Deriving Multi-Agent Coordination through Filtering Strategies

Eithan Ephrati* and Martha E. Pollack*' and Sigalit Ur t
Department of Computer Science* and Intelligent Systems Programt
University of Pittsburgh
tantush@cs.pitt.edu, pollack@cs.pitt.edu, sigalit@pogo.isp.pitt.edu

Abstract

We examine an approach to multi-agent coordination that builds on earlier work on enabling single agents to control their reasoning in dynamic environments. Specifically, we study a generalization of the *filtering* strategy. Where single-agent filtering means tending to bypass options that are incompatible with an agent's own goals, multi-agent filtering means tending to bypass options that are incompatible with other agents' known or presumed goals. We examine several versions of multi-agent filtering, which range from purely implicit to minimally explicit, and discuss the trade-offs among these. We also describe a series of experiments that demonstrate initial results about the feasibility of using multi-agent filtering to achieve coordination without explicit negotiation.

1 Introduction

Distributed Artificial Intelligence (DAI) is concerned with effective interactions, and with the mechanisms by which these interactions can be achieved. Broadly speaking, two main approaches have been proposed in the literature. The first involves explicit coordination; agents are designed to reason about their potential interactions, and negotiate with one another as needed. Examples of this approach include [Ephrati and Rosenschein, 1991; 1993; Kraus, 1993; Zlotkin and Rosenschein, 1993]. A difficulty with explicit coordination and negotiation is that it can be extremely time-consuming, and in dynamic environments, agents may not be able to afford the time required. The second approach involves implicit coordination; agents are designed to follow "local" rules of behavior that lead to their acting in apparently coordinated ways; see, for example, [Shoham and Tennenholtz, 1992; Goldman and Rosenschein, 1994]. This approach is motivated in part by a belief that one can design simple rules that are easy for an agent to follow, yet result in coordination.

In this paper, we take the second approach, examining an implicit coordination strategy. The strategy we study, *multi-agent filtering,* is an extension of a single-agent strategy for controlling reasoning in dynamic environments. The notion of single-agent filtering derives from the work of Bratman [Bratman, 1987]; it involves an agent committing to the goals it has already adopted, and tending to bypass (or "filter out") new options that would conflict with their successful completion [Bratman *et* a/., 1988; Pollack, 1992; Pollack *et* a/., 1994]. We and others have studied the effectiveness of filtering in domains with various characteristics[Kinny and Georgeff, 1991; Pollack *et* a/., 1994].

The original filtering strategy was designed as a method for an individual agent to focus its reasoning in a dynamic, but not necessarily multi-agent, environment. Here, we generalize this strategy to multi-agent environments. Where single-agent filtering means tending to bypass options that are incompatible with an agent's *own* goals, multi-agent filtering means tending to bypass options that are incompatible with *any* agent's known or presumed goals.

We examine several forms of multi-agent filtering, which range from implicit, in which agents have rules of legal action that lead to their avoiding conflict without ever reasoning explicitly about one another's goals, to minimally explicit, in which agents perform shallow reasoning to assess whether their actions are incompatible with the likely intended actions of other agents. In no cases do the agents engage in any explicit negotiation.

It seems clear that if one agent, call it A, avoids interfering with the goals of a second agent, call it *B,* then *B* will be better able to achieve its goals. But what about *A?* Won't its performance be worse, because it is subject to additional constraints on its behavior? If *A* is the only multi-agent filterer, it seems likely that its performance will suffer, but if *A* and *B* are both multi-agent filterers, then the effect is less obvious. What we need to ask is whether the advantage that *A* derives from B's multi-agent filtering is sufficient to override any penalties *A* receives from its own multi-agent filtering. And we need to ask the same thing about *B.* The central questions we address in this paper are thus: What happens in multi-agent environments in which all (or most or few or none) of the agents are multi-agent filterers? And do these effects depend, in any interesting and identifiable ways, on properties of the domain? To address these questions, we conducted a series of experiments using a multi-agent version of the Tileworld system [Pollack *et* a/., 1994; Joslin *et al.,* 1993], an abstract testbed for studying behavior in dynamic environments.

