Integration of Reactive and Telerobotic Control in Multi-agent Robotic Systems

Ronald C. Arkin and Khaled S. Ali

Mobile Robot Laboratory
College of Computing
Georgia Institute of Technology
Atlanta, GA, USA 30332-0280
e-mail: arkin@cc.gatech.edu, kali@cc.gatech.edu

Abstract

Multi-agent schema-based reactive robotic systems are complemented with the addition of a new behavior controlled by a teleoperator. This enables the whole society to be affected as a group rather than forcing the operator to control each agent individually. The operator is viewed by the reactive control system as another behavior exerting his/her influence on the society as a whole. Simulation results are presented for foraging, grazing, and herding tasks. Teleautonomous operation of multi-agent reactive systems was demonstrated to be significantly useful for some tasks, less so for others.

1 Introduction

Reactive multi-agent robotic societies can be potentially useful for a wide-range of tasks. This includes operations such as foraging and grazing (e.g., [1,9,6]) which have applicability in service (vacuuming and cleaning), industrial (assembly) and military (convoy and scouting) scenarios.

Although promising results have been achieved in these systems to date, purely reactive systems can still benefit from human intervention. Many purely reactive systems are myopic in their approach: they sacrifice global knowledge for rapid local interaction. Global information can be useful and it is in this capacity that a teleoperator can interact with a multi-agent control system.

A related problem in teleoperation is that a human operator is potentially overwhelmed by the large amount of data required to control a multi-agent system in a dynamic environment. This phenomenon is referred to as cognitive overload. The approach described in this paper provides a mechanism to significantly reduce the teleoperator's cognitive and perceptual load by allowing the reactive system to deal with each robot's local control concerns. Two principal mechanisms to achieve this are

by allowing the operator to act either as a constituent behavior of the society or to allow him/her to supervise the societal behavioral sets and gains, acting only as needed based upon observable progress towards societal task completion.

In this research, the teleoperator is allowed to control whole societies of agents; not one robot at a time, but rather controlling global behavior for the entire multi-agent system. This is a straightforward extension of our work in both multi-agent robotic systems [1] and teleautonomy [2]. The end product is a simple way for a commander to control large numbers of constituent elements without concern for low-level details (which each of the agents is capable of handling by themselves). In essence, the teleoperator is concerned with global social *strategies* for task completion, and is far less involved with the specific behavioral tactics used by any individual agent.

2 Single agent teleautonomous control

Our previous results [2] in the integration of reactive and telerobotic control in the context of single agents provide the basis for our extension of this concept into multi-agent societies. In this earlier work we have shown that a teleoperator can interact with a reactive robot in at least two different ways:

- **Teleoperator as a schema**: Here the human acts as an additional behavior in the already existing collection of behaviors that are active within the robot. Using a schema-based methodology [3], each active behavior contributes a vector that is related to the agent's intentions - such as to get to a particular object, not crash into something, etc. The teleoperator's intentions are introduced at the same level - as another schema contributing forces in the same manner as all the other behaviors do.

- **Teleoperator as a supervisor**: In this case, the teleoperator changes the behavioral settings of the robot as it moves through the world, essentially

changing its "personality". For example, the robot can become more aggressive by increasing its attraction towards a desirable object or decreasing its repulsion from obstacles.

In schema-based reactive control [3], each active behavior (schema) provides its own reaction to the environment by creating a vector response to a specific perceptual stimulus. The entire set of vector outputs created by all active schemas is summed and normalized and then transmitted to the robot for execution. No arbitration is involved, rather a blending of all active concurrent behaviors occurs. The system at this level is completely reactive, not retaining knowledge of the world or the agent's past performance.

3 Multi-agent Teleautonomous Control

Our laboratory is conducting extensive research in multi-agent robotic systems [1,4,5] both in simulation and on our 3 Denning Mobile Robots. Robotic systems are specified as a finite state acceptor that specifies the behavioral (schema) assemblages [7,8] and the transitions between them. An example state machine for a foraging task appears in Figure 1. In this figure there exist three distinct high-level behavioral states for each agent:

- *Wander* - which consists of a high gain and long persistence **noise** schema that is used to produce wandering while having moderate inter-robot repulsion to produce dispersion coupled with significant obstacle repulsion (**avoid-static-obstacle** schemas).

- *Acquire* - which consists of using a **move-to-goal** schema to move towards a detected or reported attractor (depending on the communication strategy used [4]) with a reduced inter-robot repulsion to allow for multi-robot convergence on attractors and continued obstacle avoidance (again provided by the **avoid-static-obstacle** schema). A small amount of **noise** is still injected into the system to facilitate navigation [3].

- *Deliver* - which occurs after acquisition of the attractor and results in delivery of the object back to home base by one or more agents. The same behaviors are used as in the *acquire* state with the goal location now being the home base.

