Designing Control Laws for Cooperative Agent Teams

Lynne E. Parker
MIT Artificial Intelligence Laboratory
545 Technology Square, Room 709
Cambridge, MA 02139
parkerl@ai.mit.edu

Abstract

The design of the control laws governing the behavior of individual agents is crucial for the successful development of cooperative agent teams. These control laws utilize a combination of local and/or global knowledge to achieve the resulting group behavior. A key difficulty in this development is determining the appropriate balance between the use of global information and the use of local information to achieve coherent cooperation without excessive communication requirements. This paper addresses this issue by presenting some general guidelines and principles for determining the appropriate level of global versus local control. These principles are illustrated and implemented in a "keep formation" cooperative task case study, which presents several alternative control strategies along the local versus global spectrum. In this case study, we present experimental data that demonstrate that local control alone is not sufficient to meet the goals of certain tasks, and that an increasing use of global knowledge can result in a steadily improving group cooperation. We conclude that the use of local control information to ground global knowledge in the current situation is perhaps the best way to achieve the proper balance between local and global control.

1 Introduction

The design of the control laws governing the behavior of individual agents is of overriding importance in the successful development of teams of cooperating, situated, autonomous agents. These control laws determine not only how each agent behaves in its own local situation, but also how the group as a whole behaves in its environment. If the agents are truly autonomous, and thus decide on their own actions independent of any centralized control, we say that the group behavior *emerges* as a side-effect of the interactions of the individual agents in the world. The question then becomes how to design the individual control laws to achieve the desired global group behavior.

A popular approach to the design of these control laws is to give each agent the ability to react solely to its own local environment, consisting of those proximate aspects of its world the agent can sense [1, 4, 6, 10]. In this approach, knowledge about the global goals of the group as a whole is not available to the agents. The hope is that local knowledge alone will be sufficient to form a cohesive, cooperative group. Indeed, in the referenced papers and elsewhere, such control laws are shown to yield intriguing group behaviors in a variety of applications.

Another, quite different, design approach is to provide the agents with knowledge about the group's global goals. The agents then use the global goals, perhaps in combination with additional global information, to select actions that are more consistent with the overall group intentions, thus yielding a more cooperative team. This approach, too, has been successful in certain applications at achieving the desired global cooperative behavior, and is usually typified by the use of communication between agents to convey partial or complete global information among agents [3, 5, 7].

Both of these approaches offer certain advantages: (1) local control laws are appealing because of their simplicity and potential to generate globally emergent functionality, whereas (2) global control laws generally allow more coherent cooperation. On the other hand, (1) it may often be unclear how (and whether it is even possible) to design local control laws to achieve the required group behavior, and (2), the use of global knowledge is usually paralleled with increased inter-agent communication. Thus, the designer of a cooperative system must determine the appropriate balance between the use of global information and the use of local information to achieve coherent cooperation without excessive communication requirements. How does one determine the proper mix? This paper seeks to answer this question by addressing the issue of local versus global control in cooperative systems. In section 2, we look more closely at global and local control, while in section 3 we discuss principles for determining the proper balance between local and global control. Section 4 presents in detail the "Keeping Formation" case study which stimulated our thoughts on the local versus global control issues, discussing the design and implementation of several alternative control strategies, and our results. The final section offers concluding comments and a summary of the general principles and guidelines put forth by this paper.

2 Descriptions of global and local control

In practice a continuum exists between strictly global and strictly local control laws. Thus, the control laws guiding an agent will probably use a mixture of local and global knowledge, rather than adhering strictly to one type alone. To simplify the discussion, however, these types are considered separately in this section, which compares and contrasts these two types of control.

2.1 Global control

Global control laws utilize the global goals of the cooperative team and/or global knowledge about the team's current or upcoming actions to direct an individual agent's actions. With these laws, an agent is able to influence its own actions toward team-level goals that cannot be sensed in its own local world. To better understand the implications of the use of global control laws, let us look individually at the two types of information utilized by these laws: *global goals* and *global knowledge*. The *global goals* of a team indicate the overall mission that the team is required to accomplish. These goals are typically imposed upon the team by a centralized controller, such as a human or another autonomous agent. Often this controller is an agent from outside the cooperative team rather than from within, although it is not uncommon to have a leading agent within the team specifying these goals.

