A Metalevel Coordination Strategy for Reactive Cooperative Planning

Ei-Ichi Osawa
Sony Computer Science Laboratory Inc.
Takanawa Muse Bldg. 3F, 3-14-13 Higashi-Gotanda
Shinagawa-ku, Tokyo 141, Japan
osawa@csl.sony.co.jp

Abstract

In this paper, we propose a metalevel coordination strategy to implement an adaptive organization for reactive cooperative planning. The adaptive organization changes its organizational scheme adaptively as a means of coping with the dynamic problem spaces. Preliminary experiments shows that an adaptive organization can be made to the increase efficiency in dynamic problem spaces. The reason for this works is that reducing the degree of freedom in the problem space, while increasing the degree of interaction, demands greater coordination. However, if the number of effective local plans decrease, it would seem likely that if the agents were to have a better metalevel strategy, they would be better able to search this reduced space efficiently. The metalevel coordination incorporates an agent-wide metalevel heuristic function. In designing the metalevel coordination strategy, we take three aspects of reactive cooperative planning into account. These aspects include: the difference in the degree of achievement in successive turns; the certainty of shared information; and the degree of freedom of choice for agent's behavior. The adaptive organization works efficiently in cases where the communication cost is relatively expensive.

Introduction

In multiagent cooperative planning for shared global goals, important tasks include the generation of promising local plans at each agent and the coordination of these local plans. Specifically, where the dynamics of the problem space, e.g. the changing rate of goals compared with the performance of problem solvers, is relatively large, reactive planning that interleaves the plan generation and execution phases is known to be an effective methodology at least for a single agent (McDermott 1978; Agre & Chapman 1987; Maes 1991; Ishida & Korf 1991).

Schemes for reactive cooperative planning in dynamic problem spaces have been proposed and evaluated based on the pursuit (predator-prey) game (Benda, Jagannathan, & Dodhiawalla 1985; Stephens & Merx 1989; Gasser *et al.* 1989; Levy & Rosenschein 1992; Korf 1992). It is known that negotiation-based organizations are generally more efficient and more flexible than other rigid organizations, such as the autonomous agent organization, the communicating agent organization, and the controlling agent organization, in adapting to the random movement of the prey (Stephens & Merx 1989).

We expect that **adaptive organizations**, which dynamically change their organizational scheme, will perform efficiently when there is a trade-off between the accuracy of the global solutions and the efficiency of problem solving, according to changes in the problem space in the course of problem solving. For instance, Decker and Lesser (Decker & Lesser 1993) has shown that a dynamic reorganization method with a fast coordination algorithm is more effective than a static method in distributed sensor networks. Furthermore, we expect that the adaptive organization will also be effective in the reactive cooperative planning. To illustrate this expectation, let us consider the heuristic efficiency[1] (Stephens & Merx 1989) of the pursuit game. (Stephens & Merx 1989) reported that in many initial configurations of the problem space, the heuristic efficiency of the autonomous agent organizations is higher than that of the controlling agent organizations. However, the latter is the better of the two with respect to the capture result. These two facts suggest that certain adaptive organization that take both the accuracy of the global solution and the cost of problem solving into account will provide efficient problem solving. Many cooperative problem solving schemes have been proposed. Few, however, are adaptive in the above mentioned sense. Gasser et al (Gasser *et al.* 1989) proposed a generic cooperative problem solving scheme that captures the above mentioned aspect, but unfortunately, no performance results were provided.

The **metalevel coordination strategy** is a strategy with which agents determine the organizational

[1] The characteristic of this measure is that if an organization can rapidly capture the prey from the initial configuration, the heuristic efficiency of the pursuit becomes fairly large.

scheme (i.e. pursuing strategy in the pursuit game) to be followed. In this paper, we propose a simple metalevel coordination strategy that enable the realization the adaptive organizations.

The organization of this paper is as follows. In Section , we revisit the pursuit game, and summarize the results of our preliminary experiment, a performance analysis of an adaptive organization under some uncertainty. Section proposes an approach to the metalevel coordination strategy. In Section , we compare the costs of local plan generation and coordination, and show the theoretical upper bound of the speed of the prey which we can guarantee to capture. Also, we show the maximum time cost incurred for computing the metalevel coordination strategy where its operation is effective. Section gives the relationship between our scheme and other work in this area. Section presents our conclusions.

