much greater than the converse, leading to a net probability of retreat-retreat outcomes for asocial-asocial encounters than for mixed encounters. Figure 6 depicts a similar outcome for the comparison between mixed and social-social encounters; mixed are more likely to lead to flee-flee outcomes than social-social. Once again there is an overall ordering from asocial-asocial to social-social in terms of beneficial outcomes.

Given the orderings of fight-fight and retreat-retreat probabilities, we can produce an ordering for mixed (beneficial) outcomes. Social-social encounters are more likely to lead to mixed outcomes than social-asocial encounters, while social-asocial encounters are more likely to lead to mixed outcomes than asocial-asocial encounters. Thus, social agents should engage in fewer expensive fight-fight or retreat-retreat outcomes and in more beneficial mixed outcomes. It is important to note, however, that these probabilities hold only when dispositions are distributed across the disposition-space (e.g., random or gaussian distributions, or distributions that map the entire space in a grid-like fashion). Other distributions may be denser in areas that benefit asocial agents.

4 Experimental Design

This section describes the implementation of the probabilistic conflict resolution mechanism, as well as the social extension that employs disposition sense. It then describes the behavior of the different kids of agents. Fi-

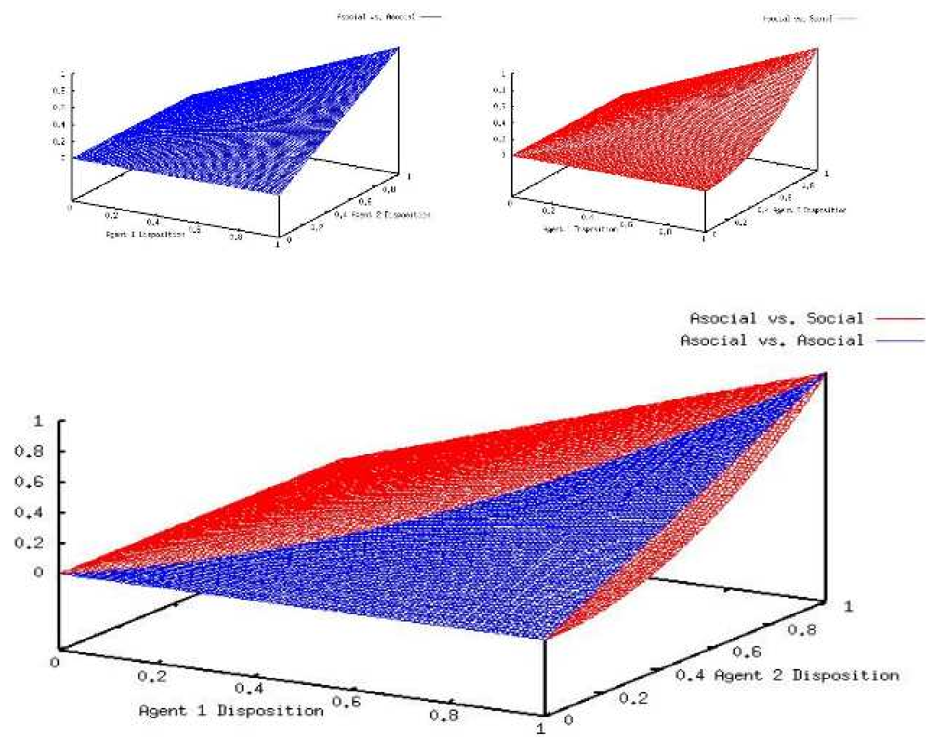

Figure 2: Homogeneous Asocial Fight-Fight vs. Heterogeneous Asocial-Social Fight-Fight Probabilities

nally, it lays out the experimental setup used to test the efficacy of the new agents.