Of particular impact on the design of cooperative teams is the time at which the global goals become known[8]. If the goals are known and fixed at design-time, then it may be possible to incorporate these goals implicitly into the control laws of each agent. Whether this can be done depends on the proper match between the sensing capabilities of the agents and the sensing requirements of the global goals. If all the information required for an agent to act consistently with the global goals can be sensed locally by that agent at run-time, then the global goals can be designed into the agent. On the other hand, if the goals are not fixed or known at design-time, then they obviously cannot be designed into the agents. In this case, the agents must possess the capability to obtain and appropriately act upon the goals provided at run-time.

The second type of information used by global control laws, *global knowledge*, refers to the additional information that may be necessary for the cooperative team to achieve the global goals. This information typically indicates what other agents in the team are doing or are going to do, or what the environment looks like in relation to the current cooperative task. By definition, all such information is normally not available to the individual agents through their sensors (other than their communication channels); if it were, then we would consider it to be local information.

How does an agent obtain this global knowledge? Several methods are possible. Perhaps the most obvious manner is for a centralized informant (either a human or an autonomous agent outside the group, or an individual agent within the group) to explicitly communicate the information directly to the team as it becomes available. The agents can then utilize this explicitly communicated information as advice, along with locally sensed data, to undertake appropriate actions which are consistent with the global goals. A second method of obtaining global knowledge, albeit in an approximate form, is for agents to interpret the actions of another agent through the use of a model of that agent's behavior. The behavioral model can be used not only to interpret an agent's current actions, but also to predict that agent's future actions. In a sense, this method utilizes implicit communication, since the observing agent receives information from the actions of another agent. Note that the behavioral model does not need to be explicitly accessible to the modeling agent. Rather, it could be learned or programmed implicitly such that certain actions by the modeled agent trigger the appropriate responses in the modeling agent.

The use of global goals and information is not without its shortcomings, however. Adequate global information may not be available to achieve the desired global goal. Even with global knowledge, an agent may still not exhibit optimal global behavior unless it utilizes all of the global knowledge available. Processing this global information requires time and resources, both of which are usually limited in real-world applications. If the global goals or information is changing often enough, the agent may not be able to act upon the global knowledge before it becomes out-of-date. Indeed, in some situations, global control of any kind will be impossible, thus mandating the use of local control.

2.2 Local control

Local control laws, on the other hand, guide an agent's actions based on the proximate environment of that agent. Such information is derived from the agent's sensory capabilities, and thus reflects the state of the world near the agent. Local control laws allow agents to react to dynamic changes in their environment without relying on preconceived plans or expectations of the world. With a careful design, global functionality can emerge from the interaction of the local control laws of the individual agents. For example, Franklin and Harmon [4] have shown that a global cooperative hunting behavior emerges from the use of three local cooperative control laws: cooperative pursuit, triangulation, and encirclement. These control laws are appealing because of their simplicity and power to generate globally emergent functionality.

However, local control laws also have their limitations — certain global goals cannot be attained through the use of local control laws alone. In some cases, it may be possible to utilize local control laws to achieve an approximation to the optimal results, which may be totally acceptable for many applications. However, since local control relies strictly on features of the environment that can be sensed, those aspects of global goals that have no physical manifestation in the world cannot be acted upon by local control laws.

3 The proper balance

Selecting the proper balance between the use of local and global control laws is not an easy task, and varies from application to application. Of central importance is determining the acceptable level of cooperation and performance of the autonomous agent team in a particular application. Some applications may be considered successfully accomplished if the team finishes the task at all, regardless of how they do it or how long it takes. Several questions arise when considering the design of cooperative control laws. What are the tradeoffs between global versus local control? Will global and local information conflict, and, if so, how does one arbitrate between them? These issues and others are discussed in the following sections.

3.1 Tradeoffs between global and local control

Assuming the availability of global goals and/or global knowledge which can be used by the cooperative team, the designer must decide whether to incorporate the use of this global information into the team, or to approach the problem with more local control. In doing this, the designer must weigh the costs of using global information with those of doing without. Several questions must be addressed. First, how static is the global knowledge? The knowledge could be known and fixed at the start of the task, thus making it an excellent candidate for use in a global control law. In general, the more static the global knowledge is, the more practical its use by a global control law.