Space prevents a full discussion of the mechanisms for reactive multi-agent control. The interested reader is referred to [1,4] for more information.

3.1 Implementation

In the results presented below, teleoperation is implemented as an additional schema in the system (the teleoperator as a schema approach). Based on the


Figure 1: Behavioral States for Foraging Task

instructions of a human agent, the **teleautonomy** schema contributes a vector in the same way as do the other schemas, such as **move-to-goal** or **avoid-static-obstacle**. Unlike the other schemas, however, which produce different vectors for each robot, the **teleautonomy** schema produces the same output for all of the robots. Thus, if the human agent tells the robots to go north, then all the robots receive the same instructions. The vector produced by the **teleautonomy** schema is summed with the vectors produced by the other schemas in each agent to produce a combined vector which determines the overall direction and rate of travel of the robot. In this way, the robots use environmental knowledge provided by the human agent in conjunction with their other goals, such as not to collide with obstacles or each other, rather than having the teleoperator's goals completely override the robots' other behaviors.

The human agent has control over both the direction and magnitude of the vector produced by the **teleautonomy** schema. The operator uses an on-screen joystick (via a mouse) to provide input to the **teleautonomy** schema.

3.2 Simulation Environment

The system is tested on a graphical simulation environment prior to its port to our Denning robots. The objects represented in the simulation environment include robots, obstacles, and attractors. Each robot's trail is depicted by a broken line. Every robot uses the same set of behaviors (a homogeneous society), but the sensory input for each is different, depending on the robot's location within the environment. The robots can sense objects within a certain radius around them. They have the ability to distinguish whether a sensed object is an obstacle, another robot, or an attractor.

The agents have a limited form of communication between themselves. A robot is capable of communicating its current behavioral state or the location of an attractor that it is acquiring or delivering [4]. The communication is simulated by using shared memory. Each agent only looks at this shared memory when there is no attractor

within its sensing range.

In tasks that require the movement of attractors, more than one robot is allowed to contribute to the transport of the object at the same time. The net effect of this cooperation is simulated by having the robots move the attractor farther during each time unit if there are more robots carrying it. The distance traveled while carrying an attractor is determined by the mass of the object and the number of robots carrying it.

3.3 Tasks

The use of teleoperation in multi-agent systems was tested for three different tasks. The tasks were foraging, grazing (vacuuming), and herding the robots into a pen. In all three tasks, a teleoperator provided input at his own discretion.

In the foraging task, the robots wander around looking for attractors. When a robot finds a target object, it communicates its location to the other agents while simultaneously moving to acquire it. After its acquisition, the robot carries the attractor back to a home base, then deposits it, and finally returns back to the task of searching for more attractors. If a robot cannot detect an attractor within its sensory radius, it checks to see if any other agent has communicated the location of another candidate goal object. If so, then the robot proceeds to acquire it.

In the grazing task, the robots are placed in an environment studded with obstacles. Initially, all of the floor that is not covered with obstacles is considered "ungrazed". Each section of the floor that is ungrazed is treated as if it had a a large number of distributed attractors on it. That is, a robot can sense an ungrazed section of floor from a distance, and it can also communicate the presence of an ungrazed section of the floor to the other robots. When an agent passes over an ungrazed region it becomes grazed (clean). The task is completed when a certain percentage of the floor, specified in advance, has been grazed. The robots normally wander randomly until an ungrazed floor area is detected.

In the herding task, there is a pen with an opening formed of obstacles in the simulation environment. All the agents are initially outside of the pen. The robots remain in the *wander* state for the duration of the run and wander aimlessly in random directions. The robots are repulsed by the obstacles and the other robots. The task is to get all of the robotic agents inside the pen at the same time.

4 Results

For the foraging and grazing tasks, tests were conducted that compared the total number of steps taken by the robots to complete the tasks with and without the help of a teleoperator. For the herding task, no comparison could be made between teleoperation and no teleoperation, because the likelihood of all the robots wandering into the pen by themselves at the same time is virtually nil. Interesting information was gained about this task nonetheless.

4.1 Foraging Results

In the tests conducted for the foraging task, three robots were used to gather six attractors. The density of obstacles in the environment was 10%. The total number of steps required to finish the task was measured both with and without a teleoperator. If teleoperation is used wisely, it can significantly lower the total number of steps required to complete the task by greatly reducing the time spent in the *wander* state (i.e., the number of steps that the robots spend looking for attractors). If none of the agents currently sense an attractor, then the teleoperator can assist by guiding the robots in one's direction. However, once the robots can sense an attractor, the teleoperator should stop giving instructions, unless the instructions are to deal with a particularly troublesome set of obstacles. In general, the robots perform more efficiently by themselves than when under the control of a teleoperator if the agents already have an attractor in sight. The human's instructions tend to hinder the robots if they are already moving to acquire or return an attractor. Indeed, when teleoperation is used at all times, the overall number of steps required for task completion often increases when compared to no teleoperation at all. However, if the human only acts to guide the robots toward an attractor when none are currently detected, significant reductions in time for task completion are possible. The average over several experimental runs of the total number of time steps required for task completion when teleoperation was used in this manner was 67% of the average task completion time when no teleoperation was used.