4.1 Agent Design

We conducted a series of experiments to test the effects of the use of dispositional cues in agents. Implementing the probabilistic mechanism itself was the first step. Collisions between agents are fatal to both, so part of the reactive layer of each agent's architecture is a retreat reflex that causes the agent to turn and move away quickly when an immanent collision is detected. There is a penalty associated with the reflex since it increases the speed of the agent as it retreats and movement cost is a quadratic function of speed. The probabilistic mechanism makes use of this reflex mechanism to control the agent's response to other affective agents. During an encounter, each agent is assigned a probability of suppressing this reflex to stay where it is in hopes of obtaining whetever resources are nearby. This probability is equal to the D-value for asocial agents, while for social agents it is determined according to the equation given in Section 2.2.

Asocial agents whose D-values are high have a high probability of deciding to fight, whereas those with low D-values have a high probability of deciding to retreat. Social agents with high D-values are also likely to fight, however this probability is increased somewhat when they are confronted by a less aggressive opponent and decreased when faced by a more aggressive opponent. Similarly, social agents with low D-values are not very likely to decide to fight, but are somewhat more so when their opponent is even less aggressive than they, and somewhat less so when their opponent is more aggressive. We also defined two other agent types which are simply special cases of the asocial agents: timid and aggressive agents. Timid agents are assigned a D-value of zero, leading them

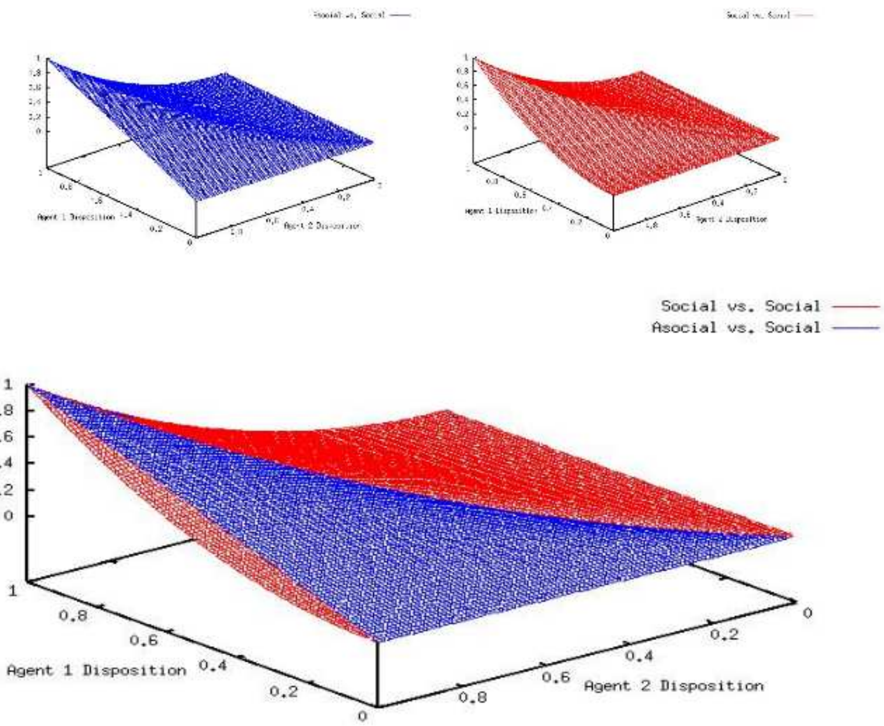

Figure 3: Heterogeneous Asocial-Social Fight-Fight vs. Homogeneous Social Fight-Fight Probabilities

to follow an always-retreat strategy. Aggressive agents are given a D-value of one, so they employ an always-fight strategy.

4.2 Experimental Setup

These probabilistic conflict resolution mechanisms were tested to determine what benefit they provide, if any. In each of the experiments described below, the world is continuous and unlimited, while the food sources are limited to a square of 1440 by 1440 units. Food is distributed randomly throughout this region, and agents are free to wander off, but must return in order to eat. Food is created with a probability of 0.5 per cycle. The measure of performance for each of the experiments is survival (i.e., the number of surviving agents at the end of a simulation run, averaged over 40 simulation runs with different random initial conditions). Experiments proceed for 10,000 simulation cycles, or roughly 30 generations. No mutation is employed; these experiments measure only the relative performance of these agents, not the potential for evolutionary trajectories between them. Each simulation starts with 20 agents of each type participating in that test (e.g., in homogeneous asocial tests, 20 asocial agents start, whereas in mixed asocial-social environments, 20 of each type start). Agents reproduce in roughly 350-cycle generations, and live roughly 500 cycles each, on average.

