Using Communication to Reduce Locality in Distributed Multi-Agent Learning

Maja J Matarić
Computer Science Department
University of Southern California
941 West 37th Place, SAL 228 MC 0781
Los Angeles, CA 90089-0781
mataric@cs.usc.edu

Abstract

This paper attempts to bridge the fields of machine learning, robotics, and distributed AI. It discusses the use of communication in reducing the undesirable effects of locality in fully distributed multi-agent systems with multiple agents/robots learning in parallel while interacting with each other. Two key problems, hidden state and credit assignment, are addressed by applying local undirected broadcast communication in a dual role: as sensing and as reinforcement. The methodology is demonstrated on two multi-robot learning experiments. The first describes learning a tightly-coupled coordination task with two robots, the second a loosely-coupled task with four robots learning social rules. Communication is used to 1) share sensory data to overcome hidden state and 2) share reinforcement to overcome the credit assignment problem between the agents and bridge the gap between local/individual and global/group payoff.

1 Introduction

This paper attempts to bridge the fields of machine learning, robotics, and distributed AI. It describes how simple communication methods can be applied to enable and speed up learning in complex, noisy, situated multi-agent systems. The domains in question do not permit reliance on complex communication protocols but can be made significantly more manageable through the use of minimalist communication schemes which require no additional alteration of the learning framework. We describe two such schemes and apply them to reduce the undesirable effects of locality in fully distributed multi-agent systems with multiple concurrent learning agents.

The paper focuses on using simple communication methods to deal with two key problems: 1) hidden state and 2) credit assignment. The hidden state problem arises because situated agents typically cannot sense all of the relevant information necessary for completing the task and learning to perform it efficiently. The credit assignment problem arises because reinforcement in a distributed system is often provided at a global level, and must somehow be divided over multiple agents whose impact differs and varies over time. Both problems can be addressed by using simple strategies that apply communication in a dual role: as sensing and as reinforcement, in each case through local undirected broadcast. In both cases, communication is treated in a manner consistent with other sensory inputs, and requires no additional overhead. We demonstrate the methodology on two multi-robot learning experiments. The first describes two robots learning a tightly-coupled coordination task (box-pushing), using communication for sharing sensory data, to overcome hidden state, and reinforcement data, to overcome the credit assignment problem between the two agents. The second experiment describes a loosely-coupled task with four robots learning social rules (yielding and sharing information), using communication to share social reinforcement in order to bridge the gap between global/group and local/individual payoff.

In both cases, the role of communication is to, locally in space and time, increase the scope of impact of a single agent. Thus, communication serves to effectively cluster agents for a period of time during which they are tightly interacting. This has the effect of making the system temporarily and locally less distributed, and consequently alleviates the hidden state and credit assignment problems inherent in distributed multi-agent learning.

The paper is organized as follows. Section 2 defines some key terms concerning communication and cooperation. Section 3 addresses the use of simple communication as a means of sensing, thus enlarging the agent's perceptual space and its capabilities for cooperation and learning. Section 4 focuses on the use of communication as reinforcement for addressing the credit assignment problem by reducing locality of the agents' actions. Section 5 describes the first of the two robot experiments and demonstrates the use of communication to enable two robots to learn a tightly-coupled cooperation task. Section 6 describes the second robot experiment, in which four roots learn the social rules for yielding and sharing information. Section 7 discussed the results, Section 8 overviews some related work, and Section 9 concludes the paper.

2 Communication and Cooperation Definitions

Communication is the most common means of interaction among intelligent agents. Any observable behavior and its consequences can be interpreted as a form of communication. For purposes of clarity, we introduce some definitions for the scope of this paper and our work.

Direct communication is a purely communicative act, one with the sole purpose of transmitting information, such as a speech act, or a transmission of a radio message. Even more specifically, *directed communication* is direct communication aimed at a particular receiver. Directed communication can be one-to-one or one-to-many, but in both cases the receivers are identified.

In contrast, *indirect communication* is based on the observed behavior, not communication, of other agents, and its effects on the environment. This type of communication is referred to as *stigmergic* in biological literature, where it refers to communication based on modifications of the environment rather than direct message passing.

Both direct and indirect communication are practiced by most species in nature. For example, bees use signals, such as the waggle dance, with the sole purpose of transmitting information and recruiting. In contrast, they also use cues, such as the direction of their flight, which transmit hive information as a by-product of their other behaviors (Seeley 1989).

Cooperation is a form of interaction, usually based on communication. Certain types of cooperative behavior depend on directed communication. Specifically, any cooperative behaviors that require negotiation between agents depend on directed communication in order to assign particular tasks to the participants.

Analogously to communication, *explicit cooperation* is defined as a set of interactions which involve exchanging information or performing actions in order to benefit another agent. In contrast, *implicit cooperation* consists of actions that are a part of the agent's own goal-achieving behavior repertoire, but have effects in the world that help other agents achieve their goals.

In order to study the role of communication in a controlled fashion and to explore how much communication is needed for the group behaviors we are studying, we adopted a minimalist approach. No directed, one-to-one communication between the agents was used in any of the experiments. Indirect communication was based on sensing the external state of neighboring agents, as well as sensing their density, and the effects of their actions. Direct communication was undirected, and limited to local broadcast: agents could transmit messages that could be received by others near by. However, agents did not have the ability to choose the receivers of their messages, and thus to engage in directed communication.

We will show that this simple form of communication is in itself a powerful mechanism for aiding various aspects of learning in distributed multi-agent system.