Preliminary Experiment – Adaptive Organization –

We take the pursuit game as a simple, but typical problem for reactive cooperative planning. In the following subsections, we first briefly revisit the game. Then, we summarize the results of a preliminary experiment of an adaptive organization. The metalevel coordination strategy discussed in this paper is based on the results obtained from this experiment.

Pursuit Game

The original version of the pursuit game was introduced by Benda et al (Benda, Jagannathan, & Dodhiawalla 1985). In the game there are two classes of agents, *red* and *blue*. There are four blue agents, and one red agent. They are allowed to move on a infinitely spread rectangular grid world (Figure 1). At each turn, each agent can move one cell horizontally or vertically, but not diagonally. Or, it may remain stationary. The goal of the four blue agents is to completely surround the red agent, by occupying the grid positions immediately north, south, east, and west of the red agent. Because of the characteristic of this goal, the four blue agents are sometimes called *predators*, and the red agent the *prey*.

In general, it is quite difficult to realize effective global control for problem solving in a distributed environment. Agents that act in a distributed environment may have limited abilities and restricted perception. Therefore, they may be able to obtain only partial information about their environment. When the agents face a problem that they cannot resolve alone, cooperation among the agents becomes significant. Research issues in cooperative problem solving involve methods of interaction, and structures of organization. The pursuit game was proposed to study the effects of organizational structures and communication on cooperative problem solving.

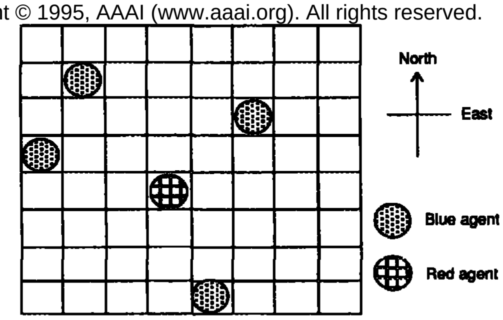

Figure 1: Pursuit game

The problem spaces of the pursuit game include locally-optimal states, because, for instance, the strategy by which the prey moves is assumed to be unknown. In other words, even though a predator selects, as its local goal, the capture position having lowest estimated cost, it is not guaranteed to be the globally optimal solution, because the prey might move into a different cell at the same turn. Also, if no organizational scheme is provided, the group of predators may not be able to achieve the capture state (the shared goal), because the combination of the optimal states for each agent may not form the capture state. They are two examples of uncertainty faced in the pursuit game. We also incorporate the notion of incomplete information by restricting the view and communication range[2] of each agent. This factor introduces another sort of uncertainty into the problem space.

Adaptive Organization and Its Performance

As we described in the previous section, we expect adaptive organizations will perform efficiently when there is a trade-off between the accuracy of global solutions and the efficiency of problem solving, according to changes in the problem space in the course of problem solving. To verify the assumption in reactive cooperative planning, we performed everal preliminary experiments, as explained below. We took the same approach as Stephens and Merx (Stephens & Merx 1989) regarding the design of the pursuit strategies for the autonomous-agent (*As*), communicating-agent (*Cs*), negotiating-agent (*Ns*), and controlling-agent (*CLs*) organizations.

Autonomous-agent Each predator estimates the cost of its reaching the capture position, which is the euclidean distance between the predator and the four cells surrounding the prey. Each predator then selects as its local goal the capture position for which the estimated cost is lowest. The estimation of the

[2]Given the location of a predator (a, b), when the range of the agent is restricted to r, the agent can detect only agents in the area surrounded by four points $(a-r, b-r)$, $(a-r, b+r)$, $(a+r, b+r)$, $(a+r, b-r)$.

distance is equivalent to the evaluation of heuristic function $h(n)$ of A^*. If the prey is not within the range of the predator, the predator explores its neighborhood in a breadth-first fashion.

Communicating-agent Each predator can communicate with the other predators concerning the current location of the prey.

Negotiating-agent Each predator estimates the cost of reaching the capture positions around the prey. The capture position and the estimated cost of reaching it are stored as a pair on the local goal blackboard. The negotiation procedure is started by transmitting the local goal blackboard of each agent to the other agents. The negotiation strategy is to award the first bid to the agent having the greatest disadvantage.