Based on the analysis given in Section 2.2, we predict that homogeneous social environments will outperform homogeneous asocial environments, that is, that there will be more survivors in social-only environments than in asocial-only environments. This is because asocial-asocial encounters are more likely to result in expensive fight-fight outcomes, leading to more prolonged encounters with greater cost. Also, the flee-flee probability is greter for asocial agents, so they are more likely to leave a resouce behind than social agents.

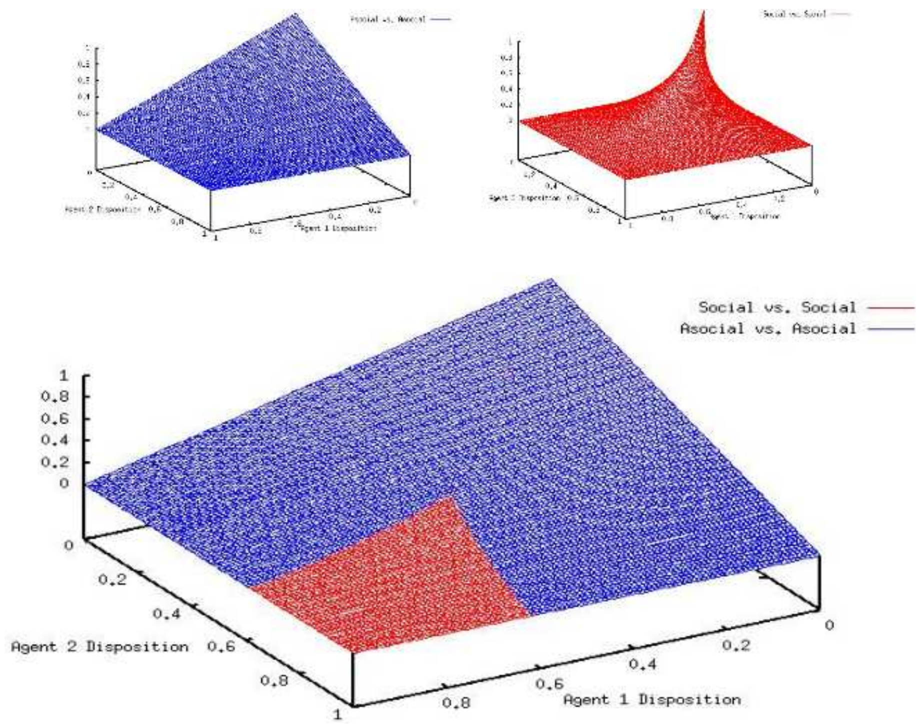

Figure 4: Homogeneous Asocial Flee-Flee vs. Homogeneous Social Flee-Flee Probabilities

We also predict that in mixed social-asocial environments, there will tend to be more social agents surviving than asocial agents. Social agents should have an advantage, because mixed encounters tend to have fewer fight-fight outcomes than asocial-asocial encounters, as well as fewer retreat-retreat outcomes, and social-social encounters are better than mixed in both cases. This means that overall, social agents should be part of fewer fight-fight and flee-flee outcomes than asocial agents, and should, therefore, perform better.

5 Results

We conducted experiments with homgeneous environments for each of the four agent types, and for each combination of two-type mixed environments. Figure 7 depicts our results. Again, these results indicate the average number of survivors over 40 experimental runs. Turning first to the homogeneous results, we find that timid, asocial, and social agents perform about the same, while aggressive agents perform much worse, with roughly one third as many survivors, on average, as the other types. This contradicts our first prediction, as social agents did not outperform asocial ones. We speculate that this is because of the environmental constraints; it may be that a more densely populated environment would lead to more conflicts, thus allowing social agents to show their advantage. Alternatively, it may be that fewer food resources would more effectively highlight inefficiencies in asocial agents' conflict resolution mechanism, allowing social agents to perform better. This is an avenue we plan to pursue.