In the next section, we review the theory of filtering, and discuss its generalization to the multi-agent case. A brief overview of the experimental platform is provided in Section 3. Sections 4 through 6 present our experimental results on multi-agent filtering to date: Section 4 describes experiments with extremely bold multi-agent filterers, Section 5 describes the effectiveness of enriching the filter with an override mechanism, and Section 6 considers the implications of multi-agent filtering in environments of self-motivated rational agents. The most interesting and surprising result is that, at least for the simple, abstract environments so far studied, multi-agent filtering is a dominant strategy: no matter what proportion of the agents in some environment choose not to filter, those that do filter perform better. We summarize our results in Section 7.

2 Filtering and Multi-Agent Filtering

Our work on multi-agent filtering derives from our earlier work on filtering as a strategy that individual agents can use to focus their reasoning. The notion of filtering derives from the work of Bratman [Bratman, 1987], who argued that it is useful for resource-limited agents to adopt and commit to plans, tending to bypass, or "filter out" from consideration new options that would conflict with the successful completion of those existing plans. On this view, an agent's existing intentions frame its subsequent reasoning: the agent can focus on ways of achieving its current goals, and can, in general, bypass deliberation about the myriad of options that are incompatible with its current goals.

Typically, filtering is augmented with some kind of override mechanism that enables the agent to deliberate about options that are *prima facie* important, even when they are incompatible with pre-existing options. A central challenge for the designer of an agent with a filtering mechanism is to construct an override component that embodies the right degree of sensitivity to the problems and opportunities of the agent's environment.

Note that the filter-plus-override mechanism does not, by itself, determine what intentions the agent will adopt. When an option survives the filter, either because it is deemed compatible with existing plans or because it triggers an override, it is then subject to a deliberation process that selects the actions towards which the agent will form intentions. In other words, it is the deliberation process that performs the type of decision-making that is the focus of traditional decision theory. The filtering mechanism frames particular decision problems, which the deliberation process solves.[1]

In fact, in his original discussion of the role of commitment in resource-limited reasoning, Bratman suggested *two* advantages that accrue to an agent who uses a filtering strategy. First, filtering can help the agent focus its reasoning, as described just above; this has been a main concern in our previous work. Second, filtering can help multiple agents coordinate their activities, because each agent can count on the other agents carrying out the plans to which they have committed. Following this observation, we hypothesized that it was possible and desirable to extend the strategy of filtering for multi-agent environments so that it could serve as a technique for coordination as well as for control. The basic idea is straightforward: not only should agents tend to bypass consideration of options that conflict with their own goals, but they should also tend to bypass consideration of options that conflict with the goals of other agents.

The precise interpretation of conflict depends on the relationships that hold among the goals of the agents in the environment. For example:

1. The agents may have one common goal, but individual and distinct subgoals. In this case, avoiding conflict means avoiding actions that make it more difficult for another agent to achieve its subgoals. This situation underlies the work on social laws [Moses and Tennenholtz, 1992], where the (implicit) common goal is the maximization of the designer's reward (through the individual activities).
2. The agents may have one common goal, but potentially overlapping subgoals. Here, "conflict" can mean achieving (or helping in the achievement of) a subgoal for another agent. To some extent, this situation underlies the work on cooperative state-changing rules [Goldman and Rosenschein, 1994].
3. The agents have distinct, possibly conflicting, goals. There may be competition not only for resources to achieve goals, but also for the goals themselves. This case would appear to pose the greatest challenge to a multi-agent filtering strategy.

As we discuss below, in conducting our experiments we addressed each of these variants.

3 Experimental Platform

The Tileworld testbed is a Lisp-based tool that was developed to support controlled experimentation with agents in dynamic environments. For the current project, we built a multi-agent version of the Tileworld system, called MA-Tileworld. We first briefly describe the MA-Tileworld, and then discuss some details of the multi-agent filtering strategies.

3.1 The Multi-Agent Tileworld System

Like the original Tileworld, MA-Tileworld is an abstract, dynamic, simulated environment, with embedded agents. It is obviously, and intentionally, a highly artificial environment. In keeping the environment divorced from any realistic application, our goal has been to provide a tool that allows researchers concerned with *any* application to focus on what they consider to be key features of that application's environment, without the confounding effects of the actual, complex environment itself. We have, in other words, traded realism—in the short run, at least—for sufficient control to allow for systematic experimentation. (Cf. [Hanks *et al.*, 1993]).

