The Effect of Action Recognition and Robot Awareness in Cooperative Robotic Teams

Lynne E. Parker
Center for Engineering Systems Advanced Research
P. O. Box 2008, Oak Ridge National Laboratory, Oak Ridge, TN 37831-6364
email: ParkerLE@ornl.gov

Abstract

Previous research in cooperative robotics has investigated several possible ways of coordinating the actions of cooperative teams — from implicit cooperation through sensory feedback to explicit cooperation using the exchange of communicated messages. These various approaches differ in the extent to which robot team members are aware of, or recognize the actions of their teammates and the extent to which they use this information to effect their own actions. The research described in this paper investigates this issue of robot awareness of team member actions and its effect on cooperative team performance by examining the results of a series of experiments on teams of mobile robots performing a puck moving mission. In these experiments, we vary the team size (and thus the level of redundancy in team member capabilities) and the level of awareness robots have of their teammates' current actions and evaluate the team's performance using two metrics: time and energy. The results indicate that the impact of action awareness on cooperative team performance is a function not only of team size and the metric of evaluation, but also on the degree to which the effects of actions can be sensed through the world, the relative amount of work that is available per robot, and the cost of replicated actions. Based on these empirical studies, we discuss the impact of action recognition and robot awareness on cooperative team design.

1 Introduction

A primary aim in the development of cooperative robotic systems is to synthesize teams of robots that are able to accomplish missions which cannot easily be achieved, if at all, using single robot solutions. The advantages of cooperative systems over single robot solutions include the potential for increased fault tolerance, simpler robot design, widened application domain, and greater solution efficiency. However, the use of multiple robots introduces additional issues of robot control that are not present in single robot solutions. Foremost among these new issues is the question of how to achieve globally coherent and efficient solutions from the interaction of robots lacking complete global information.

Much existing cooperative robot work addresses the problem of global coherence and efficiency by designing robotic teams that use sensory information to glean implicit information on the activities of other robot team members and/or on the current state of the world (e.g. [3, 5]). In these approaches, no explicit communication among robots is utilized. A more difficult approach calls for robots to use passive action recognition to observe the actions of their teammates and modify their own actions accordingly (e.g. [4]). A third, quite common, approach involves explicit cooperation among team members by employing direct communication between robots to relay information on robot goals and/or actions to other team members (e.g. [1, 7]). These three approaches define a continuum in a robot team member's awareness, or recognition, of the actions or goals of its teammates, from implicit awareness through a teammate's effect on the world, to passive observation of a teammate's actions or goals, to explicit communication of a teammate's actions or goals.

Each of these approaches to cooperation has its own advantages and disadvantages. *Implicit cooperation* (also called *cooperation through the world*) is appealing because of its simplicity and its lack of dependence upon explicit communications channels and protocols. However, it is limited by the extent to which a robot's sensation of the world reflects the salient states of the mission the robot team must accomplish. *Passive action recognition* is appealing because it does not depend upon a limited-bandwidth, fallible communication mechanism. As with implicit cooperation, however, it is limited by the degree to which a robot can successfully interpret its sensory information, as well as the difficulty of analyzing the actions of robot team members. Finally, the *explicit communication* approach is appealing because of its directness and the ease with which robots can become aware of the actions and/or goals of its teammates. However, it is limited in the area of fault tolerance and reliability, in that it can be highly dependent upon the presence of a reliable communications medium for the successful accomplishment of a cooperative mission, and it usually depends upon a limited-bandwidth communications channel.

Together, these various approaches raise interesting questions concerning the impact of a robot team member's awareness (or lack thereof) of the actions and/or goals of its teammates. For implicit cooperative systems and those using passive action recognition, the question is: What is the impact of a limited ability to sense the effect of robot actions through the world? For explicit communication systems, the question is: What is the impact of communication failure, which leads to the lack of awareness of team member actions/goals? Or, conversely, what benefits can be gained by using explicit communication to increase robot awareness of team member actions/goals?

This paper investigates this issue of robot awareness of team member actions and its effect on cooperative team performance by examining the results of a series of experiments on teams of mobile robots performing a puck moving mission. In these experiments, we vary the team size (and thus the level of redundancy in team member capabilities) and the level of awareness robots have of their teammates' current actions and evaluate the team's performance using two metrics: time and energy. The results indicate that the impact of action awareness on cooperative team performance is a function not only of team size and the metric of evaluation, but also on the degree to which the effects of actions can be sensed through the world, the relative amount of work that is available per robot, and the cost of replicated actions. Based on these empirical studies, we discuss the impact of action recognition and robot awareness on cooperative team design.