Timid agents failed to survive at all in both asocial and social mixed environments. In each case, the timid agents produced a total of only 100 or fewer agents on average, whereas their competitors produced nearly 1000, indicating that the timid agents died out fairly early in the simu-

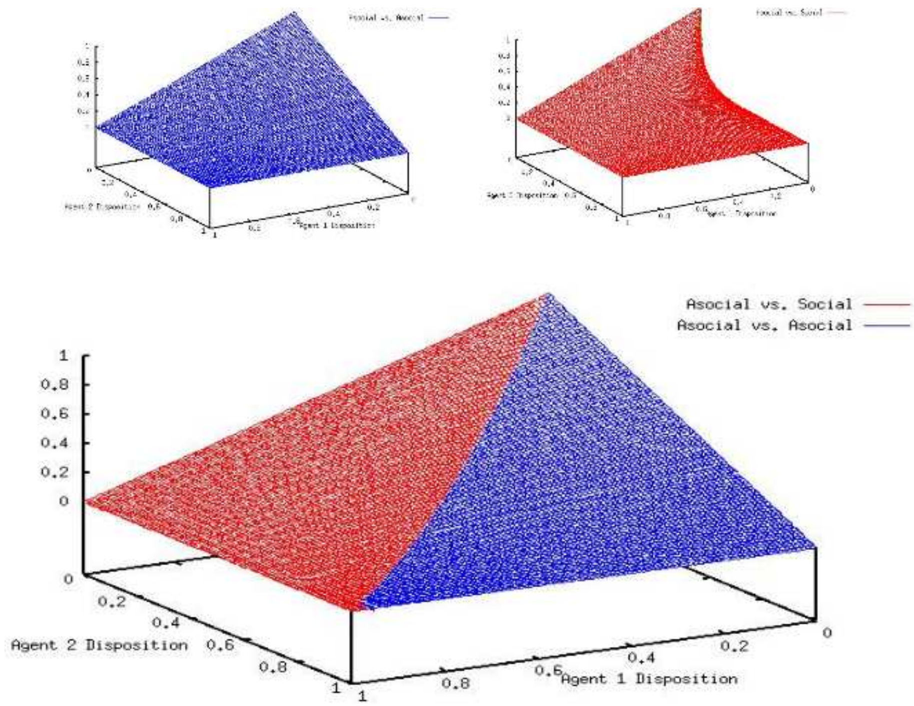

Figure 5: Homogeneous Asocial Flee-Flee vs. Heterogeneous Asocial-Social Flee-Flee Probabilities

lation runs. The fact that their competitors performed similarly to their homogeneous performance also indicates that they were alone in the environment for a large portion of the simulations. Against aggressive agents, however, timid agents performed very well. They achieved survival very near their homogeneous rates, whereas the aggressive agents averaged only a handful of survivors (around 3), far fewer than in the homogeneous environment.

The asocial-social mixed environment is of most interest here. On average, nearly twice as many social agents survived these simulations as did asocial agents. The difference is significant. This supports our prediction, and indicates that there is a strong advantage to social conflict resolution mechanisms. The high standard deviations for both agent types indicate that there were likely some environments in which very few social agents survived, and even more in which very few asocial agents survived. Table 3 gives the results for each simulation run of asocial-social environments. In 16 of the simulations fewer than 10 asocial agents survived, while in 5 of the simulations fewer than 10 social agents survived. In each case, the other agent type had 39 or more survivors.

Finally, aggressive agents fared poorly against both asocial and social agents in mixed environments. Their performance was slightly better against these types than the timid agents, but in both cases they averaged only about one survivor over the 40 simulation runs.