[1] The filtering mechanism is meant to be only one part of a rich meta-level control structure. Thus, for example, agent designers may also want to include mechanisms to filter from full consideration options that are especially unpromising, even if they are compatible with existing plans.

The MA-Tileworld environment consists of a rectangular, two-dimensional grid, on which are located a variety of objects, including holes, tiles, obstacles, etc., and simulated agents. A *trial* is a single run of the MA-Tileworld system, defined by its duration and its *experimental condition,* user-specified parameter settings that establish agent and environment conditions. Trials for the same experimental condition will, in general, differ from one another, because the system's performance depends stochastically on the parameter settings.

An example of an experimental condition is the rate at which objects (tiles, etc.) appear and disappear during the trial. When a MA-Tileworld agent successfully fills a hole with tile(s), it receives an reward, the size of which depends on the type of tiles that were used to fill the hole. The agent may carry one or many tiles at a time; however, the more tiles it carries, the more energy (or "gas") it expends; agents must be concerned not only with filling holes but also with maintaining sufficient energy. For further details see [Pollack *et al,* 1994; Joslin *et* a/., 1993].

The agents that are embedded in the MA-Tileworld observe a filtering strategy. That is, they bypass consideration of options that are incompatible with their own goals except when those options trigger an override. The question of how easy it is for an option to trigger an override is put under the control of the experimenter, who specifies a override threshold t. When the agent recognizes an option *o* that is incompatible with its existing goals, it computes an estimated value V_0 for o. For *o* to trigger a filter override, V_0 must exceed the computed value of the current intentions by at least *t.* Thus, the lower the threshold, the more likely the agent is to allow options to pass through the filter, and hence the more likely it is to engage in deliberation: in the terminology of [Bratman *et* a/., 1988], the more *cautious* the agent will be. In contrast, we say that an agent with a high threshold is *bold.* The threshold can be negative, to allow for full deliberation even about options that appear, upon original estimation, to be less valuable than the intention(s) with which they conflict. In all the experiments described in this paper, the estimated value of a fill-hole option was set equal to the raw score of the hole (i.e., the highest score that will be achieved if the hole is successfully filled with the best tiles); the estimated value of getting gas is a function of the agent's current gas level; and the estimated value of the other options, such as stockpiling tiles and wandering, is a low constant.

3.2 Operationalizing Multi-Agent Filtering

The filtering process as just described disposes agents to filter from consideration options that are incompatible with their *own* existing intentions. For our current purposes, we generalized this strategy to the multi-agent case. Recall that our goal is to investigate filtering as an *implicit* coordination strategy, i.e., a set of easily followed rules for behavior that result in coordinated action by multiple agents inhabiting some environment. Thus, our implementation of multi-agent filtering had to observe strict limits on the amount of reasoning that each agent needs to do about the others. We implemented and investigated three different filtering methods:

1. *Static geographical boundaries:* The environment is divided into geographical regions. Each agent is assigned a particular region, and filters out options to fill holes in other regions. Because no two agents are assigned the same region, filtering automatically leads to conflict avoidance.
2. *Dynamic geographical boundaries:* The environment is not partitioned *a priori.* Instead, each agent filters out options to fill holes that are nearer to some other agent than to itself. This leads to conflict avoidance, since every hole is nearest to a single agent.
3. *Intention posting:* The first two cases are clearly "implicit": the agents can follow those filtering strategies without any computations that directly take into consideration the goals of other agents. This third approach is slightly more explicit: here, agents post to a globally accessible data structure each intention they form to fill a hole. Agents then filter from consideration hole-filling options that have already been declared by other agents.

Note that in all cases, hole-filling is the only type of option that may lead to multi-agent filtering in these experiments; options like getting gas and stockpiling are never seen to be incompatible with another agent's goals. It is also important to remember that just because an option is subject to filtering, it does not mean that that option will *necessarily* be discarded from consideration. What it does mean is that it will be further considered only if it triggers a filter override, i.e., it is deemed *prima facie* to be worthy of deliberation despite the fact that it conflicts with another agent's goals. If deliberation does occurs, it may result in adopting a new intention towards the option, but it may also result in bypassing the option.