The following section gives a brief overview of the previous work in this area. Section 3 describes the experimental setup for this research, including descriptions of the robots used in these studies and an overview of the puck moving mission. Section 4 presents the results of our empirical studies, followed by a discussion of these results in Section 5. We offer concluding remarks in section 6.

2 Related Work

Previous research concerning the effect of robot awareness, or recognition, of team member actions has usually been described in terms of the effect of communication in cooperative robot teams (e.g. [2]). However, the current paper uses the phrase "robot awareness, or recognition, of team member actions" to describe precisely the issue of interest (awareness of teammate actions), rather than the plethora of information that could possibly be communicated between robot team members (e.g. a robot's bid on an activity in a negotiation system, the current local state of the environment near a given robot, the sensed location of an intruder, etc.), and to emphasize the fact that a robot may become aware of a team member's actions without the use of explicit communication (i.e. through passive action recognition).

MacLennan [6] investigates the evolution of communication in simulated worlds and concludes that the communication of local robot information can result in significant performance improvements. Werner and Dyer [9] examine the evolution of communication which facilitates the breeding and propagation of artificial creatures. The closest work related to our study is [2], in which Balch and Arkin examine the importance of communication in robotic societies performing forage, consume, and graze tasks. They found that communication can significantly improve performance for tasks involving little implicit communication through the world, and that communication of current robot state was almost as effective as communication of robot goals — results that are indeed consistent with the results discussed in the current paper. Our research, however, looks at a number of additional factors influencing the effect of robot awareness and action recognition, including the level of redundancy in team member capabilities, the relative amount of work available per robot, the cost of replicated actions, and the use of a performance metric other than time (namely, energy). In addition, our research was performed primarily on physical mobile robots, rather than on simulated robots.

3 Approach

3.1 The Robots

Our empirical studies were conducted on teams of R-2 robots purchased commercially from IS Robotics. Each of these robots is a small, fully autonomous wheeled vehicle measuring approximately 25 centimeters wide, 31 centimeters deep, and 35 centimeters tall. The R-2 has two drive wheels arranged as a differential pair, two caster wheels in the rear for stability, and a two-degree-of-freedom parallel jaw gripper for grasping objects. The robot sensor suite includes eight infrared proximity sensors for use in collision avoidance, piezoelectric bump sensors distributed around the base of the robot for use in collision detection, and additional bump sensors inside the gripper for use in measuring gripping force. Figure 1 shows three R-2s performing the puck moving mission, which is described in the next subsection.

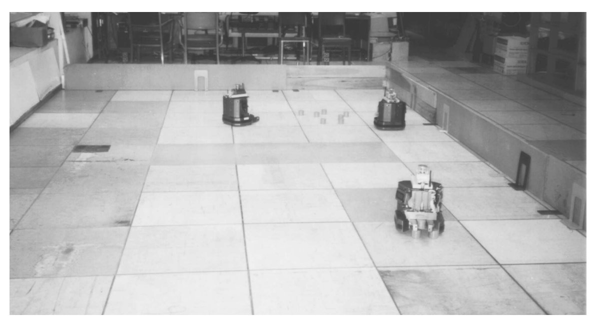

Figure 1: Three R-2s performing the puck moving mission.

A radio communication system allows robot team members to communicate with each other. This radio system is integrated with a global positioning system, which consists of a transceiver unit attached to each robot plus two sonar base stations for use in triangulating the robot positions. The positioning system is accurate to about 15 centimeters and is useful for providing robots with information on their own position with respect to their environment and with respect to other robot team members.

3.2 The Puck Moving Mission

The experimental mission used to study the issue of robot awareness of team member actions is a puck moving mission. Illustrated in Figure 2, this mission requires a collection of pucks to be moved by a team of robots. This mission requires robot team members to perform three distinct tasks: (1) the robot team must locate the pucks, (2) move them to a safe location, while also (3) periodically reporting the team progress to humans monitoring the system. These three tasks are referred to in the remainder of the paper as *find-locations*, *move-puck*, and *report-progress*.

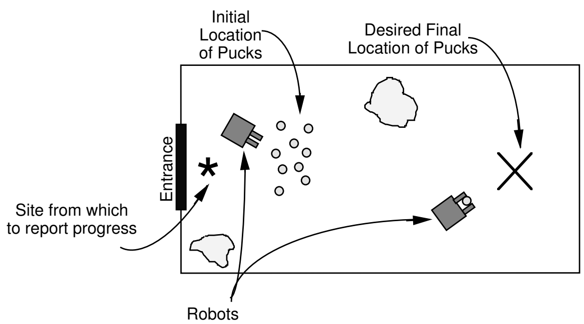

Figure 2: The experimental mission puck moving mission.