An example trace of a forage task without teleoperation is shown in Figure 2(a). Another trace of the same forage task with a human teleoperator helping the robots find the attractors when they did not have one in sensing range is shown in Figure 2(b). The robots all started at the home base in the center of the environment. In the run without teleoperation, the robots immediately found the two closer attractors at the lower right. Then they quickly found the two closer attractors at the upper right. At this point, the robots did not immediately detect the remaining two attractors. Two of the three agents proceeded by chance to the left and upper left sides of the environment, wandering unsuccessfully while seeking an attractor. Eventually, the other robot found the attractor in the lower right corner, and the other two robots moved to help with its return. After delivering it to the home base, the robots wandered again for a while without finding the last attractor. Finally, the last attractor

was detected and successfully delivered to home base. In the same world with the help of a human teleoperator, the two protracted periods of wandering while searching for attractors are avoided. This indicates the types of environments where the use of teleoperation for the forage task is most beneficial. The greatest benefit from teleoperation can be seen when there are one or more attractors that are far from both the home base and the start locations of the robots. Typically, this is when the robots do not sense the target objects without wandering for a while.

4.2 Grazing Task Results

For the grazing task, five robots were used. A sample run of a grazing task is shown in Figure 3. In this case, the robots performed poorly when a large amount of teleoperation was involved. Teleoperation only proved useful when the robots had difficulty in locating a section of ungrazed floor. When the robots had already detected an ungrazed area, they performed better without any input from the teleoperator. The agents' performance degraded considerably, often taking several times longer to complete the task, if teleoperation was used when a robot had already located an ungrazed floor area. Moreover, since remaining untreated areas tend to be clustered together in large patches, the agents typically do not need to spend long periods of time looking for another ungrazed spot (which is opposite the case of the foraging task discussed above). Therefore, the use of teleoperation did not help significantly with the grazing task. When teleoperation was used solely to help the robots find ungrazed floor area when they were not already cleaning, only a 4% improvement in average task completion time performance was observed when compared to not using teleoperation. Thus, when used wisely, teleoperation helped somewhat but not to a large extent.

4.3 Herding Task Results

For the herding task, five robots were herded into a pen that was 36 units long by 18 units wide, with a 12 unit long door in one of the longer sides. All of the robots started at one spot on the side of the pen with the door. In most test runs, the teleoperator encountered no difficulty with this task. He was able to herd the robots into the pen without problems. In some of the test runs, there were a few minor difficulties, such as robots wandering back out of the pen after having been herded in. However, the teleoperator was still able to complete the task without much frustration and in a reasonable amount of time. The results of a test run for the herding task are shown in Figure 4.

(a)

(b)

Figure 2: Foraging task.
(a) Without Teleoperation (b) With Teleoperation

5 Analysis

Some conclusions can be ascertained from the studies conducted thus far. It should be remembered, however, that these are preliminary studies, and there are many variables that have not yet been explored. For instance, we intend to explore the effects of teleoperation while varying the number of robots for a particular task, to study the role and impact of different inter-agent communication methods on teleoperation, and to conduct an analysis of what types of environments teleoperation is most suited for.

The use of the **teleautonomy** schema in conjunction with the robots' other behaviors proved particularly effective for the foraging task, while being less so for the grazing task. Herding the robots into a pen was also feasible using this method. During foraging, the best results were observed when teleoperation was used only to guide the robots in the direction of an attractor if one had not been previously sensed. For the grazing task, teleoperation was not significantly better than no teleoperation, although minor improvements were observed. The best results were again seen when teleoperation was used in guiding the robots towards dirty areas that were outside the sensor (or communication) range of the agents.

Two conceivable improvements can be implemented for the herding task regarding teleoperation. The first is to allow the teleoperator to turn off the input from the teleoperation schema for specific robots but not for others, allowing the operator to concentrate on the outside robots without worrying what effects his actions will have on robots already inside the pen. The other improvement is to allow the teleoperator to completely stop a robot's movement when it is inside the pen. In this way, the output of the teleoperation schema could be thought of as producing a vector that nullifies the vectors produced by the robot's other schemas. However, both of these strategies involve producing different output for the **teleautonomy** schema for different robots. This means that the teleoperator would have a greater burden, defeating the purpose of this research in reducing the cognitive workload.