Controlling-agent One of four predators is selected as the controlling agent. The movement of the other three agents is controlled entirely by that agent. The controlling agent tries to realize the Lieb configuration[3].

Also, each agent can, at each turn, move one cell horizontally or vertically. The prey moves randomly.

Applying the above experimental settings, we have performed the following two kinds of experiments.

- Observation of the pursuit process at each turn, provided that incomplete information about the prey and other predators can be supported (since the range in which the predators can detect the prey is restricted, there is a possibility of this occurring).
- Relationship between total solution cost and ratio between unit communication cost and unit movement cost.

The following was observed.

Figure 2 shows that for organization *As*, if the range is restricted, convergence from the initial configuration to a near capture position becomes extremely inefficient[4]. The average euclidean distance from the predators to the prey in the experiments is approximately 4.0. When the range falls below that value, the efficiency of convergence to a near capture position rapidly decreases. Especially, convergence in the early phase of the pursuit is very bad. On the other hand, the *Cs* organization converges from the initial configuration to a near capture position fairly well, even though the range is restricted. Also, the number of communications needed to inform each predator of the current position of the prey is very small (an average of 4.0 communication units).

[3]Blocking of the prey and reaching the final state can be guaranteed once the Lieb configuration(Gasser *et al.* 1989) has been realized.

[4]Note that the enclosure distance on the vertical axis of the figures is the sum of the euclidean distance between each predator and the prey, normalized by the sum in the initial configuration.

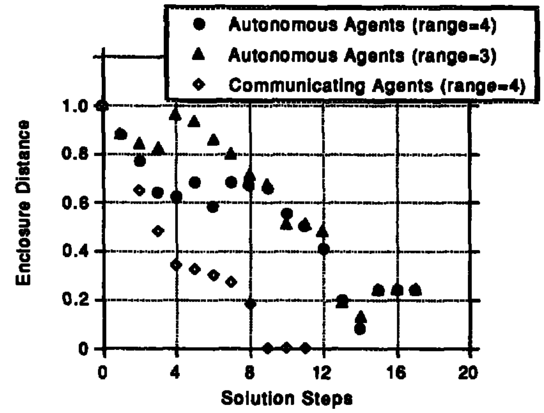

Figure 2: Effect of restricted range

By observing the changes in the enclosure distance of the *As* and *Cs* organizations in the near final state, we see that once a particular capture position has been occupied by several predators, the situation can be improved no further (a locally-optimal state). This is because the *As* and *Cs* organizations substantially have no effective pursuing strategy in the near final state. We also observed that, in those cases where the *As* and *Cs* organizations fail to capture the prey, the enclosure distance almost stabilizes. In such a situation, the predators cannot improve their pursuit states any further. Therefore, if the organization devises some functionality to recognize the situation, the predators can improve their pursuit process by adaptively changing the organizational scheme. This is the outline of **the metalevel coordination strategy**, discussed in the following sections. Figure 3 shows the pursuit processes of both an adaptive organization ($Cs \rightarrow CLs$ at turn 8) and *Ns* organization. Both organizations have the restricted ranges. This figure shows that if the above mentioned situation can be recognized, by adaptively changing the organizational control, the organization can successfully capture the prey.

Both the adaptive organization and *Ns* organization can successfully capture the prey. Also, both organizations exhibit similar convergence, from the initial configuration to a near capture position. Although not depicted in this figure, the performance *CLs* organization with restricted range is very bad. This is due to the inability of the predators to move effectively until the controlling agent has successfully found the prey.

In Figure 4, the horizontal axis indicates the ratio between unit communication cost and unit movement cost (the c/m ratio). The vertical axis indicates the total solution cost, the sum of the communication cost and movement cost. One unit cost of communication is the cost of one communication. In the *Ns* organization each predator sends its preference list to the other three predators. Therefore, it incurs a communication cost

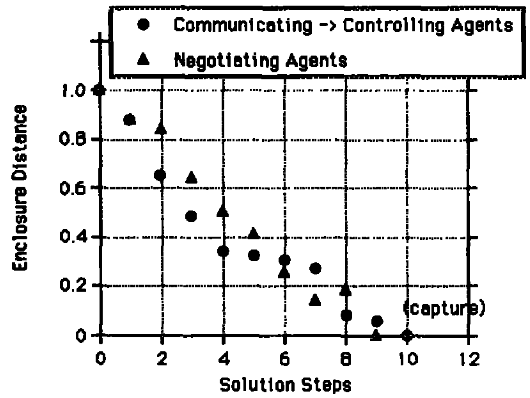

Figure 3: Effect of restricted range

of three. One unit cost of movement is the cost of one movement.