In this mission, the human monitor does not know the exact location of the pucks in robot coordinates, and can only give the robot team qualitative information on the initial location of the pucks and the final desired location to which the robots must move the pucks. Thus, the robots are told that the initial location is in the center of the front third of the room, and that the desired final location is in the back, center of the room, relative to the position of the entrance. A robot team member can then find the location of the pucks (i.e. execute the *find-locations* task) by first noting its starting (or home) x, y position and then following the walls of the room using its side infrared sensors until it has returned to its home location while tracking the minimum and maximum x and y positions it reaches. It then uses these x, y values to calculate the coordinates of the initial location of the pucks and the desired final location using the qualitative information given by the human. (To simplify the experiments, the area in which the robots work is rectangular, with sides parallel to the axes of the global coordinate system.) To prevent interference among robots, ideally only one robot at a time attempts to find the pucks, broadcasting the computed locations to the other team members once the task is complete.

A robot team member transports a puck to the goal location (i.e. executs the *move-puck* task) by moving to the general area of the initial puck location, wandering through that area looking for pucks, grasping and lifting a puck once it is located, moving to the final puck destination, and releasing the puck. In our experimental setup, the initial pile consisted of 10 pucks. Any number of robots could work on moving the pucks at the same time.

Finally, a progress report (i.e. the *report-progress* task) is performed by having an available robot team member return to the room entrance to radio the team's current mission completion status to a human monitor.

3.3 Robot Control

The robots in these experiments were controlled using the ALLIANCE software architecture [7, 8] — a behavior-based, fully distributed architecture that utilizes adaptive action selection to achieve fault tolerant cooperative control in robot missions involving loosely coupled, largely independent tasks. Robots under this architecture possess a variety of high-level functions that they can perform during a mission, and must at all times select an appropriate action based on the requirements of the mission, the activities of other robots, the current environmental conditions, and their own internal states. Since cooperative robotic teams often work in dynamic and unpredictable environments, this software architecture allows the team members to respond robustly and reliably to unexpected environmental changes and modifications in the robot team that may occur due to mechanical failure, the learning of new skills, or the addition or removal of robots from the team by human intervention. This is achieved through the interaction of mathematically modeled motivations of behavior, such as impatience and acquiescence, within each individual robot. These motivations allow robots to take over tasks from other team members if those team members do not demonstrate their ability — through their effect on the world — to accomplish those tasks. Similarly, it allows a robot to give up its own current task if its sensory feedback indicates that adequate progress is not being made to accomplish that task.

To enhance the robots' perceptual abilities, ALLIANCE utilizes a simple form of broadcast communication that allows robots to inform other team members of their current activities. Thus, at some prespecified rate, each robot broadcasts a statement of its current action. This one-way broadcast communication could be replaced by the use of passive action recognition, if such a capability were provided to the team members. No two-way conversations are employed in this architecture. Refer to [8] for more details on the ALLIANCE software control architecture.

3.4 Experiments

In this study, we used five experimental setups of the puck moving mission that varied the number of robots on the team and the level of awareness the robots had of the actions of their teammates. The five versions of this experiment were:

I. Two-robot team, full awareness of teammates' actions.

II. Three-robot team, full awareness of teammates' actions.

III. Two-robot team, no awareness of teammates' actions.

IV. Three-robot team, no awareness of teammates' actions.

V. One-robot team.

The one robot experiment was executed as a basis for comparison with the other four experiments. To achieve versions III and IV — those involving no awareness — the broadcast communications of each robot were turned off. Since these broadcasts were the sole mechanism in these experiments allowing robots to detect the actions of other robots whose effects could not otherwise be sensed through the world, the effect was to cause each robot to "think" it was working alone. For each of these experimental setups, we ran 10 missions to completion on the robots and collected data on the actions selected by each robot at each point in time and the length of time they required to complete those actions.

The outcomes of these experiments were evaluated based on their impact on the amount of time and energy required to complete the entire mission. To measure the energy usage, we made the approximation that a robot that is turned on but is idle (i.e. it is not moving either its wheels or its gripper) uses zero energy, whereas a robot that is using any of its four motors (i.e. right wheel, left wheel, grip, or lift) uses a unit quantity of energy per unit time.