In our first set of experiments (Section 4), we studied agents that were extremely bold: they can be viewed either as having no override mechanism at all, or, equivalent,, as having an override mechanism with an infinite threshold value. (This is, in fact, how they were implemented.) For these agents, options that were deemed incompatible with the options of other agents never triggered an override and thus were never subject to deliberation. Note that the agents had a single override mechanism, which did not distinguish between options that conflict with their own existing plans and those that conflict with the plans of other agents. Thus, the extremely bold agents also filtered out all options that were deemed incompatible with their own goals. As we will describe, the use of a multi-agent filtering strategy, even such a rigid (bold) one, improved the agents' overall performance. In a second set of experiments (Section 5), we explored the effect of making the multi-agent filtering process more flexible, by including particular override strategies. The results demonstrate cases in which overriding is important.

In all these experiments, what we measured was the total *effectiveness* of the agents. Effectiveness, for the single-agent Tileworld, is defined to be a normalized measure of an agent's score: the score it actually received during a trial divided by the total full score of all the holes that appear during the trial. In the multi-agent case, we compute (average) effectiveness by summing the scores received by all the agents and dividing by the total score of all the holes that appear during the trial. Thus defined, (average) effectiveness is an appropriate measure for first two types of multi-agent settings mentioned above: those in which all the agents share a common goal, and may or may not have potentially overlapping subgoals. For the third type of setting, in which the agents may have competing goals, we measured the average effectiveness of agents with each strategy.

4 Rigid Filtering

In the first set of experiments, the strategy of extremely bold multi-agent filtering was studied. We examined each of the three filtering mechanisms defined above (static geographical boundaries, dynamic geographical boundaries, and intention posting). A fourth experimental condition involved extreme caution: agents deliberated about all options that appeared in the Tileworld grid, regardless of whether they were potentially or actually incompatible with the goals.

This set of experiments was aimed at examining the effectiveness of multi-agent filtering in various environmental conditions. In the Tileworld environment, the most influential parameters are the average rate of change in the world (the "world speed"), and the average number of holes available to fill at any time. We therefore conducted two experiments; in the first one we varied the rate of change in the world, while in the second we held world speed constant and varied the average number of holes.

Both experiments involved four agents on a 20x20 grid. Their "thinking" speed was set to a baseline rate established in our earlier experiments [Pollack *et al.*, 1994]. The agent's "moving" speed was varied directly with the rate of world change, because our interest is in the relation between the agent's computation cycle time and the degree of dynamism in the world, not between the speed at which the agent can move and the degree of dynamism in the world. Generated holes were were randomly assigned a score of between 25 and 75, again consistent with baselines established in our earlier experiments.

For each experimental condition, we ran 51 trials, where the length of each trial was 80,000 clock ticks (which is equivalent to the amount of time it takes an agent to move 400 units of distance). The number of trials per condition and the length of each trial are the same as in our earlier, single-agent experiments.

Figure 1 describes the case where the the world speed was changed. There were a total of 44 experimental conditions (4 filtering strategies and 11 rates of dynamism). The x-axis shows the world speed: experimental results for the least dynamic worlds are shown at the origin, and speed increases across the x-axis. Average effectiveness is plotted on the y-axis. As can be seen, all three multi-agent filtering strategies result in better performance than no filtering, regardless of the rate of change

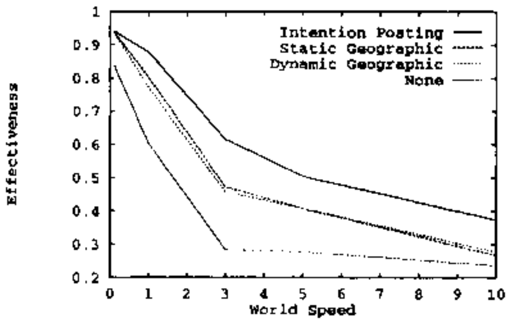

Figure 1: Bold Filtering with Varied World Speed

in the world. Among the filtering strategies, a society of intention posting agents performs better than a society of agents using either of the geographic strategies, both of which result in quite similar performance.