This figure shows that, in those cases where the unit communication cost is negligible in comparison with the unit move cost, the total solution cost of the *Ns* and adaptive organizations are small. However, if the c/m ratio reaches approximately 0.1, the total solution cost for the *Ns* and adaptive organizations become equivalent. As the c/m ratio contributes to increase, the total solution cost of the *Ns* organization rapidly becomes large.

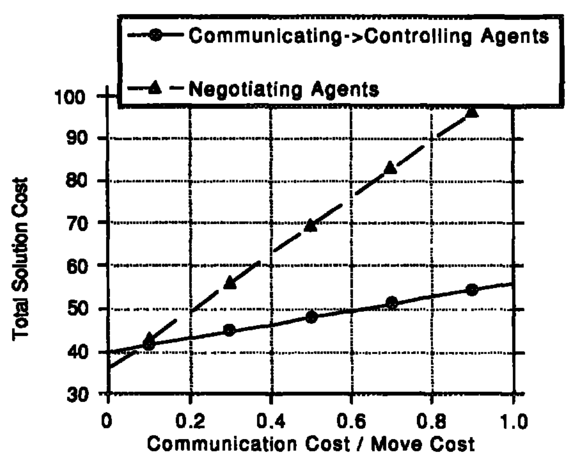

Figure 4: c/m ratio vs. total solution cost (range = 4)

In the early phase of cooperative problem solving, even though each agent behaves almost autonomously with data communication (*Cs* organization), its operation is fairly efficient since there is a large degree of freedom of choice in this phase. Also, communicating agents do not incur extensive communication like negotiating agents. However, when all agents approach the final state (a shared global goal), freedom of choice decreases, meaning that the number of possible effective local plans also decreases. Therefore, in near-final states it is much more efficient to commit the behavior of all agents to a single control scheme.

From the above, we can note the following: 1) an adaptive organization can be made to increase their efficiency, and 2) the reason for this is because reducing the degree of freedom in the problem space while increasing the degree of interaction demands greater coordination. However, if the number of effective local plans decreases, (thus decreasing the population of goal states in the agent-wide planning space) it would seem likely that, if the agents had a better metalevel strategy, they would be better able to search this reduced space more efficiently.

Metalevel Coordination Strategy

Based on the observations explained in the previous section, we can list those factors that are important in designing the metalevel coordination strategy.

Difference in achievement degree When an organization fails to capture the prey in the near capture position, the sum of the euclidean distance becomes nearly stable around some value. This means that, in such a situation, the predators cannot improve their pursuit performance any further.

Certainty When all agents approach the final state (a shared global goal), it is often possible for them to share sufficient information to achieve that shared goal. Once such information is shared (for the pursuit game, if the location of the other predators and the prey has been known), committing agent behavior to a efficient organizational scheme, such as single control scheme, would appear to be effective.

Freedom If the organization is not able to determine an effective global plan for achieving the shared goal, it is important to give each agent a large degree of freedom of choice in its behavior. This is because doing so could enable a greater opportunity to explore better possibilities. On the other hand, when all agents share sufficient information to achieve a shared global goal, it will be possible to decrease the overhead for coordination by decreasing freedom of choice, i.e. the number of possible possible local plans.

Below, we will introduce a metalevel coordination strategy for the pursuit game. Several terms and some aspects of the following strategy are specific to the game because of the purpose of this paper (i.e. they hopefully make this paper easier to understand). However, slight modification of the strategy produces a generic version of the metalevel coordination strategy for reactive cooperative planning.