4 Results

Our experimental investigations revealed a number of factors that play a role in the effect of action recognition and robot awareness on cooperative team performance. Along with the factors of team size and performance metric which were investigated directly in our experiments, our studies also identified three additional factors that influence the effect of robot awareness: the degree to which the effects of actions can be sensed through the world, the amount of work that is available to be performed per robot, and the cost of redundant actions. The results are therefore presented in terms of these influencing factors.

In this mission, each instance of the *move-puck* task is an action whose effects can be sensed through the world; that is, the robots do not try to move pucks that are no longer at the initial location. On the other hand, the *find-locations* task and each instance of the *report-progress* task are information-gathering or information-broadcasting types of actions whose effects cannot be sensed through the world by this robot team. Thus, we find that the lack of awareness of the actions of other robots causes a replication of the *find-locations* task and the *report-progress* task, but does not cause a replication of effort in the *move-puck* task. Figure 3 shows the average total time a team of three robots spent on each of the three tasks, both with and without awareness. Notice that much more time is spent collectively on the *find-locations* and *report-progress* tasks when the teams are not aware of the actions of their teammates. (A similar situation is seen with a two robot team.)

On average, the mission completion time of one run of version vsn of this mission, is given by:
$T(vsn) = t_{fl} + [m \times (t_{mp} + p(vsn) \times t_{rp})]/n$
where t_{fl}, t_{mp}, and t_{rp} are the average times to perform one instance of *find-locations*, *move-puck*, and *report-progress*, respectively, m is the number of moves required to move the spill, $p(vsn)$ is the number of

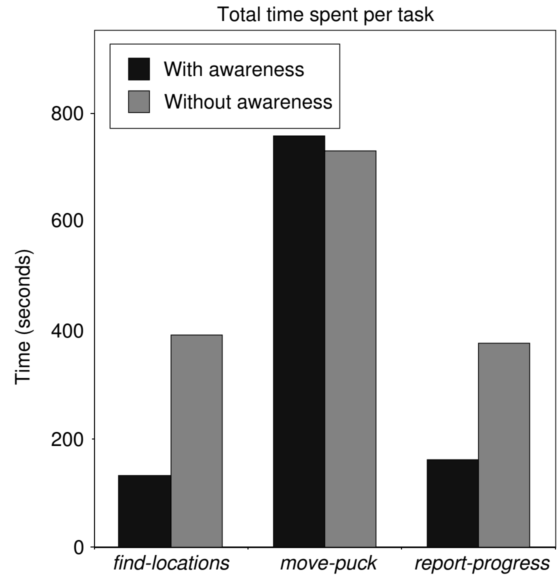

Figure 3: Total time spent collectively by the robot team on each of the three tasks both with and without awareness of team member actions.

progress reports performed in that version, and n is the number of robots.

On average, the total energy required to complete one run of version vsn for versions I and II is given by:
$E(vsn) = e_{fl} + m \times (e_{mp} + p(vsn) \times e_{rp}))$
and for versions III and IV is given by:
$E(vsn) = n \times e_{fl} + m \times (e_{mp} + p(vsn) \times e_{rp}))$
where e_{fl}, e_{mp}, and e_{rp} are the average energies to perform one instance of *find-locations*, *move-puck*, and *report-progress*, respectively,

To analyze how serious the replication of effort due to limited awareness can be, we define a relative cost measure that allows us to quantify the effect of limited awareness on the team performance as the cost of the replicated actions varies. This relative cost measure for a given task is defined simply as the ratio of the cost of performing that replicated task to the cost of performing each task of the entire mission.

5 Discussion

The puck moving mission explored in these experiments proved to be a very illuminating test environment for our cooperative robotics research. The three tasks required by this mission varied sufficiently in their requirements so as to bring to light many issues in cooperative team performance. We now discuss these issues in the context of the puck moving task, and generalize where possible.