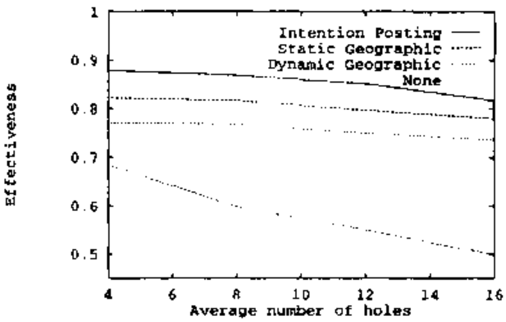

Figure 2: Bold Filtering with Varied Number of Holes

Figure 2 shows the experiment in which world speed is held constant at a baseline level, while the average number of holes in the world is varied between four and sixteen. Effectiveness is again plotted on the y-axis. The results are similar to our first experiment: regardless of the average number of holes in the environment, the multi-agent filterers do best, and the best of the multi-agent filtering strategies is intention posting.

We believe that the intention posting outperforms the other two filtering strategies because it is more accurate: agents will avoid goals that other agents actually intend to pursue, rather than avoiding goals that others might pursue. In addition, intention posting is computationally simpler than the dynamic geographic strategy: instead of calculating whether the goal is within the the agent's territory, the filter is triggered by an immediate lookup operation. (The additional computational overhead of the dynamic geographic strategy may also explain why it performs somewhat worse than the static geographical strategy.) These results suggests that in some cases, minimally explicit coordination strategies like intention posting may outperform implicit coordination.

5 Filtering with Overriding

The first set of experiments involved extremely bold agents, who never deliberated about options that were deemed to be incompatible with the goals of other agents. But, as we noted above, such unconditional acceptance of other agents' goals may lead to inefficiencies in performance, just as unconditional acceptance of one's own goals may lead to inefficiencies. This is why we include an override mechanism along with filtering. In [Pollack *et al.*, 1994], we described conditions in which overriding was beneficial for the single-agent case. Our next step was thus to investigate how overriding affects performance in the multi-agent setting.

What exactly does overriding amount to in the MA-Tileworld? Recall that multi-agent filtering results in agents bypassing consideration of options that they believe would interfere with the goals of other agents; in other words, multi-agent filterers avoid "stepping on one another's toes." But sometimes they *should* step on one another's toes. The reason is that the world is dynamic: opportunities don't last forever. Sometimes the possibility of successfully filling a hole may disappear by the time it could be filled by the responsible agent—either the one in whose geographical area the hole lies or the one who has declared an intention to fill it. In some such cases, the opportunity may be captured by another agent who happens to be nearby. But this will only occur if that nearby agent can override the normal filtering of the option to fill the hole in question.

To operationalize the override mechanism in the MA-Tileworld, we used a threshold technique similar to that used in our single-agent experiments. The threshold t is set by the experimenter. The value of the conflicting (hole-filling) option o is then computed as the maximal score that will be awarded if the hole is filled, divided by the Manhattan distance between the agent and the hole. We included overriding both in a static geographic strategy and an intention posting strategy; because the earlier experiment suggested that the overhead associated with the dynamic geographic strategy was too high, we did not include that in these experiments.

We further considered the question of whether the original agent—the one whose toes are being stepped on—should be notified of this fact. In the first two experimental conditions, agents are not notified when another agent takes over one of their goals. In the third experimental condition, *intention posting with preemption,* conflict is determined via intention posting, but, when an override occurs, the original agent is notified that its goal has been taken over, and so drops the goal and looks for an alternative.

In our earlier work, we determined that extreme boldness was a surprisingly good strategy in a wide variety of single-agent Tileworld environments. Overriding became beneficial in environments which can be characterized as presenting many opportunities that have relatively small payoff, and occasional critical opportunities, which have high payoff but short deadlines. Under those circumstances, it seems natural to think that extreme commitment to existing goals would not be a good strategy, because the high-payoff opportunities, if they are to be successfully acted on, require a quick response. We therefore constructed an environment that had those characteristics. In particular, it had two types of holes. "Common" (C-type) holes were quite numerous, had low scores and long lifetimes, and took a long time to fill. "Special" (S-type) holes were rare, had high scores and short lifetimes, and took a short time to fill. We studied similar environments for the multi-agent case.