First, we must introduce several predicates and variables. Let $\mathcal{E}$ denote a proposition that is true if a shared global goal is achieved (for the pursuit game, if the four predators completely surround the prey, the proposition is true). Let $\mathcal{V}$ denote a unary predicate over agents (i.e. the predators or the prey) that is true if and only if its argument is visible (i.e. within the sight range). Let B_1, B_2, B_3, B_4, and R denote each predator, and the prey, successively. We assume a symmetrical visibility, i.e. if $\mathcal{V}(B_i)$ is true for B_j, then $\mathcal{V}(B_j)$ is also true for B_i. Let $d_i(B_j)$ denote the distance between predator B_j and the prey R at turn i. The capture distance at turn i is denoted as D_i, and is defined as $\sum_{j=1,\cdots,4} d_i(B_j)$. The difference between the previous capture distance and the current capture distance ($\Delta_i = D_{i-1} - D_i$) is called the capture distance difference. S_k denotes an organization scheme, while $[S_0, S_1, \ldots, S_n]$ denotes a list of organization schemes. Note that if j is greater than i, then the freedom of agents in scheme S_i is greater than that of S_j.

[Metalevel coordination strategy] (The strategy described below is that of predator B_1. The strategy of all other predators is similar.)

1. Let turn $i \leftarrow 0$, counter $C \leftarrow 0$, $D_0 = \infty$, $\Delta_0 = \delta(\geq 0)$, and initial index of organization scheme $j \leftarrow 0$.
2. If $\mathcal{E}$ is true, then halt.
3. If $\mathcal{V}(R) \wedge \mathcal{V}(B_2) \wedge \mathcal{V}(B_3) \wedge \mathcal{V}(B_4)$ is false, then execute S_j, let D_i be ∞, and $i \leftarrow i+1$, and go to step 2.
4. If $\Delta_i > \delta$, then execute S_j, let $i \leftarrow i+1$, and $C \leftarrow 0$, and go to step 2.
5. Let $C \leftarrow C+1$. If $C < \tau$, then execute S_j, let $i \leftarrow i+1$, and goto step 2.
6. While $j < n$, let $j \leftarrow j+1$. Execute S_j, let $i \leftarrow i+1$, and $C \leftarrow 0$, and go to step 2.

Next, we we provide a brief explanation of constants δ and τ, introduced in the above metalevel strategy. In the metalevel coordination strategy, if the capture distance difference of the current turn Δ_i is smaller than δ, it is assumed that the current organization scheme may not be effective. If such a state is recognized in successive τ turns, the next less free (i.e. more coordinated) organization scheme is chosen. Therefore, Δ_i, δ, and τ can be regarded as constituents of an agent-wide metalevel heuristic function.

Since the symmetrical visibility is assumed for all predators, the value of Δ_i is common to all predators. Therefore, it is guaranteed that all predators synchronously, i.e. at the same turn, select the next organizational scheme without metalevel communications.

The adaptive organization scheme discussed in the previous section can be realized by letting S_0 be Cs organization, S_1 be CLs organization, δ be 0, and τ be some small integer (e.g. 3).

Maximum Changing Rate of Achievable Shared Goals

In Section , the advantage of the adaptive organization w.r.t. efficiency relative to the negotiating organization was derived from the total solution cost versus the ratio between unit communication cost and unit movement cost. In our experiment, we assumed that all agents can, at each turn, decide their next destination cell and move into that cell. In the pursue game, that four predators can successfully capture the prey can not be guaranteed unless the speed of each predator is faster than the average speed (changing rate) of the prey (the proof of this theorem is analogous to the proof of the completeness of the moving target search (Ishida & Korf 1991)). The faster all predators can move, the faster prey they can capture. Below, we will assume that the speed of predators depends mostly on the time cost being consumed to decide the next destination cell. In the following subsections, we first estimate the maximum speed of the prey that can be captured by each basic organization scheme. Then, we discuss the time cost incurred for computing the metalevel coordination strategy.

Performance of Basic Organization Schemes

The time cost incurred to decide the next destination cell is primarily the sum of the time cost for communication between predators and that for reasoning which capture position has the lowest estimated cost. Let p be the time cost incurred to autonomously, i.e. without communication, select a local goal – the next destination cell – out of four possible moves based on the capture position having the lowest estimated cost. Also, let c be the time cost of one unit communication for negotiation, communication, and control. For simplicity, we assume that the unit communication cost doesn't depend on the distance between agents. Furthermore, let q be the time cost incurred in deciding the next destination cell out of a predator's preference list in the negotiating organization (details of this process are given in Section).