5.1 Effects on Time Requirements

Figures 4 and 5 show the effect of varying the relative costs of the *find-locations* and *report-progress* tasks on the average time required to complete the mission for each experimental version. For both of these tasks, the worst time performance by far occurred when using only a single robot. Thus, adding robots is beneficial in this mission regardless of the level of awareness. The relative ordering of the remaining four scenarios in terms of time is the same in both cases, with the three robot team with full awareness performing the best, followed by the three robot team with no awareness, the two robot team with full awareness, and the two robot team with no awareness. Note that the three robot teams always give a better time performance than two robot teams for this mission regardless of the presence or absence of awareness. Clearly, dividing up a given amount of work across more robots without replicating any tasks would be expected to result in the mission being completed more quickly than for a two-robot team either with or without awareness. However, why would a three-robot team *without* awareness perform more quickly than a two-robot team *with* awareness? The answer lies in the tradeoff between the beneficial effects of redundancy for tasks whose effects can be sensed through the world and the adverse effects of redundancy for tasks whose effects cannot be sensed through the world. When the cost of the tasks replicated due to lack of awareness is relatively low, the redundancy in task coverage for those actions whose effects *can* be sensed through the world is the dominating factor, and thus a noticeable improvement occurs. In this specific example, then, we see that when the *report-progress* relative cost is low, the three-robot team, even without awareness, provides a decided advantage because it moves the pucks more quickly without incurring much of a penalty for repetitive progress reports. As the cost of *report-progress* increases, however, this advantage dwindles. The lesson learned here is that although awareness is helpful for a fixed-sized robot team, a larger team without awareness may actually be able to perform the mission more quickly if a significant proportion of the mission consists of tasks whose effects can be sensed through the world.

Also observe that that as the relative cost of the *find-locations* task increases, the relative benefits of awareness and team redundancy decrease. In the case of the *find-locations* task, lack of awareness causes no time penalty — it simply leads all the robots to perform the *find-locations* task once at the beginning of the mission. Thus, when the *find-locations* cost is low, a significant proportional benefit in preventing the repetitive execution of the tasks *move-puck* and *report-progress* can be drawn from awareness and team redundancy. On the other hand, if the fixed, non-shareable startup cost of the mission is large compared to those portions of the mission that can be shared across the team, then it makes little time difference, proportionally, whether the team has awareness or task redundancy. The lesson learned here is that awareness and task redundancy can help with the time requirements of the mission only if the mission includes a fair percentage of ¡shareable tasks, especially those whose effects cannot otherwise be sensed through the world.

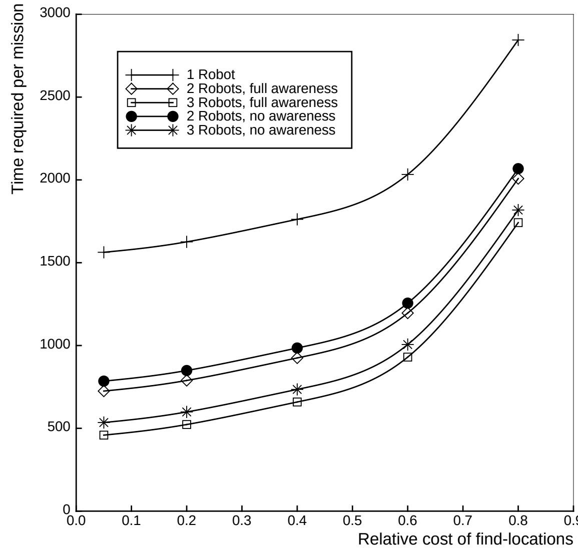

Figure 4: The average time required to complete the mission as a function of the relative time cost of the *find-locations* task.

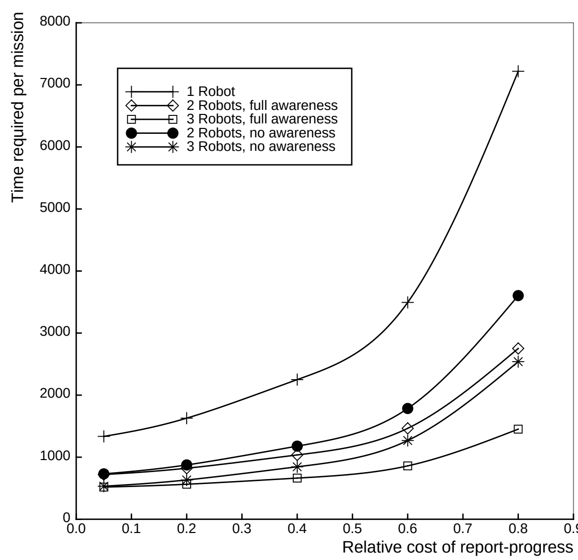

Figure 5: The average time required to complete the mission as a function of the relative time cost of the *report-progress* task.