6 Conclusions and Future Work

Our results indicate that there is benefit to employing social strategies (i.e., strategies that take into account information about others in the environment) when making decisions in conflicts. This benefit accrues as a result of social agents making better decisions in conflicts based on the information they obtain from cues of their opponents' dispositions. It would seem as though the advantage is substantial enough that it is likely that asocial environ-

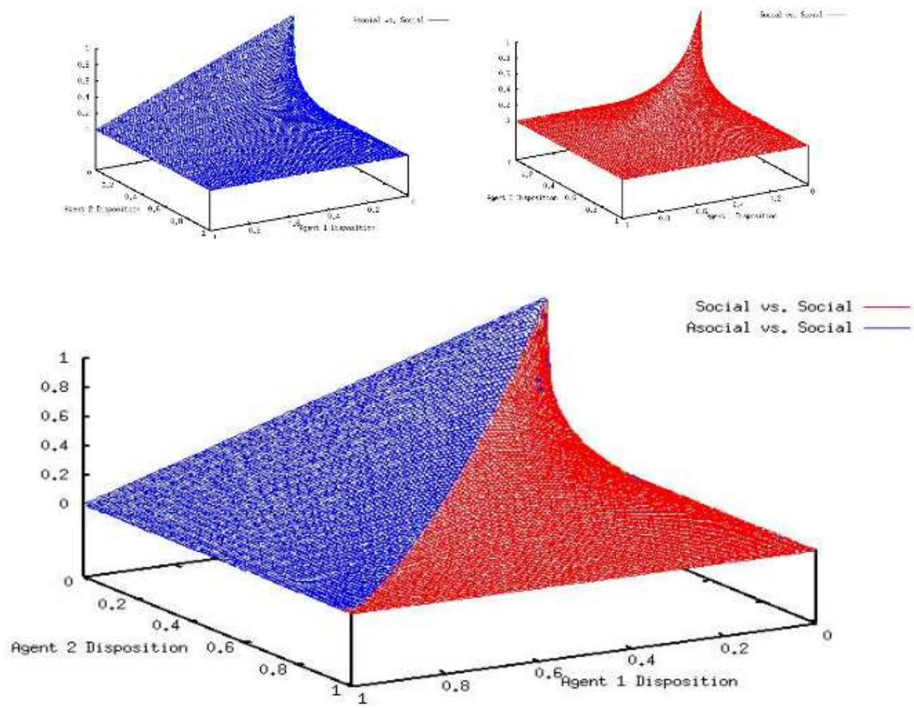

Figure 6: Heterogeneous Asocial-Social Flee-Flee vs. Homogeneous Social Flee-Flee Probabilities

Exp.	Asocial	Social	Exp.	Asocial	Social	Exp.	Asocial	Social	Exp.	Asocial	Social
1	44	5	11	30	23	21	32	19	31	43	7
2	3	51	12	4	54	22	40	14	32	23	20
3	0	50	13	40	13	23	19	40	33	49	3
4	1	48	14	10	37	24	4	40	34	13	38
5	0	50	15	6	39	25	35	21	35	31	18
6	4	44	16	0	47	26	50	2	36	0	52
7	16	33	17	15	32	27	28	24	37	36	11
8	2	54	18	6	48	28	0	45	38	40	14
9	0	50	19	9	39	29	35	14	39	14	40
10	16	32	20	22	27	30	49	2	40	0	48

Table 3: Individual Experiment Results for Asocial-Social Environments

ments would be invaded by social agents. It is interesting to note that these mechanisms are social in another sense, as well: although each agent is working selfishly to further its own ends, the social strategy leads to a form of cooperation in which less aggressive agents "help" more aggressive opponents by not forcing them to enter into