An additional issue concerns how difficult it is to approximate global knowledge by comparing observations of an agent's actions with a model of that agent's behavior. This type of approximation can be quite challenging, depending upon the complexity of the autonomous agents and the environment. When possible, behavioral observation is more robust and dynamic than the use of global knowledge that may change unexpectedly. As global knowledge becomes more unreliable, an agent team must increase its dependence on behavioral observation and interpretation. Good results with behavior observation and interpretation should be expected particularly for teams of agents possessing a fixed set of discernible actions. One of the primary difficulties with behavior observation, however, lies in the limited ability of agents to sense the current actions of other agents. In cases where the sensing capabilities are not sufficiently extensive, the team can utilize communication to inform other agents of their current actions.

Other issues that must be addressed include: How badly will the performance degrade without the use of global knowledge? How difficult is it to use global knowledge? How costly is it to violate the global goals? How accessible is the global knowledge? How much local information can be sensed? Answers to these questions must be application-dependent, and considered in light of the capabilities of the specific agents to be used, the environment they will be operating in, and the scope of the application. In general, the more unknown the global information is, the more dependence a team must have on local control, perhaps combined with approximations to global knowledge based on behavioral and environmental observation and interpretation.

3.2 Conflicts between global and local control information

A combination of local and global control in the same agent may lead to conflicts if the control laws are designed to compete with one another by having the global control laws utilize strictly global information, while the local control laws utilize strictly local information. A better way to design the system is to view the global information as providing general guidance for the longer-term actions of an agent, whereas the local information indicates the more short-term, reactive actions the agent should take within the scope of the longer-term goals. This can often be achieved by combining the use of local and global information into a composite control law that more intelligently interprets the local information in the context of the global knowledge.

Problems may still arise if an agent using global knowledge is also trying to react appropriately to an agent that is not using global knowledge. In this case, the designer must provide the agents with the ability to arbitrate between certain aspects of global or local information when the need arises. Perhaps the best way to achieve the interaction of the two types of knowledge is by using local control information to ground global knowledge in the current situation. In this manner, the agents are able to remain focused on the overall goal of their group while reacting to the dynamics of their present contexts.

4 Experimental results

We have implemented and evaluated several control strategies along the local versus global spectrum by performing a wide range of experiments in simulation. For each of the control strategies, we measured the results quantitatively by collecting data on the mission completion time and amount of agent error in performing the mission. The section describes these results, first defining the mission performed by the agents, and then discussing the results of experiments with four control strategies that vary in the amount of global and local information.

4.1 Task description

The "Keep Formation" task requires a group of agents to stay in formation with one another (i.e. remain aligned

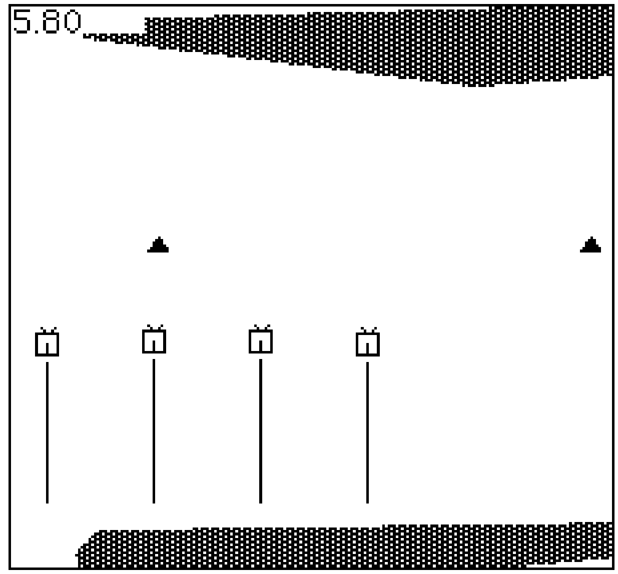

Figure 1: Four agents keeping formation side by side.

side to side) while the leader of the group follows a prespecified route and while all agents avoid obstacles as they appear (see figure 1). Each of these agents has the ability to sense the location of its neighboring agents relative to itself (local knowledge) and is physically constrained by the inability to move backwards.