Using these costs, the maximum time cost incurred to decide the next cell at each turn for each basic organization scheme can be calculated as follows (hereafter, it is simply called the maximum time cost).

Table 1: **Maximum time cost of each agent at each turn**

Organization	Maximum time cost
autonomous-agent	$C_{auto} = p$
Communicating-agent	$C_{comm} = p + c$
Negotiating-agent	$C_{nego} = p + 3c + q$
Controlling-agent	$C_{ctrl} = 4p + 3c$

The reciprocals of the costs in the above table are the maximum speed of the prey for which capture is possible, provided that reasoning, communication, and other processing are sequentially performed at each agent.

In the preliminary experiments, we observed that q is approximately equal to $2p$. Therefore, the maximum time cost of the organizations is, in ascending order, *As*, *Cs*, *Ns*, and *CLs*.

Effectiveness of Adaptive Organization

Table 1 shows that the maximum time cost incurred by organization *Cs* is smaller than that of either organization *CLs* or *Ns*. Therefore, it can be seen that using organization *Cs* will be effective, provided that its performance from the initial configuration to a near capture position is relatively good.

As we have already seen in Section , organizations *CLs* and *Ns* are promising in the near final state. Stephens and Merx (Stephens & Merx 1989) reported that organization *Ns* is generally more efficient and more flexible than other organizations in adapting to the random movement of the prey. However, as we mentioned in Section , the total solution cost incurred by organization *Ns* is significantly greater than that of the adaptive organization where the communication cost is relatively expensive. Additionally, if the time cost incurred to compute the metalevel coordination strategy – it is the overhead of the adaptive organization – is smaller than $C_{ctrl} - C_{comm}(= 3p + 2c)$, the adaptive organization can capture all the prey that can be captured by organization *CLs*.

Relation to Other Work

Benda et al (Benda, Jagannathan, & Dodhiawalla 1985) evaluated the average number of moves and communications required to capture the prey for nine four-predator-agent organizations that consist of basic three links (data communication, negotiation, and control). They concluded that organizations that feature a hierarchical control structure are effective with respect to the communication cost, since in these organizations unnecessary communication can be sufficiently suppressed. However, they didn't deal with incomplete information, however. According to our preliminary experiment, the single control organization with restricted range shows very bad performance while predators are far from the prey.

Stephens and Merx (Stephens & Merx 1989) evaluated four homogeneous organizations – autonomous-agent, communicating-agent, negotiating-agent, and controlling-agent organizations – against six particular initial configuration. Their experimental results showed:

- a controlling-agent organization is the only organization that is always successful in capturing the prey,
- a negotiating-agent organization is generally more efficient and more flexible than other organizations in adapting to the random movement of the prey.

The second result is somewhat inconsistent with Benda's result. This is because Stephens evaluated organizations according to the heuristic efficiency. Their experiments determined that a negotiating-agent organization is in average more flexible than a controlling agents organization. The adaptive organization proposed in this paper extends the use of those basic organization schemes, and illustrating the effectiveness of the adaptive organization when handling incomplete information . The negotiating-agent organization investigated by Stephens and Merx can be classified as as an adaptive organization. However, since such an organization requires extensive communication, the adaptive organization proposed in this paper is much more efficient where the communication cost is relatively expensive.

The previous two approaches adopted a fixed cooperation mechanism. Gasser et al (Gasser *et al.* 1989) studied cooperation schemes through the pursuit game, introducing a generic scheme. Their scheme has six phases of cooperative problem solving. It is an adaptive organizational scheme, however, no performance analysis of the scheme were given. Our approach to adaptive organizations is closely related to their idea. We incorporated the metalevel coordination strategy with an agent-wide metalevel heuristic function. Also, We demonstrated effectiveness of the idea of dynamically changing an organization through experiment, and estimated those cases where the adaptive organization works effectively.

Levy and Rosenschein (Levy & Rosenschein 1992) apply a game-theory approach to the pursuit game. Contrary to previous approaches, they avoid explicit cooperation, and focus instead on the utility functions of rational predators to produce the necessary behavior. Their approach is relatively complex and deliberation-oriented, and seems to require intensive computation. However, no cost analysis has yet been reported. We regarded the pursuit game as being a type of reactive cooperative planning, and have analyzed adaptive organizations from the viewpoint of time cost and efficiency.