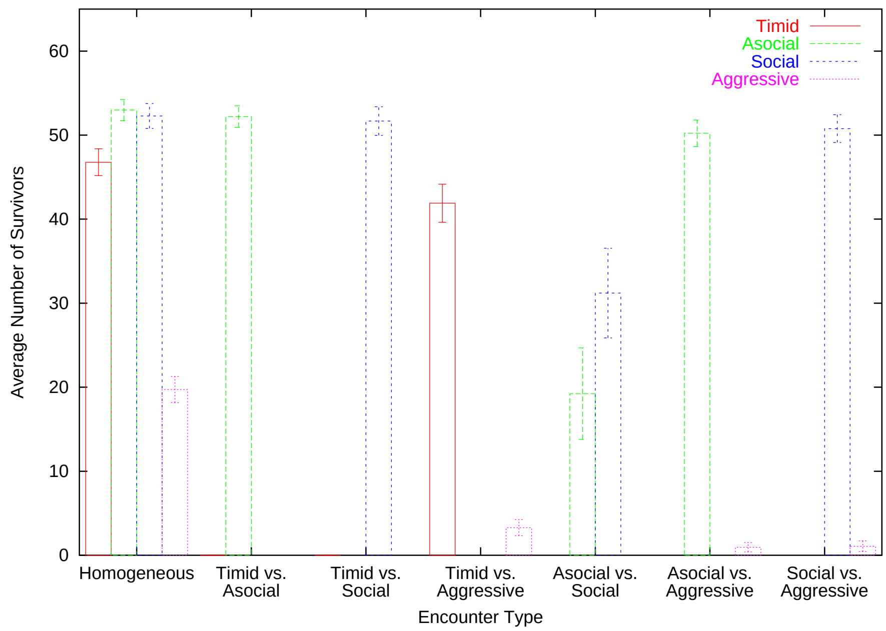

Figure 7: Experimental Results (Homogeneous and Heterogeneous Environments)

prolonged encounters. The less aggressive agent still benefits (by not paying as much for an encounter it is likely to lose), but the more aggressive opponent stands to benefit more.

Purely timid strategies work very well in isolation, but any agent that decides to fight part of the time will successfully invade timid environments and, given the results shown above, may decimate the timid population. Excessive aggression, however, appears to be too expensive in any mixed environment. This is probably a result of aggressive agents killing each other off via long encounters; encounters with agents even slightly likely to retreat will most often end before one of the agents die, but since aggressive agents will always choose to fight, their encounters will continue until one dies.

The work described here is preliminary. We intend to further explore the effect of the environment on agent performance, as mentioned in Section 5. Allowing agents' behavioral dispositions to change as a result of their experiences in conflicts could lead to interesting behaviors; agents who win encounters could become more aggressive, while those who lose could become more fearful. We are interested in attempting to evolve the social mechanism, perhaps by having the social agents start out being very bad at estimating their opponents' dispositions and seeing if they get progressively better. It would also be interesting to compare the perfomance of other conflict strategies, such as Tit-for-Tat (adapted for the kind of game described above), against the social agents described here. Finally, we are currently investigating cost in the context of agent architectures [6], and will be examining the issue of cost of the social mechanisms described above. While we have shown that they are beneficial, they will also come with a cost. Assessing that cost will be crucial to determining whether the benefits of such mechanisms will outweigh their costs.

References

[1] Shelley A Adamo and R. T. Hanlon. Do cuttlefish (cephalopoda) signal their intentions to conspecifics

during agonistic encounters? *Animal Behavior*, (52):73–81, 1996.

[2] Kristi R. DeCourcy and Thomas A. Jenssen. Structure and use of male territorial headbob signals by the lizard *Anolis carolinensis*.

[3] Hans A. Hofmann and Klaus Schildberger. Assessment of strength and willingness to fight during aggressive encounters in crickets. *Animal Behavior*, (62):337–348, 2001.

[4] John Maynard Smith. *Evolution and the Theory of Games*. Cambridge University Press, Cambridge, 1982.

[5] Matthias Scheutz. The evolution of simple affective states in multi-agent environments. In Dolores Cañamero, editor, *Proceedings of AAAI Fall Symposium*, pages 123–128, Falmouth, MA, 2001. AAAI Press.

[6] Matthias Scheutz and Paul Schermerhorn. Steps towards a theory of possible trajectories from reactive to deliberative control systems. In *Proceedings of the 8th Conference of Artificial Life*.

[7] A. Sloman and R. Poli. Sim_agent: A toolkit for exploring agent designs. In M. Wooldridge, J. Mueller, and M. Tambe, editors, *Intelligent Agents Vol II (ATAL-95)*, pages 392–407. Springer-Verlag, 1996.