The global goal of this task is twofold: first, the agents should reach their destination as quickly as possible without increasing their maximum speed, and, second, they must maintain the specified formation in a manner that appears to a casual human observer to be human-driven, meaning that the agents should not allow huge or "unnatural" (an admittedly subjective measure) perturbations in the desired group formation[1]. This subjective measure is quantified by defining the notion of *normalized cumulative formation error*, which is calculated as follows: at a given time t, the formation error, fe, is given by

$$fe_t = \sum_{i \neq leader} d_i$$

where d_i is the distance between the current position of agent i and the proper formation position of agent i, based on the leader's current location. The cumulative formation error, cum_fe is then given by:

$$cum\_fe = \sum_{t=0}^{t_{max}} fe_t$$

for integral values of t, meaning that the formation error is sampled and accumulated at discrete points in time up to t_{max}, which is the mission completion time. Since this cumulative formation error is dependent on the total time of mission completion, it is divided by the total mission

[1]Of course, we are not requiring that the Turing test be passed by these agents. The point is not to fool humans, but to display human-like strategies toward staying in formation.

time to result in the normalized cumulative formation error, which is used as a basis of comparison between the control strategies.

The agents in this mission, designed using layers of behaviors [2, 9], are provided with competences to avoid obstacles and to follow a specified route. The following experiments vary the design of the KEEP_FORMATION behavior to determine the level of performance we are able to achieve with different levels of local versus global control.

4.2 Implementation

Two simulation systems were utilized for these experiments. The first (developed for a more extensive simulation training program) consisted of a Sun-4 workstation running the main cooperative agent code (which was written primarily in C), connected to a Symbolics machine for experiment creation, and to a VP-1000 system for graphical display. The second simulation system, written in Common Lisp, ran entirely on a Macintosh II[2]. The experiments varied in the route the agents were instructed to follow, the character of the route (i.e., sharp versus smooth turns, following a road or traveling through open terrain, etc.), the number of agents in the team, the formation the agents were to maintain, and the presence of static or dynamic obstacles in the paths of the agents. Typical experiments involved from 1 to 14 agents instructed to follow a specified route while staying in a side-by-side formation. Often, an additional team of agents simultaneously performed a similar task along an intersecting route, requiring the agents in both teams to avoid dynamic obstacles (other agents) as they maintained their formation.

Each of the control strategies described below was implemented and tested separately to determine the group behavior that resulted from each of the strategies. These strategies were evaluated based on the quantitative measures described earlier (i.e. mission completion time and normalized cumulative formation error). To collect this data, each experiment for each control strategy was run ten times. Figure 2 plots the results, which are discussed in the next section[3].

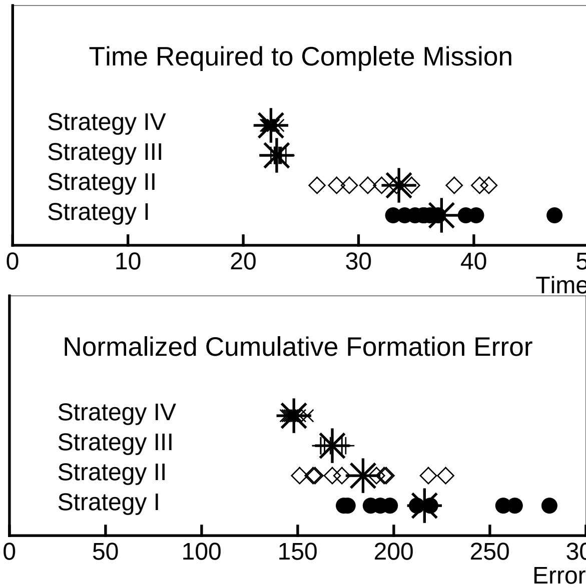

Figure 2: Time and error results of 10 runs of each control strategy. (Asterisks denote the mean values.)