Another important point is that if the teleoperator is given unrestricted control of the magnitude of the vector produced by the teleoperation schema, it is possible for the teleoperator to force a robot to collide with obstacles and other robots. The teleoperator must be careful when increasing the gain of the **teleautonomy** schema so that this does not occur. It can be a delicate task to override the output of the **noise** schema, which is necessary to cause the robots to quickly move in a particular direction, while not overriding the **avoid-static-obstacle** behaviors.


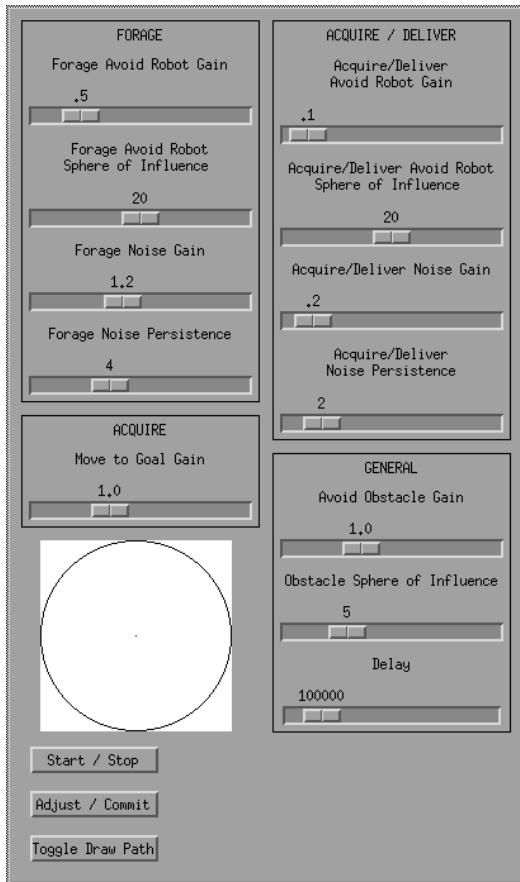Figure 3: Grazing Task

Figure 4: Herding task

Figure 5: User interface

6 Summary

A method by which multi-agent reactive robotic societal task execution can be influenced via human intervention has been demonstrated. This has been shown for a range of tasks including: improving the efficiency of foraging behavior; limited impact on improving grazing (vacuuming) activity; and the ability to congregate agents in a small confined area (herding) under human guidance.

The next phase of this research has involved extending the simulations to include the second type of teleoperation described in Section 2. We have recently developed an interface, shown in Figure 5, that allows the teleoperator to act as both a schema and as a supervisor of schema parameters. The next step is to port the results of this simulation onto our multi-agent robotic testbed consisting of three Denning Mobile Robots. An additional aspect of future research might involve the generation of an additional autonomous agent (e.g., another more informed robot) that could ultimately supplant some of the activities of the teleoperator.

Acknowledgments

References

[1] Arkin, R.C., "Cooperation without Communication: Multi-agent Schema Based Robot Navigation", *Journal of Robotic Systems*, Vol. 9(3), April 1992, pp. 351-364.

[2] Arkin, R.C., "Reactive Control as a Substrate for Telerobotic Systems", *IEEE Aerospace and Electronics Systems Magazine*, Vol. 6, No. 6, June 1991, pp. 24-31.

[3] Arkin, R.C., "Motor Schema-Based Mobile Robot Navigation", *International Journal of Robotics Research*, Vol. 8, No. 4, August 1989, pp. 92-112.

[4] Arkin, R.C., Balch, T., and Nitz, E., "Communication of Behavioral State in Multi-agent Retrieval Tasks", *Proc. 1993 IEEE International Conference on Robotics and Automation*, Atlanta, GA, May 1993, Vol. 3, pp. 588-594.

[5] Arkin, R.C. and Hobbs, J.D., "Dimensions of Communication and Social Organization in Multi-Agent Robotic Systems", *From animals to animats 2: Proc. 2nd International Conference on Simulation of Adaptive Behavior*, Honolulu, HI, Dec. 1992, MIT Press, pp. 486-493.

[6] Brooks, R., Maes, P., Mataric, M., and More, G., "Lunar Base Construction Robots", *IEEE International Workshop on Intelligent Robots and Systems (IROS '90)*, pp. 389-392, Tsuchiura, Japan, 1990.

[7] Lyons, D.M., and Arbib, M.A., "A Formal Model of Computation for Sensory-Based Robotics", *IEEE Transactions on Robotics and Automation*, Vol. 5, No. 3, June 1989.

[8] MacKenzie, D. and Arkin, R.C., "Formal Specification for Behavior-based Mobile Robots", *Mobile Robots VIII*, Boston, MA, Nov. 1993, pp. 94-104.

[9] Mataric, M., "Minimizing Complexity in Controlling a Mobile Robot Population", *1992 IEEE International Conference on Robotics and Automation*, Nice, pp. 830-835.