For this experiment, we again used a 20x20 grid with four agents, and held the world speed constant at a baseline value. The number of C-type holes varied in the range of 30-80 with a score of 25-75, while the number of S-type holes varied in 0-4 with score range of 500-1500). The results are summarized in Figure 3, which shows the overall effectiveness (on the Y axis) as a function of the filtering threshold (on the X axis).

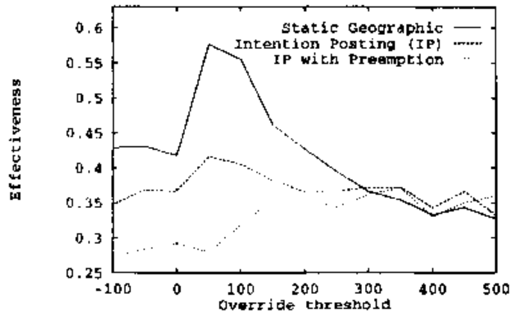

Figure 3: Overriding in the Multi-Agent Tileworld

As the graph shows, overriding is indeed beneficial, at least in the environmental parameters that we have considered. Moreover, the filtering threshold has a significant influence on the effectiveness of a specific filtering strategy. Although there are several local maximas, for each filtering strategy there is a unique threshold value where the maximal effectiveness is attained. In none of the experimental conditions was either extreme boldness or extreme caution the optimal strategy.

Perhaps more surprising was the relative performance of the various operationalizations of filtering in this environment. Recall that in the uniform environments studied in the first set of experiments, intention posting was always the best method of filtering. Here, the geographical-boundaries method is best, except when the override threshold is very high. This result led us to re-evaluate what is significant about the alternative filtering methods. What we realized is that what is particularly important in all the environments we studied is for the agents to maintain reasonable geographical separation from one another. In the uniform environments of the first set of experiments, all the filtering strategies lead to geographical separation. With geographical boundaries filtering, the agents focus on holes in distinct areas, and so tend to stay separated. With intention posting filtering, agents dynamically create separate regions because they tend first to form plans to fill nearby holes, and thus create territories that are avoided by other agents. Moreover these territories tend to stay fixed, because filtering is absolute.

However, once overriding is introduced, the stability of these local areas decreases. With a relatively low threshold, posted intentions will frequently be overridden by other agents, and there is nothing to prevent the agents from becoming clustered in one area of the grid and thus missing many remote opportunities. In contrast, although geographical-boundaries filtering will also be subject to frequent overrides, once any particular out-of-region goal is completed, the agent will return to its original territory.

Although this result in some sense is quite specific to the MA-Tileworld environment, it can be related to a much more general claim about the importance of resource distribution in coordination. What is interesting about this case is that the resources are both goals and objects needed to satisfy those goals.

6 The Effect of Defectors

The first two sets of experiments aimed at giving at least preliminary evidence that multi-agent filtering can be an effective means to achieve collaboration: it is an implicit strategy, requiring only that agents observe local rules of behavior, and it leads to overall improved performance, at least within the simulated environments we investigated. However, the fact that global performance is better if all the agents adopt a filtering strategy does not, in and of itself, guarantee that each agent will choose this strategy, if it has the choice. In some, and perhaps most, settings agents will have individual goals and utility functions: their concern is not with global performance, but with maximization of their private utility. We addressed these settings with experiments that abandon the assumption that agent will benevolently adopt the filtering coordinating mechanism. We assumed instead simply that agents are self-motivated (rational).