4.3 Control strategies

4.3.1 Strategy I: Using local control alone

At first glance, it appears that KEEP_FORMATION could be achieved using local control laws alone. Each agent could be assigned a leader and then use a simple control law that directs it toward a prespecified offset and direction from its leading agent. As the group leader moves forward along the path (which is known only to the group leader), the other agents follow along to stay in formation. Indeed, in experiments involving relatively few agents traversing smooth routes in the absence of obstacles, we found that this law would perform adequately well. However, a problem arises if the group leader makes a sharp turn along the path, as illustrated in figure 3[4]. In this snapshot of the simulation, agent B is the overall leader, agents A and C are following agent B, and agent D is following agent C. In following its leader, agent A seeks to always locate itself a preset distance to the left of B, while agents C and D strive to be located the same distance to the right of their respective leaders. In this figure, the group leader, B, is making a right-hand turn. Since the followers are using strictly local information in this case, they continue to follow the same rules as before, maintaining a specified distance and offset from its leader. Agent A performs satisfactorily, aiming toward the location the appropriate distance to the left of B. Since these agents cannot back up, agent C turns around and aims toward a location to the right of B. Now, however, we have a problem with agent D. It aims as usual toward the right of C, but this position is out of formation with the rest of the group. Here we see that local control information is not sufficient to achieve the desired global goals. Figure 2 shows that this strategy resulted in the worst quantitative performance of all the control strategies studied.

4.3.2 Strategy II: Using local control augmented by a global goal

An improvement on the situation provides the agents with knowledge of the global goal of the group. Now, since the agents are "aware" that they should achieve a global linear formation, they select their positions after agent B's right-hand turn based on the global formation, while still remaining responsive to the local dynamics of the agents adjacent to them. With this information, agents A and C aim toward the same positions as in the previous case, but agent D now heads toward a more globally appropriate location, as shown in figure 4. Unfortunately, these movements could still be inappropriate if the leader is just avoiding an obstacle, rather than making a turn along the path. In spite of this, it is clear that knowledge and use of the global goal can yield improved group coordination. In figure 4 we see that this strategy resulted in an average 10% reduction in mission completion time and an average

[2] The simulation data and snapshots in this paper were collected from the Macintosh version.

[3] The variation in results for control strategies I and II is due to unpredictable interference among agents when they stray significantly out of formation.

[4] In figures 3– 6, the bold arrows, when present, indicate the intended direction of travel of the agents, the thin lines show the paths already traversed by the agents, and the leader's path goes from its starting location to the small triangle directly in front of it, and then to the small triangle on the right.

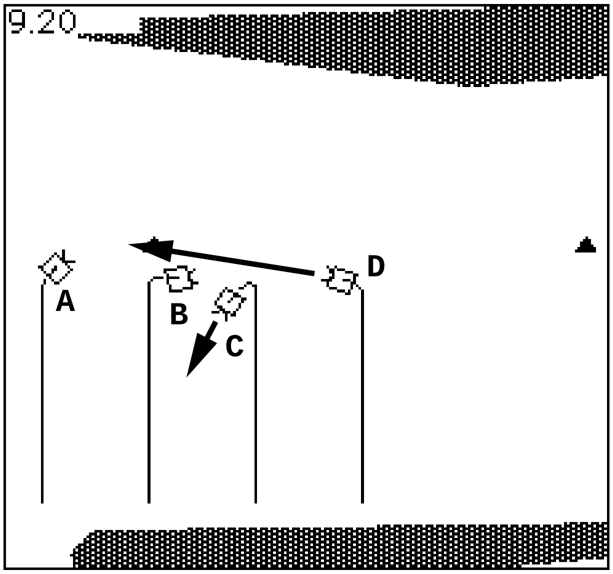

Figure 3: Team behavior using Strategy I.

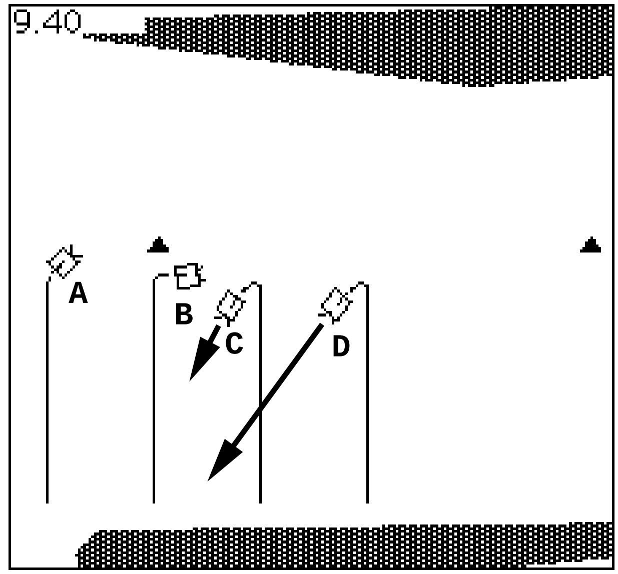

Figure 4: Team behavior using Strategy II.

