MASSACHUSETTS INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

A.I. Memo No. 1357 March 1992

Local versus Global Control Laws for Cooperative Agent Teams

Lynne E. Parker

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 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 required to achieve the desired emergent group behavior. 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 demonstrate that local control alone is not sufficient to meet the goals of this particular task, 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 for cooperative applications of this type.

Copyright © Massachusetts Institute of Technology, 1992

This report describes research done at the Artificial Intelligence Center at Hughes Research Laboratories and at 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.

1 Introduction

The design of the control laws governing the behavior of each agent is of overriding importance to the successful development of teams of cooperating, situated, autonomous agents. These control laws determine not only how each agent will behave in its own local situation, but also how the group as a whole will behave 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, 7, 12, for example]. 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, in combination with any additional global information that may be available, 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. In the literature, this approach is usually typified by the use of communication between agents to convey partial or complete global information to other agents [3, 6, 8]. For example, in [8], Noreils describes how the use of communication can be used to synchronize group activity. Other methods of global information use are also possible, and are detailed in section 2.

Since successful applications have been demonstrated both with and without global information, the designer of a cooperative system must determine the appropriate balance between the use of global information and the use of local information. 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. The final section offers concluding comments and a summary of the general principles and guidelines put forth by this paper.

2 Global Control versus Local Control

We shall say that a collection of autonomous agents *cooperates* when the group accomplishes a task or achieves a goal that is beyond the capabilities of the agents individually. Implicit in this definition is the assumption that the team of agents is brought together for a common purpose, rather than just cohabiting. Striving toward a common goal, however, does not imply that every action by every agent must be a cooperative action, since conflict at a local level may occur without compromising the global goals. Our assumptions are thus somewhat weaker than the *benevolent agent* assumption, which assumes the absence of conflict among agents [11]. Within our framework, agents exhibit cooperation through the use of global and/or local control laws. In practice, of course, a continuum will exist between strictly global control laws 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 below. The following sections compare and contrast 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. Agents can exhibit higher levels of cooperation from the use of global control laws than would be possible with local control alone.

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 will be 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 influence on the design of cooperative teams is the time at which the global goals become known[9]. 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. Otherwise, a run-time use of global information will be required to achieve the desired group behavior. If the goals are not fixed or known at design-time, then they obviously cannot be designed into the agents. The agents must be provided with 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. For our purposes, explicit communication refers to those actions agents undertake that have no purpose other than to provide information to other agents. 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 utilize behavioral analysis of other agents. This involves interpreting 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.

As an example of the use of a global control law, consider a team of office-cleaning robots that have as their highest priority global goal to minimize the disruption of the humans in the office, and a next priority goal of minimizing energy use. The first goal requires that the robots clean only one room at a time, staying together until all robots have finished their tasks in that room. This goal can be achieved easily enough by using only local control laws. However, the agents cannot minimize their energy use without knowing the floor plan of the building they are cleaning and what the other robots are going to do. Without this knowledge, the agents are unable to predict their cleaning path, and thus cannot optimize their energy use. Another approach would be to design global information for a particular cleaning environment into the agents such that they minimize their energy use for that environment; however, the global goal of minimal energy use would probably not be achieved in a different environment. 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. At times (such as in the office-cleaning robots example), however, global information is absolutely required to achieve the desired results.

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, (referred to as *open systems* [5]), 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 usually 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. As described in the previous section, certain global goals cannot be attained through the use of local control laws alone. 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 get it done. For example, a group of robots could be designed to clean floors overnight, with the goal of having the floors cleaned by morning. In this application, it does not matter how the robots go about getting the floors cleaned, just as long as they get it done. The robots could be programmed with three competences: an obstacle avoiding behavior, a wander behavior, and a cleaning behavior. The wander behavior selects random directions for the robot to traverse, while the obstacle avoiding behavior avoids objects in the robot's path, and the cleaning behavior cleans the floor it is passing over. These behaviors, in combination with the critical number of robots, will allow the entire floor to be cleaned after a sufficient period of time. Here, we do not care that the robots are using an inefficient cleaning method that could be done much quicker using a more cooperative algorithm. However, other applications could require the robots to be more productive in their mission, thus forcing the agents to use more knowledge about the activities of the other team members.