Decker and Lesser (Decker & Lesser 1993) proposed a one-shot dynamic coordination algorithm. It reorganizes the areas of responsibility for a set of distributed sensor network agents. While their domain of application is different from our domain, they showed that a dynamic reorganization based on a fast coordination algorithm can often outperform a static method in distributed in a particular environment. Their algorithm uses one metalevel communication action, while in our algorithm all agents are able to synchronously change the organizational scheme without metalevel communication.

Conclusion and Future Work

We have presented our initial experimental results for an adaptive organization aimed at reactive cooperative planning. The adaptive organization changes its organizational schemes adaptively to cope with dynamic problem spaces. Also, we presented a metalevel coordination strategy, based on which cooperating agents change their organization.

The initial results of our experiments can be summarized as follows. An adaptive organization can be made to increase efficiency. The reason for this works is because reducing the degree of freedom in the problem space while increasing the degree of interaction demands greater coordination. However, if the number of effective local plans decreases, (thus decreasing the population of goal states in the agent-wide planning space) it would seem likely that if the agents were to have a better metalevel strategy, they would be better able to search this reduced space more efficiently.

The metalevel coordination strategy reflects three significant aspects of the preliminary experiments. These includes: 1) the difference in the degree of achievement in successive turns, 2) the certainty of shared information among agents, and 3) the degree of freedom of choice for the agent's behavior. Primarily, the metalevel coordination strategy was designed to recognize those situations where the current organization scheme is not effective. Once such a situation has been recognized, a more coordinated organization scheme can be chosen. Therefore, the metalevel coordination strategy can be seen as to include an agent-wide metalevel heuristic function.

It is well known that some previously proposed organizations, such as the negotiating agents organization and the game-theoretic organization, is generally flexible in adapting to the random movement of the goal. In this sense, it can be characterized as an adaptive organization. However, since such organizations requires extensive communication or intensive computation, the adaptive organization proposed in this paper works much more efficiently where the communication cost is relatively expensive.

We are currently working on the following extension. The proposed metalevel coordination strategy has been evaluated in a dynamic domain with two kinds of uncertainties, i.e. an unknown variable characteristic of the shared goal state, and incomplete information because of the restricted range. We are now extending the metalevel coordination strategy, such that it can cope with more complex spaces, e.g. spaces that contain many unpredictable obstacles. In this case, the design of agent-wide metalevel heuristic function needs a consideration for situations where each predator might frequently falls into locally-optimal states (e.g. heuristic depressions).

References

Agre, P. E., and Chapman, D. 1987. Pengi: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, 268–272.

Benda, M.; Jagannathan, V.; and Dodhiawalla, R. 1985. On Optimal Cooperation of Knowledge Sources. Technical Report BCS-G2010-28, Boeing AI Center.

Decker, K., and Lesser, V. 1993. A One-shot Dynamic Coordination Algorithm for Distributed Sensor Networks. In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, 210–216.

Gasser, L.; Rouquette, N.; Hill, R. W.; and Lieb, J. 1989. Representing and Using Organizational Knowledge in Distributed AI Systems. In Gasser, L., and Huhns, M. N., eds., *Distributed Artificial Intelligence, Volume II*. Morgan Kaufmann Publishers, Inc. 55–78.

Ishida, T., and Korf, R. E. 1991. Moving Target Search. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, 204–210.

Korf, R. E. 1992. A Simple Solution to Pursuit Games. In *Proceedings of the Eleventh International Workshop on Distributed Artificial Intelligence*.

Levy, R., and Rosenschein, J. S. 1992. A Game Theoretic Approach to Distributed Artificial Intelligence. In Werner, E., and Demazeau, Y., eds., *Decentralized A.I. 3*. Elsevier/North Holland.

Maes, P. 1991. Situated agents can have goals. In Maes, P., ed., *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*. The MIT Press. 49–70.

McDermott, D. 1978. Planning and Action. *Cognitive Science* 2:71–110.

Stephens, L. M., and Merx, M. 1989. Agent Organization as an Effector of DAI System Performance. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, 263–292.