15% reduction in normalized formation error.

4.3.3 Strategy III: Using local control augmented by a global goal and partial global information

Yet another improvement can be attained by providing the team with partial global knowledge about the path the group is to take. In the previous two cases, the right-hand turn by agent B prompted the other agents to change their alignments. However, B could have just been avoiding an obstacle, and thus the other agents should have continued along their present path without realignments. Without knowing anything about the route that the leader is following, the agents cannot always react properly to B's actions. Now, however, at the time of agent B's right-hand turn, let us assume that all the agents are told that the group should be headed toward waypoint X. With this partial global information, agents C and D can avoid the needless backtracking present in the previous case, and instead aim forward along the route toward the upcoming waypoint, as shown in figure 5, moderating their speeds as required to remain in alignment with their neighbors. In this manner, the agents achieve a much more more efficient cooperation, in which we attain average improvements of 38% in time and 22% in error over local control alone, and 32% and 9% average time and error improvements, respectively, over strategy II.

4.3.4 Strategy IV: Using local control augmented by a global goal and more complete global information

Yet another improvement can be achieved with the use of additional global information. Global knowledge of the route the group leader is tracking allows the agent followers to accurately predict future actions of the team members. In this example, knowledge of the global path being followed allows the agents to anticipate the right-hand turn, thus enabling the agents to the right of the leader to stop earlier in preparation for this turn (see figure 6). With such predictions, each agent can modify its actions to better maintain the formation. Using this strategy, we found an additional average error improvement of 12% over strategy III, which is an overall average improvement of 32% in normalized cumulative formation error over local control alone. However, we see little improvement in the mission completion time over strategy III, which is due to the fact that the agents making an error in formation in strategy III have time to correct their errors before the leader reaches the goal, thus not impacting the overall mission completion time.

5 Summary and conclusions

The design of the control laws governing the behavior of individual agents is crucial for the successful development of cooperative agent teams. These control laws may utilize a combination of local and/or global knowledge to achieve the resulting group behavior. A key difficulty in this development is deciding the proper balance between local and global control to achieve the desired emergent group behavior. This paper has addressed this issue by presenting some general guidelines and principles for determining the appropriate level of global versus local control, developed from quantitative studies of the keep formation case study. To summarize, the basic general principles and guidelines proposed in this paper are as follows:

- Global goals: If the global goals are known at design-time and all the information required for an agent to act consistently with the global goals can be sensed locally by the agent at run-time, these goals can be designed into the agents.
- Global knowledge: The more static, reliable, completely known, and easy-to-use the global knowledge is, the more practical its use in a global control law. The more unknown the global information, the more dependence the team will have on local control, perhaps combined with behavioral and environmental analysis to approximate global knowledge.

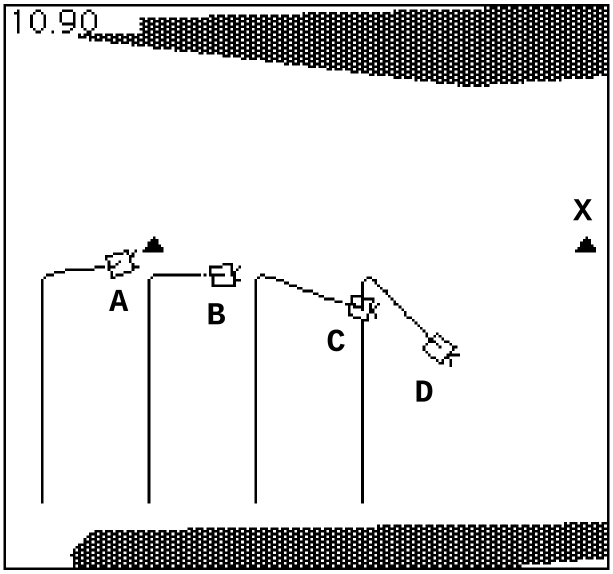

Figure 5: Team behavior using Strategy III.