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. For instance, a team of robots could have the mission to follow a specified road to a certain destination. During the exercise, these robots may have to wander off the road to avoid obstacles, but they will still be generally following the road. Thus, this global information is very static for this application. On the other hand, a different mission for these robots could be to investigate certain interesting locations to be determined during the mission. Here, global knowledge is still available, but is not known as far in advance. Nevertheless, it can be used successfully by the team to improve its cooperative performance. At the other extreme, a team of robots could be employed to chase down a dangerous prey. Since the actions of the prey are totally unknown, the robots will have no global knowledge about the path they will be taking. 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 using behavioral analysis. This type of analysis can be quite challenging, depending upon the complexity of the autonomous agents and the environment. When possible, behavioral analysis 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 analysis. Good results with behavior analysis should be expected particularly for teams of creatures possessing a fixed set of discernible actions. One of the primary difficulties with behavior analysis, however, lies in the 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 analysis.

3.2 Conflicts Between Global and Control Information

One may wonder whether the combination of local and global control will result in situations where the two types of control laws vote for conflicting actions. For instance, a global control law may tell a road-following robot to turn left while a local control law votes for a right turn. This type of problem may arise if the control laws are designed to compete with one another by having the global control law vote strictly according to global information, and the local control law vote strictly according to 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. This is a true case of global versus local conflict, and the designer must provide the agents with the ability to indicate which goals are more important — local or global. For instance, in the case study below, a robot trying to stay in formation with other robots will wait forever for its neighboring robot to catch up on the assumption that all the members of the robot team are following the specified route. If the straying robot has actually abandoned this route-following task in favor of some other task unknown to the rest of the team, it may be better for the other team members to ignore the straying robot (i.e. the local information) and proceed ahead in favor of the global information. The robots must thus be given 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 vehicles are able to remain focused on the overall goal of their group while reacting to the dynamics of their present contexts.

4 A Case Study: Keeping Formation

4.1 Task Description

We now turn to an illustration of the tradeoffs between local and global control in the context of a particular cooperative task: the "Keeping Formation" task (see figure 1) [1]. In this task, a group of robots is required stay in formation with one another while the leader of the group follows a prespecified route and while all agents avoid obstacles as they appear. These robots are "familiar" with two types of formations: (1) the line formation, in which the vehicles are aligned side-by-side, and (2) the column formation, in which the vehicles are aligned front to back. Each of these robots has the ability to sense the location of its neighboring vehicles relative to itself (local knowledge) and is physically constrained by the inability to move backwards. The global goal of this task is for the robots to maintain the specified formation in a manner that appears to a casual human observer to be human-driven, meaning that the robots should not allow huge or "unnatural" (an admittedly subjective measure) perturbations in the desired group formation [2]. These robots, designed using layers of behaviors [2, 10], are already provided with competences to avoid obstacles and to follow a specified route. Our job is to design the KEEP_FORMATION behavior to achieve the above global goal. The following sections describe the implementation of this case study along with various possible control strategies which illustrate the spectrum from strictly local to increasingly global control.

[1] In this discussion, the terms *robot*, *agent*, and *vehicle* are used interchangeably.