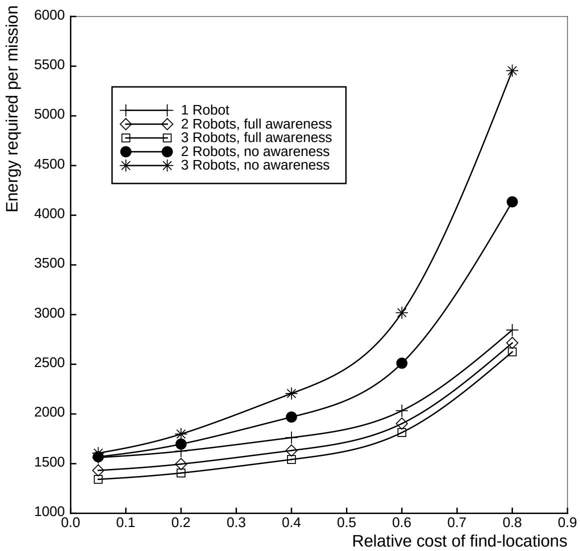

Figure 6: The average energy required to complete the mission as a function of the relative *find-locations* cost.

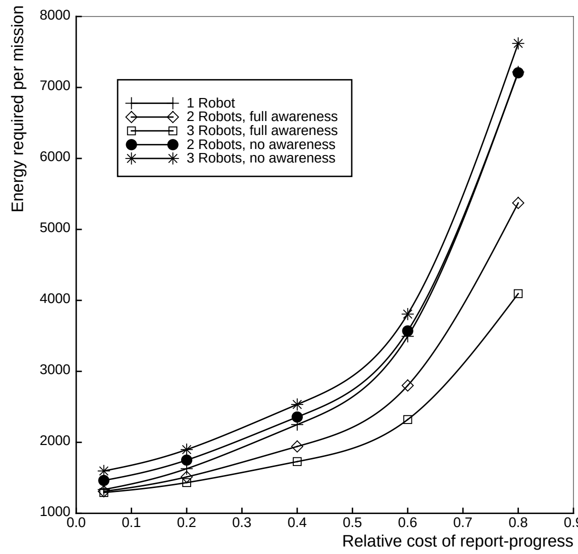

Figure 7: The average energy required to complete the mission as a function of the relative energy cost of the *report-progress* task.

5.2 Effects on Energy Usage

Figures 6 and 7 plot the effect of varying the relative costs of the two replicated tasks, *find-locations* and *report-progress*, on the average energy required to perform the mission for each of the experimental setups. Whereas the one robot experiment proved to be the worst performer in terms of time, we find here that the three robot team with no awareness performs the worst in terms of energy as the relative cost of either *find-locations* or *report-progress* increases. Clearly, the potential for wasting energy increases as the team size increases without awareness, since more robots are present to needlessly duplicate actions.

Also note that the team performance according to the energy metric improves with awareness, regardless of the team size because replication of actions is prevented. Note that for any level of team redundancy, the degree of energy improvement with awareness increases as the relative cost of the redundant action increases, since the energy saved with awareness is a direct function of the energy required to perform the redundant action.

In the case of the *find-locations* task, the energy performance of the two-robot team without awareness (version III) is better than the three-robot team without awareness (version IV), whereas for the *report-progress* task, the energy performances of these two teams were nearly identical. This again occurs because of differences in the extent to which the effect of the actions of other robots can be sensed through the world. In this situation, both teams lack awareness of the actions of other robots. The tasks that they might replicate due to this lack of awareness are the *find-locations* task and several instances of the *report-progress* task. Although neither of these task effects can be sensed through the world by this team, the *report-progress* task is closely tied to the *move-puck* task which *is* detectable through the world. Since the only time a robot tries to initiate the *report-progress* task is after it has completed the transport of a puck to the goal, and since there are a fixed number of pucks to be moved, the question of whether to report the progress only arises a fixed number of times for the team as a whole. Also, since the time required for one robot to find and move a puck is approximately the same as the time allowed between progress reports, robots in both versions III and IV are motivated to report their progress almost every time that they deliver a puck. Thus, regardless of whether the team consists of two or of three robots without awareness, the *report-progress* task is activated a fixed number of times, which means that the energy requirements remain the same.

On the other hand, the *find-locations* task, whose effects cannot be sensed through the world, is replicated by each robot on the team having the ability to perform that task, which means that as redundancy across robots increases, so does the energy usage. The lesson here, then, is that in the absence of robot awareness, redundancy across robots is detrimental for those redundant robot tasks whose effects cannot be sensed through the world.