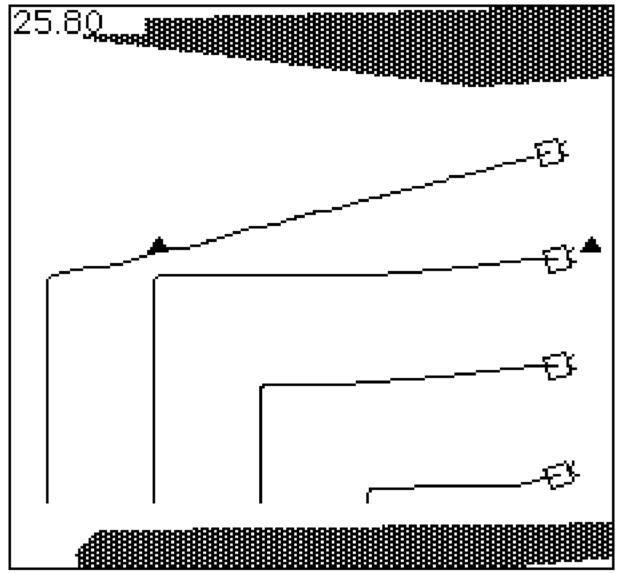

Figure 6: Team behavior using Strategy IV.

- Behavioral analysis: Behavioral analysis may provide a suitable approximation to global knowledge, and can thus be utilized to improve group cooperation. This method should be particularly useful when the agents possess a fixed set of discernible or communicable actions.
- Local knowledge: In many applications, particularly those in which *accomplishing* the task is more important than *how* the agents accomplish the task, local control may provide a suitable approximation to the optimal group behavior, thus eliminating the need for the use of global knowledge.
- Proper balance: Global knowledge should be used to provide general guidance for the longer-term actions of an agent, whereas local knowledge indicates the more short-term, reactive actions the agent should perform within the scope of the longer-term goals. This leads to the following basic principle:

> *Local control information should be used to ground global knowledge in the current situation. This allows the agents to remain focused on the overall goals of their group while reacting to the dynamics of their current situations.*

6 Acknowledgements

This research was funded in part by the Artificial Intelligence Center at Hughes Research Laboratories and in part by the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the MIT laboratory's artificial intelligence research is provided in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-85-K-0124.

References

[1] Brooks, R. A., P. Maes, M. J. Mataric, and G. More, "Lunar Base Construction Robots", *Proceedings of the IEEE International Workshop on Intelligent Robots and Systems*, Tsuchiura, Japan, pp389-392, 1990.

[2] Brooks, R. A., "A Robust Layered Control System for a Mobile Robot", *IEEE Journal of Robotics and Automation*, RA-2 (1), 1986.

[3] Connah, D., and P. Wavish, "An Experiment in Cooperation", *Decentralized A.I.*, Y. Demazeau and J. Muller, eds., Elsevier Science (North-Holland), 1990.

[4] Franklin, R. F., and L. A. Harmon, *Elements of Cooperative Behavior*, ERIM Internal Research and Development Final Report 655404-1-F, August, 1987.

[5] Maruichi, T., M. Ichikawa, and M. Tokoro, "Modeling Autonomous Agents and Their Groups", *Decentralized A.I.*, Y. Demazeau and J. Muller, eds., Elsevier Science (North-Holland), 1990.

[6] Miller, D. P., "Multiple Behavior-Controlled Micro-Robots for Planetary Surface Missions", *IEEE Systems, Man and Cybernetics Conference*, Studio City, CA, (November 1990).

[7] Noreils, F. R., "Integrating MultiRobot Coordination in a Mobile-Robot Control System", *Proceedings of the IEEE International Workshop on Intelligent Robots and Systems*, Tsuchiura, Japan, pp389-392, 1990.

[8] Payton, D. W., "Intelligent Real-Time Control of Robotic Vehicles", *Communications of the ACM*, 34(8), 1991.

[9] Payton, D. W., "An Architecture for Reflexive Autonomous Vehicle Control", *Proceedings of the IEEE Robotics and Automation Conference*, San Francisco, CA, (April 1986) pp1838-1845.

[10] Steels, L., "Cooperation Between Distributed Agents through Self-Organization", *Decentralized A.I.*, Y. Demazeau and J. Muller, eds., Elsevier Science (North-Holland), 1990.