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

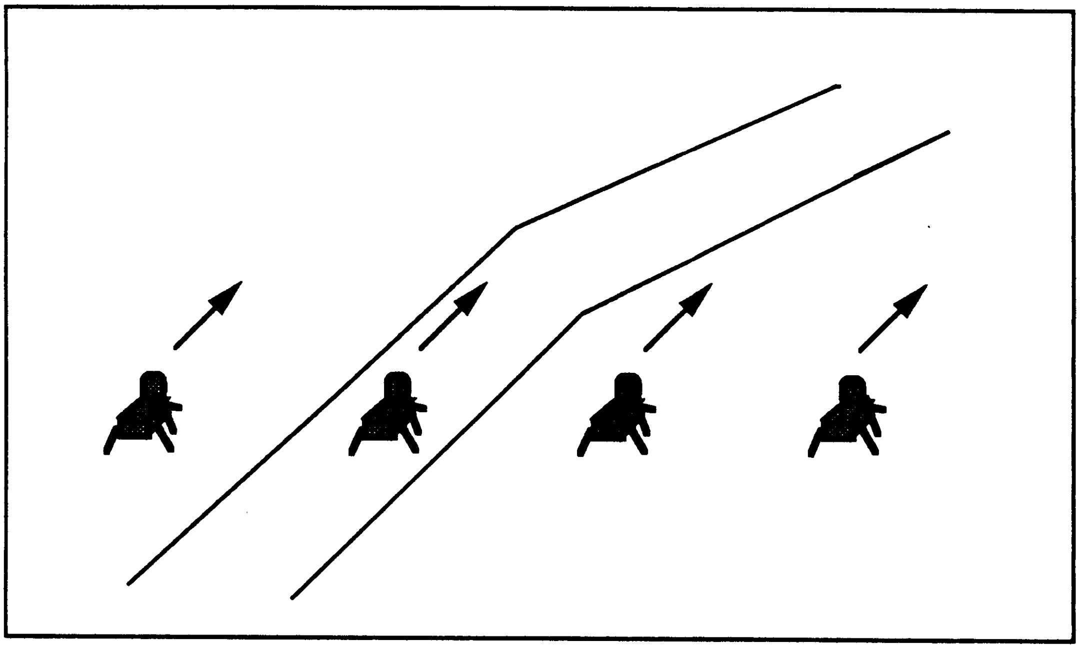

Figure 1: Keeping formation side-by-side.

4.2 Implementation

We have implemented and evaluated the various control strategies discussed in the next section by performing a wide range of experiments in simulation. The simulation system utilized for these experiments (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 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 either a column or 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. The main consideration in evaluating these strategies was the robustness with which the emergent group behavior satisfied the global goal of staying in formation while appearing to be human-driven.

4.3 Control Strategies

We now present five alternative control strategies for accomplishing the KEEP_FORMATION task, describing the type of local and/or global control information used and the results of our implementation of each strategy.

4.3.1 Using Local Control Alone

At first glance, it appears that KEEP_FORMATION can be achieved using local control laws alone. Each vehicle could be assigned a leader and then use a simple control law that directs it toward a prespecified offset and direction from its leading vehicle. As the group leader moves forward along the path, the other vehicles will 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 could law would perform adequately well. However, consider the case illustrated in figure 2. Robot B is the overall leader, robots A and C are following robot B, and robot D is following robot C. In following its leader, robot A seeks to always locate itself a preset distance to the left of B, while robots C and D strive to be located the same distance to the right of their respective leaders. In figure 2, the group leader, B, is making a right-hand turn. This turn could be interpreted by its followers without global information as either an attempt to avoid the obstacle or as a turn along the route being followed. The follower can only guess which of these options holds in this situation. Since the followers are using strictly local information in this case, they continue to follow the same rules as before. Robot A performs satisfactorily, aiming toward the location the appropriate distance to the left of B. Since these vehicles cannot back up, robot C turns around and aims toward a location to the right of B. Now, however, we have a problem with robot D. It aims as usual toward the right of C, but this position is out of formation with the rest of the group. Figure 3 illustrates the problem. These vehicles are not performing in a manner humans would consider intelligent. Thus, in this case, local control information is not sufficient to achieve the desired global goals.

4.3.2 Using Local Control Augmented by a Global Goal

An improvement on the situation is to provide the robots with knowledge of the global goal of the group. Now, since the robots are "aware" that they should achieve a global linear formation, they select their positions after robot B's right-hand turn based on the global formation, while still remaining responsive to the local dynamics of the vehicles adjacent to them. With this information, robots A and C will aim toward the same positions as in the previous case, but robot D will now head 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