Also note that in the case of the *report-progress* task, the energy performance of the three-robot team with awareness (version II) was better than the two-robot team with awareness (version I), whereas in the *find-locations* task, the energy performances of these two teams were nearly identical. This arises due to differences in the required number of instances of tasks and how well those instances can be distributed across the robot team. In this situation, both robot teams in question (versions I and II) have full awareness of the actions of other robots. The differences lie in the number of instances required by the puck moving mission of the *find-locations* task versus the *report-progress* task and how well they can be distributed across robots. While only a single instance of the *find-locations* task is required, many more instances of the *report-progress* task are necessary to complete the mission. Although each instance of these tasks can be distributed to any robot team member with the required capabilities, no single instance can be broken down into parts to be shared by more than one robot. Thus, in the case of the one required instance of the *find-locations* task, once one of the robots has selected that action, the other robots have to just wait patiently for the first robot to finish that action (of course, another robot may take over the task due to failure by the first robot, but that is another issue). In this case, an increased degree of redundancy across the team for this action does not provide any advantage, and so the performance of the two-robot team is not different from that of the three-robot team when both have awareness.

The benefit provided by the three-robot team over the two-robot team (both with full awareness) in the case of the *report-progress* task is obtained via a reduction in the required number of instances of the task. Since the three-robot team can complete the fixed amount of puck-moving required by the mission faster than the two-robot team, the time required to complete the mission is shortened. This is turn leads to a reduction in the number of progress reports required by the mission, which leads to less work for each robot to obtain the proper number of reports. Thus, the lesson learned here is that increased redundancy of robot capabilities in the presence of full robot awareness helps when the mission requires several instances of tasks that can be distributed across the robot team; it does not help, however, for single instances of tasks that cannot be shared.

5.3 Impact on Cooperative Team Design

Although numerous researchers have studied various approaches to cooperative control (e.g. [1, 2, 3, 4, 5, 8]), there is still a lack of adequate quantitative information that would allow a human designer to construct an appropriate multi-robot system that can achieve a specified level of performance when applied to a given set of tasks. This lack of design guidelines contributes to the difficulties in applying cooperative robotics to real-world applications. Thus, as a start to the development of such general design guidelines, we interpret some of our robot awareness results as they impact the design of cooperative robot teams.

In a typical real-world application, a human designer would be presented with a mission that must be performed and a metric of evaluation that should be used to judge the quality of the solution. The human designer's task would then be to develop a cooperative team of the proper size and mixture of capabilities to solve that mission, and to determine the appropriate use of explicit communication and/or implicit action recognition for that application. Of course, we must recognize that at the current level of technology, a human designer does not have simple routines for passive action recognition that can be plugged into robot team members to make them aware of the actions of their teammates. At present, the human designer can usually only provide this ability through the use of explicit communication among robots. Thus, the discussion below assumes that action recognition, if available, is provided through the use of explicit communication.

In developing an appropriate robotic team, the designer must recognize that several of the factors that have an impact on cooperative team performance are out of his or her control. Specifically:

- Since current sensory technology is still rather limited, the designer would typically not have control over the degree to which robots can sense the effect of other robots' actions through the world.
- Since the mission to be performed is specified in advance, the designer would have little or no control over the cost of replicated actions.
- Since the mission environment is identified in advance, the designer would likely have little control over factors such as radio interference or limited bandwidth that affect the reliability of an explicit communication mechanism available for the robot team's use.

While much work remains to be done to provide a designer with principles prescribing the optimal team composition for a given set of constraints, observations facilitating this design can be made from our studies of action recognition and robot awareness. Many of these observations were noted in the previous subsections. A couple of additional observations from our studies are as follows:

Increased heterogeneity can offset unreliable communication.

Our studies indicated that for robot actions whose effects cannot be sensed through the world, the lack of awareness causes an increase in the energy requirements of the mission. This increased energy requirement was shown to worsen as the level of robot redundancy increases and as the cost of the redundant actions increases. The implication, therefore, is that for applications in which the communication mechanism is unreliable relative to the reliability of the individual robot (which means that robot awareness will be low), and for missions in which a variety of tasks must be performed (such as the puck moving mission described in this paper), better energy performance can be achieved by making the robots more heterogeneous in their capabilities. The increased heterogeneity reduces the possibility of duplication of effort of a non-aware team, since robots will be more restricted in the tasks they are able to perform. However, the designer must carefully weigh the tradeoffs in the unreliability of the communication with the potential unreliability of the robots. One would not want to increase team heterogeneity to the point that one robot failure dooms the entire mission to failure.

Redundancy can be more important than awareness.