Game theory has addressed many interactions similar to the ones considered here. Such interactions have been analyzed to determine what an agent's chosen strategies would be, given the rules of the interaction. Our aim is complementary; it is to design rules that would induce the agents to adopt some specific strategy that we consider to be desirable. In our case we would like to make cooperation be the individually desired strategy. That is, each agent should prefer, out of "selfish" (rational) considerations, the filtering strategy over the other alternatives she might have. If all agents find cooperation to be their superior alternative it becomes an equilibrium point. In particular, a very strong claim would be that, regardless of whether the other agents are multi-agent filterers, each agent should itself choose to be one, i.e., multi-agent filtering is dominant to not filtering:

Definition 1 The strategy s^* is a dominant strategy if it is an agent 's strictly best response to any strategies that the other players might pick, in the sense that whatever strategies they pick, his payoff is highest with s_i .

A dominant strategy equilibrium is a strategy combination of each player's dominant strategy.

Thus, a strategy combination that is a dominant strategy equilibrium is very desirable. The fact that there is no importance to the other agents' behavior does away with the need to reason about the other agents' strategies, knowledge, or even computational capabilities. The behavior of an agent depends solely on its own characteristics.

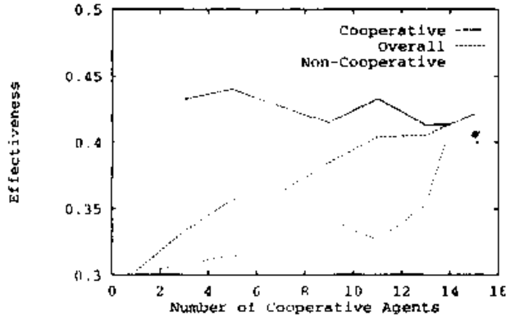

Figure 4: Stability of the Filtering Mechanism

To explore the question of whether the filtering strategy is dominant, we conducted another experiment, using MA-Tilewords populated by fifteen agents, only some of whom were cooperative. We varied the number of cooperative agents across trials, and measured the performance of the cooperative agents, the performance (average effectiveness) of the non-cooperative agents, and the global performance. We used the relatively weaker intention-posting method of filtering, and held world speed fixed at a baseline rate.[2] The experimental results are summarized in Figure 4. The X-axis indicates the number of filtering agents (out of the total population of fifteen), and the Y-axis shows the effectiveness.

As can be seen from the graph, the higher the percentage of filtering agents, the better the global performance is. But more importantly, the graph shows that at any given ratio of filtering to non-filtering agents, the filtering agents are doing better. That fact implies that regardless of the other agents' behavior, each agent should choose to cooperate and thus guarantee itself a higher utility. That is, at least for the range of environments that we have examined, intention-posting multi-agent filtering is in dominant strategy equilibrium.[3]

[2] Other experimental parameters were: Number of C-type holes=30-80, Score of C-Type holes=25-75, Number of S-type holes= 0-4, Score of C-Type holes=500-1500. Also, the grid was enlarged to 30x30, to accommodate the increased number of agents.

[3] Another phenomenon worth noting is that the effectiveness of the non-filtering agents improves as their proportion of the population decreases. This fact recalls the well-known parasite phenomenon of evolutionary game theory, and deserves further study.

7 Conclusions

In this paper, we have described an implicit approach to deriving coordination in multi-agent environments. The approach, the use of multi-agent filtering, is an outgrowth of earlier proposals that filtering is a useful technique for the control of reasoning by single agents. Multi-agent filtering is a natural extension to single-agent filtering, and has several highly desirable characteristics, such as simplicity and efficiency. To date we have conducted experiments, reported on in this paper, that provide preliminary evidence that multi-agent filtering is a good candidate for achieving coordination, and may even be a dominant strategy. Clearly the experiments conducted so far raise at least as many questions as they answer, and our experimental work, which aims to refine many of the hypotheses made in this paper, is ongoing.

The previous work that our approach is probably most closely related to is that on "social laws" [Moses and Tennenholtz, 1992; Shoham and Tennenholtz, 1992]. The idea behind social laws is that multi-agent systems may be designed so that, in any given situation, only a subset of the physically possible actions are designated as "legal". The specification of legality is intended to lead to cooperation: action restrictions should be defined so that, as long as the agents perform only legal actions, their behavior will be cooperative, i.e., their interactions will tend to lead to overall improved performance. If the restrictions on legal actions are principled, rather than ad hoc, then they can be described as social laws.