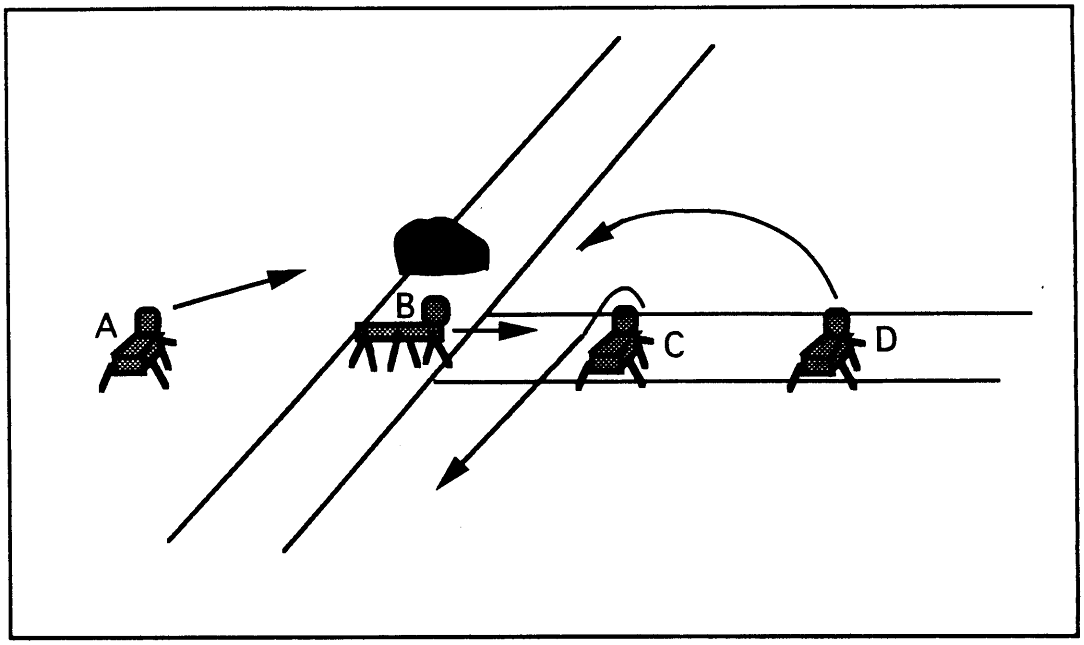

Figure 2: Using local control alone.

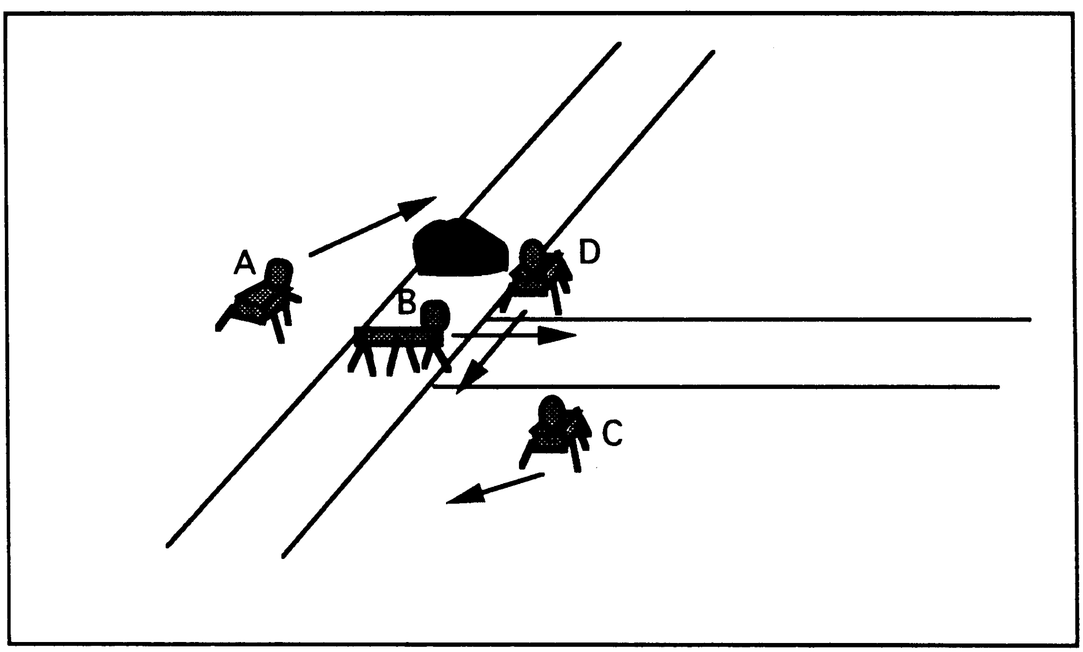

Figure 3: Problem resulting from local control alone.