Another result of our studies indicated that a team without awareness may be able to perform a mission more quickly than a team with a lower level of redundancy if a significant proportion of the mission consists of tasks whose effects can be sensed through the world. The implication of this for robot team design is that a fully homogeneous team of sufficiently large size has no need for explicit communication as long as the effects of robot actions can be sensed through the world. This result is clearly consistent with the large body of swarm robotics research that is studying emergent cooperation in teams with no explicit communication (e.g. [3]). Such approaches, of course, rely on the ability of robots to detect the side-effects of other robots' actions to achieve success. Thus, for these types of missions, the designer can achieve better results through the use of increased team size than through the incorporation of explicit communication.

6 Conclusions

In this paper, we have examined the effect of robot awareness of team member actions on the performance of actual robot teams performing a puck moving mission. Although we studied these issues in the context of the puck moving mission, we believe our results can be generalized to other cooperative robotics domains. The loss of awareness leads robots to select their actions based purely on feedback from the world through their remaining sensors (other than the radio communication sensor) and on their own internal motivations and priorities. Since robots cannot always detect the effects of the actions of their teammates through the world, the lack of awareness can lead to the redundant execution of certain tasks required by the mission. We began our study of the extent of this effect on mission performance as a functions of the degree of redundancy in robot capabilities and the performance metric of interest (either time or energy). Our studies revealed that additional factors influence the effect of awareness on team performance, including the ability of the robots to detect the actions of other robots through the world, the amount of work available per robot, and the cost of the redundant tasks. We analyzed the effects of these factors on team performance and discussed their impact on cooperative team design.

One of our longer-term goals is to derive general principles that can serve as prescriptive guidelines for engineers in designing cooperative teams for given applications. By analyzing the characteristics specific to a given mission, such as the likelihood of communication failure, the importance of time or energy efficiency, and the relative costs of replicated actions, a design engineer can use such principles to help construct a robot team with the proper degree of redundancy and mixture of capabilities to effectively execute its mission. Our future work includes examining this issue of awareness from a theoretical basis to determine if it can be accurately addressed and solved as an optimization problem, and thus provide additional useful design principles and assist in evaluating the tradeoffs of the numerous factors influencing robot team performance.

Acknowledgements

This research was funded in part by the Advanced Research Projects Agency of the Department of Defense under Office of Naval Research contract N00014-91-J-4038 at the Massachusetts Institute of Technology's Artificial Intelligence Laboratory, and in part by the Office of Engineering Research, Basic Energy Sciences (directed by Dr. Oscar Manley), of the U.S. Department of Energy, under contract No. DE-AC05-84OR21400 with Martin Marietta Energy Systems, Inc.

References

[1] H. Asama, K. Ozaki, A. Matsumoto, Y. Ishida, and I. Endo. Development of task assignment system using communication for multiple autonomous robots. *Journal of Robotics and Mechatronics*, 4(2):122–127, 1992.

[2] Tucker Balch and Ronald C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):1–25, 1994.

[3] J. Deneubourg, S. Goss, G. Sandini, F. Ferrari, and P. Dario. Self-organizing collection and transport of objects in unpredictable environments. In *Japan-U.S.A. Symposium on Flexible Automation*, pages 1093–1098, 1990.

[4] Marcus Huber and Edmund Durfee. Observational uncertainty in plan recognition among interacting robots. In *Proceedings of the 1993 IJCAI Workshop on Dynamically Interacting Robots*, pages 68–75, 1993.

[5] C. Ronald Kube and Hong Zhang. Collective robotic intelligence. In *Proceedings of the Second International Workshop on Simulation of Adaptive Behavior*, pages 460–468, 1992.

[6] Bruce MacLennan. Synthetic ethology: An approach to the study of communication. In *Proceedings of the 2nd interdisciplinary workshop on synthesis and simulation of living systems*, pages 631–658, 1991.

[7] Lynne E. Parker. ALLIANCE: An architecture for fault tolerant, cooperative control of heterogeneous mobile robots. In *Proceedings of the 1994 IEEE/RSJ/GI International Conference on Intelligent Robots and Systems (IROS '94)*, pages 776–783, Munich, Germany, September 1994.

[8] Lynne E. Parker. *Heterogeneous Multi-Robot Cooperation.* PhD thesis, Massachusetts Institute of Technology, Artificial Intelligence Laboratory, Cambridge, MA, February 1994. MIT-AI-TR 1465 (1994).

[9] Gregory M. Werner and Michael G. Dyer. Evolution of communication in artificial organisms. In *Proceedings of the 2nd interdisciplinary workshop on synthesis and simulation of living systems*, pages 659–687, 1991.