The strategy of multi-agent filtering that we describe in this paper can be cast as a particular class of social laws. However, there are some key differences between the approaches. Social laws are designed in part to guarantee that once an agent adopts a goal, no other agent will interfere. As a result, social laws are typically very difficult to generate and are very complex. In contrast, the filtering mechanism is intended to lead to improvement in the expected performance of the agents in a society, but not to guarantee success for each specific goal. The filtering strategy can be generated in a straightforward manner, based only on abstract properties of the environment and interaction. Moreover, while previous approaches assume that agents will follow the coordination strategy benevolently or through an explicit enforcing mechanisms, we expect the filtering mechanism to be self-enforcing, in which case, no particular assumption must be made about the agents' motivation, and a group of multi-agent filterer may either all be pursuing the same global goal or may have individual goals, some of which may conflict.

Acknowledgments

This work has been supported by the Air Force Office of Scientific Research (Contract F49620-92-J-0422), by the Rome Laboratory (RL) of the Air Force Material Command and the Defense Advanced Research Projects Agency (Contract F30602-93-C-0038), and by an NSF Young Investigator's Award (IRI-9258392).

References

[Bratman et al., 1988] Bratman, Michael E.; Israel, David J.; and Pollack, Martha E. 1988. Plans and resource-bounded practical reasoning. Computational Intelligence 4:349-355.

[Bratman, 1987] Bratman, Michael E. 1987. Intention, Plans and Practical Reason. Harvard University Press, Cambridge, MA.

[Ephrati and Rosenschein, 1991] Ephrati, E. and Rosenschein, J. S. 1991. The Clarke Tax as a consensus mechanism among automated agents. In Proceedings of the Ninth National Conference on Artificial Intelligence, Anaheim, California. 173-178.

[Ephrati and Rosenschein, 1993] Ephrati, E. and Rosenschein, J. S. 1993. Multi-agent planning as a dynamic search for social consensus. In Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, Chambery, France. 423-429.

[Goldman and Rosenschein, 1994] Goldman, Claudia V. and Rosenschein, Jeffrey S. 1994. Emergent coordination through the use of cooperative state-changing rules. In Proceedings of the Twelfth National Conference on Artificial Intelligence, Seattle, Washington. 408-413.

[Hanks et al., 1993] Hanks, Steve; Pollack, Martha E.; and Cohen, Paul R. 1993. Benchmarks, testbeds, controlled experimentation, and the design of agent architectures. AI Magazine 14(4): 17-42.

[Joslin et al., 1993] Joslin, David; Nunes, Arthur; and Pollack, Martha E. 1993. Tileworld users' manual. Technical Report 93-12, University of Pittsburgh Department of Computer Science.

[Kinny and Georgeff, 1991] Kinny, David N. and Georgeff, Michael P. 1991. Commitment and effectiveness of situated agents. In Proceedings of the Twelfth International Joint Conference on Artificial Intelligence, Sydney, Australia. 82-88.

[Kraus, 1993] Kraus, S. 1993. Agents contracting tasks in non-collaborative environments. In Proceedings of the Eleventh National Conference on Artificial Intelligence. 243-248.

[Moses and Tennenholtz, 1992] Moses, Y. and Tennenholtz, M. 1992. On computational aspects of artificial social systems. In Proceedings of the 11th International Workshop on Distributed Artificial Intelligence, Glen Arbor, Michigan. 267-283.

[Pollack et al, 1994] Pollack, Martha E.; Joslin, David; Nunes, Arthur; Ur, Sigalit; and Ephrati, Eithan 1994. Experimental investigation of an agent-commitment strategy. Technical Report 94-31, Univ. of Pittsburgh Dept. of Computer Science, Pittsburgh, PA.

[PoUack, 1992] Pollack, Martha E. 1992. The uses of plans. Artificial Intelligence 57:43-68.

[Shoham and Tennenholtz, 1992] Shoham, Y. and Tennenholtz, M. 1992. On the synthesis of usful social laws for artificial agent societies. In Proceedings of the Tenth National Conference on Artificial Intelligence, California. 276-281.

[Zlotkin and Rosenschein, 1993] Zlotkin, G. and Rosenschein, J. S. 1993. A domain theory for task oriented negotiation. In Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence, Chambery, France. 417-422.