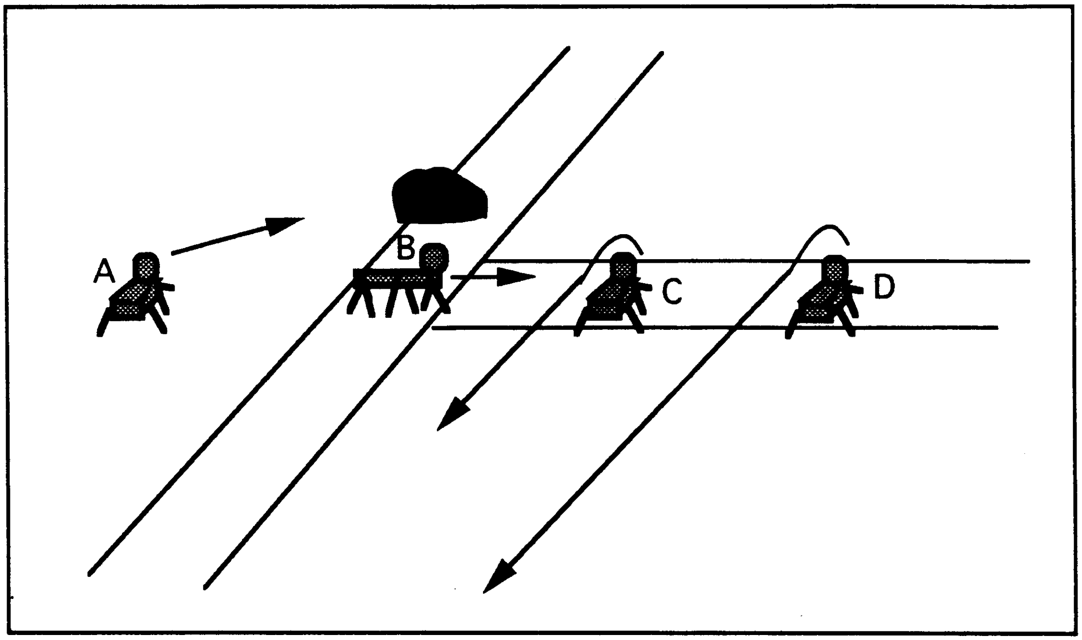

Figure 4: Local control augmented by a global goal.

is clear that knowledge and use of the global goal can often yield improved group cooperation.

4.3.3 Using Local Control Augmented by a Global Goal and Behavioral Analysis

Yet another control strategy is to endow the agents with the ability to approximate global knowledge by interpreting the actions of other agents in the group. Under this method, the agents use a model of the actions of other group members to determine the best response to those actions. This approach is best illustrated with a different type of formation, the column formation, in which the robots are required to follow directly behind their leaders at a prespecified offset. In this case, a robot uses information about its leader's leader along with a heuristic rule to determine whether, for example, a right-hand turn by its own leader is intended to avoid an obstacle, or indicates a turn along the path being followed (see figure 5). If the leader's leader has also turned right, then a bend in the path is indicated; otherwise, the robot can assume its leader is simply avoiding an obstacle. Unfortunately, such modeling cannot handle the right-hand turn in the side-by-side formation discussed earlier, as no additional leader is available to provide more information, and thus the vehicles have no way of disambiguating a right-hand turn by the group leader.

Our implementation of this control strategy used a number of behavior analysis heuristics such as the one above to achieve a fairly adequate group behavior of keeping formation. Unfortunately, the resulting control law was not based on any

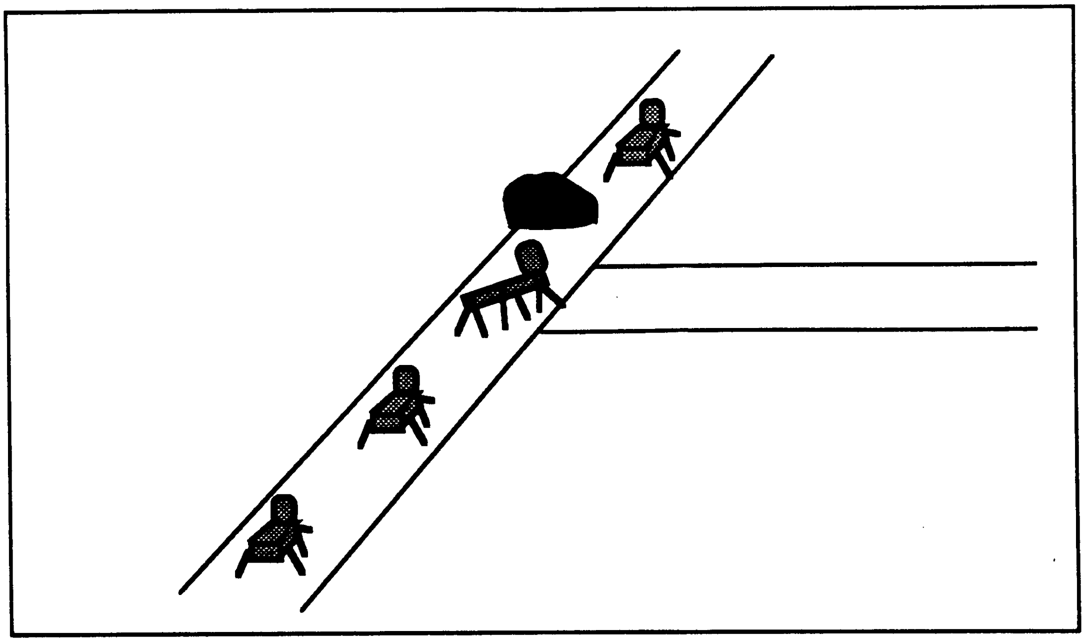

Figure 5: Local control augmented by a global goal and behavioral analysis.

obvious general principles, but instead was based on a number of special cases that occurred in specific situations and was very difficult to understand. New situations easily broke this control law, because not all of the dynamics of the environment and possible group applications could be captured. Thus, this method failed to meet the application requirements for robustness and efficiency throughout the mission. From this experience, it became clear that more global knowledge must be used to satisfy the global goal.

4.3.4 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 robot B prompted the other vehicles to change their alignments. However, B could have just been avoiding an obstacle, and thus the other vehicles should have continued along their present path without realignments. Without knowing anything about the route that the leader is following, the vehicles cannot always react properly to B's actions. Now, however, at the time of robot B's right-hand turn, let us assume that all the vehicles are told that the group should be headed toward waypoint X. With this partial global information, vehicles 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 6, moderating their speeds as required to remain in alignment with their neighbors. In this manner, the vehicles achieve an even more efficient cooperation.

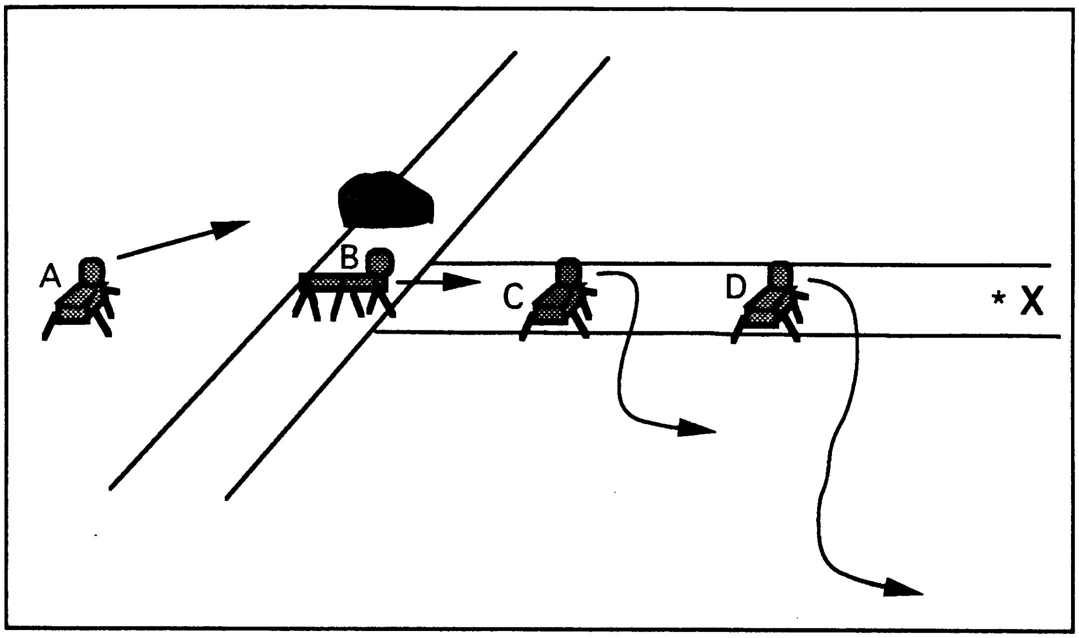

Figure 6: Local control augmented by a global goal and partial global information.

4.3.5 Using Local Control Augmented by a Global Goal and More Complete Global Information

A final improvement can be achieved with the use of additional global information. Global knowledge of the route the group leader is tracking allows the robot followers to accurately predict future actions of the team members. In this example, knowledge of the global path being followed allows the robots to anticipate the right-hand turn, thus enabling the vehicles to the right of the leader to stop earlier in preparation for this turn (see figure 7). With such predictions, each vehicle can modify its actions to yield a more efficient, and seemingly more intelligent, global cooperation.

The implementation of this control strategy resulted in a very robust, easy to understand routine. It performed very well in all of our experiments, covering a wide range of varying situations. Unlike the behavior in the other control strategies, the emergent group behavior with this strategy appears to be very human-like, thus satisfying our original global goal of staying in formation while appearing to be human-driven.

5 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

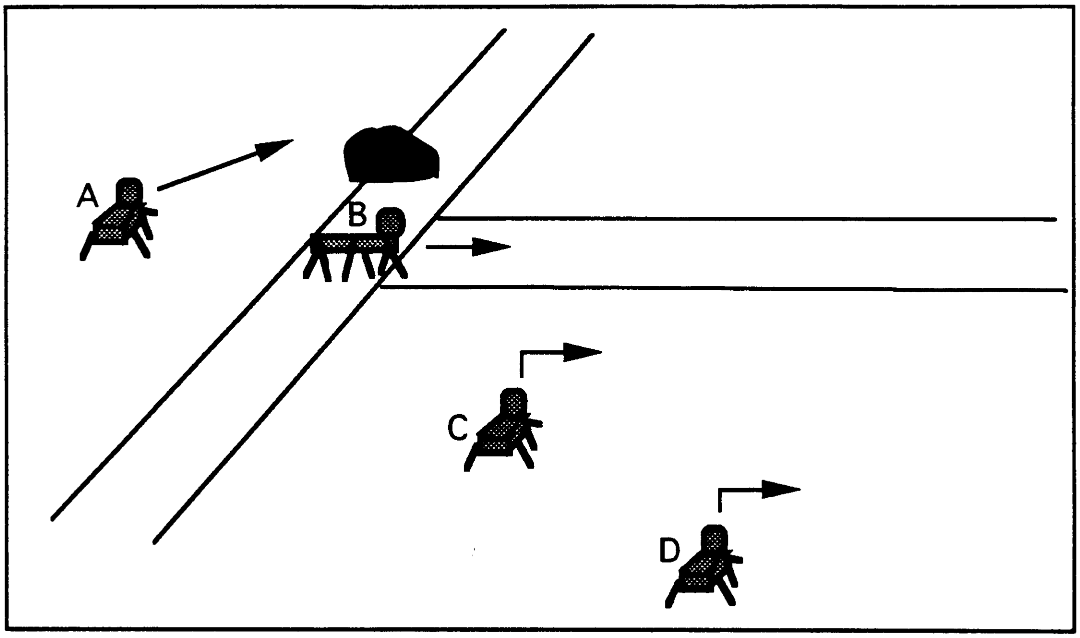

Figure 7: Local control augmented by a global goal and more complete global information.

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. To summarize, the basic general principles and guidelines proposed in this research 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.

- 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.

These principles and guidelines were illustrated and implemented in a "keep formation" case study, which presented several alternative control strategies along the local versus global spectrum. In this case study, we demonstrated that local control alone was not sufficient to meet the goals of this task, and that increasing use of global knowledge resulted in a steadily improving group cooperation.

6 Acknowledgements

My thanks and gratitude go to Scott Harmon, Dave Payton, Sam Yang, and Dave Keirsey whose comments and insight contributed greatly to this research and to this paper. Additional thanks go to Sam Yang, who developed a large number of the experiments and who implemented the control strategy involving behavioral analysis. I am also appreciative to Rod Brooks, Lynn Stein, and Maja Mataric for their helpful comments on earlier drafts of this paper.

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] Hewitt, C., "Offices are Open Systems", *ACM Transactions on Office Information Systems*, 4(3), 1986.

[6] 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.

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

[8] 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.

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

[10] 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.

[11] Rosenschein, J., and M. R. Genesereth, "Deals Among Rational Agents", *Proceedings of the 1985 International Joint Conference on Artificial Intelligence*, pp91-99.

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