3 Communication as Sensing

The lack of accurate and reliable sensors is arguably the most common complaint of researchers in situated agent control and learning. In robotics in particular, sensors have been targeted as one of the limiting factors in the way of progress toward more complex autonomous behavior. Most of the commonly used sensors provide noisy data and are difficult to accurately characterize and model, presenting a major challenge for real-time robot learning (Matarić 1996). Yet, from the multi-agent perspective, the ability to sense and correctly distinguish members of one's group from others, from obstacles, and from various features in the the environment is crucial for most tasks. While the inability to make such distinctions does not preclude symbiotic relationships (Axelrod 1984), the lack of sophisticated perceptual discrimination is a critical limitation in multi-robot work. Non-visual sensors such as infra-red, contact sensors, and sonars are all of limited use in the social recognition task. Vision typically offers best discrimination capabilities, but much of traditional image analysis involves computational overhead that is prohibitive for a collection of moving, dynamically interacting robots. Some effective low-overhead approaches to robot vision have been implemented on single robots (e.g., Horswill (1993)) but have not yet been scaled up to groups.

The inability to obtain sufficient sensory information to properly discriminate results in perceptual aliasing or the hidden state problem (Whitehead & Ballard 1990). Due to sensory limitations, multiple world states are perceived as the same input state, inducing serious problems for learning in any domain. These problems are particularly acute in multi-agent/robot domains, where the behavior of the system is the result of the interactions between agents, each of whose limitations can thus become amplified.

If viewed as a form of sensing, communication in multi-agent systems can be used to effectively deal with the hidden state problem. Like other sensors, radio receivers and microphones perceive signals (i.e., messages being communicated) and pass those on for further processing. Specific properties of sensors vary greatly, and these differences can be usefully exploited: some information that is very difficult to directly sense and visually discriminate can be easily communicated, a tradeoff that is frequently unitized by biological systems (Wehner 1987). We will demonstrate that agents that locally broadcast elements of state that may otherwise be difficult or impossible to access, have a significant advantage in a learning scenario. However, it is important to keep in mind that communication suffers from similar noise and inaccuracy properties as other sensors; messages may be corrupted, dropped, or received out of order. These features may not play a key role in disembodied multi-agent work, but impact multi-robot learning strongly, as discussed below.

4 Communication as Reinforcement

A key challenge of distributed multi-agent systems is to achieve group-level coherence. Such coherence is best guaranteed by top-down control but often requires prohibitive computational overhead that scales poorly with increased group size. A central controller must maintain updated information about the entire group, perform optimizations over the global state space, and send commands to the group; the necessary information is typically not available, the computation cannot be completed in real-time, and the communication imposes a significant bottleneck (Matarić 1995). Completely distributed alternatives in which all agents learn independently and concurrently scale much better, but in turn introduce other difficulties: non-stationary environments and the credit assignment problem.

A group of agents learning in parallel creates a non-stationary world: as the agents learn, their behavior changes, resulting in inconsistencies. Furthermore, multi-agent systems also face the credit assignment problem both at the level of the individual and at the level of the group. At the individual level, the interaction with other agents often delays an agent's payoff, aggravating the temporal credit assignment problem. At the group level, local individual behavior must be appropriately associated with global outcomes. We describe the use of communication as reinforcement, a strategy that enables the agents to locally share reward in order to overcome the credit assignment problem at both levels. In effect, *locally sharing reinforcement between agents decreases locality in the learning system.*

In summary, we have argued that communication can be effectively treated as a part of the input state. It is complementary to direct sensing, can be similarly abstracted, and has similar noise properties, thus effectively augmenting the agent's sensory space. Furthermore, communication can then be used as a powerful source of reinforcement, which can appropriately bias and shape the agent's reward function. The rest of this paper illustrates such use of communication in two examples of multi-robot learning. The first addresses the use of communication to extend local sensing of each of the robots, impacting both the hidden state and the credit assignment problem. The second addresses the use of communication to share reinforcement, impacting credit assignment at the individual and group levels. In both cases, communication extends individual sensing, action, and reinforcement by binding two or more agents together locally, for a short amount of time, while their behavior is immediately related. This simple idea has a powerful effect on distributed learning.

5 Communication for Shared Sensing

This section describes the use of communication to handle hidden state and credit assignment in an experiment pairing two mobile robots concurrently learning to perform a box-pushing task. The achievement of this task requires tightly-coupled coordination because the box is chosen to be large enough so that neither of the agents can push it to the goal alone. The

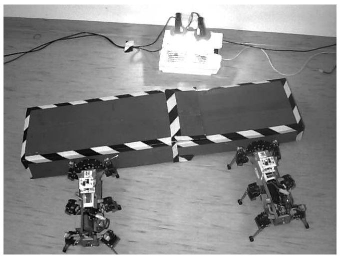

Figure 1: Two Genghis-II six-legged robots used to learn cooperative box-pushing, equipped with whiskers for detecting contact with the box, pyroelectric sensors for detecting the light at the goal, and radios for communication.

agents, six-legged mobile robots, are equipped with a radio communication mechanism, whiskers that detect contact with the box, and an array of five pyroelectric sensors that detect the direction of the goal marked with a bright light (Figure 1). The sensors do not provide information about the robots' absolute location or relative location along the box. Furthermore, the available sensory information is noisy and detailed sensory characteristics are not provided.

The goal of the learning experiment was to have the robots automatically discover a strategy that allows them to coordinate their efforts into a coherent policy capable of delivering the box to the goal regardless of the initial positions of the robots or the box and the robots' inability to directly sense each other or the movement of the box. The experiment was set up within a reinforcement learning (RL) framework, so that each of the robots was learning a reactive mapping between its sensory perceptions and its repertoire of pre-programmed fixed-duration behaviors based on reinforcement received over time (Kaelbling 1990).

The sensory data was based on the inputs from the whiskers and the pyroelectric (light) sensors. Each of the two whiskers returned a one-bit contact/no-contact value whenever probed. The light sensors consisted of a curved array of five units centered in front of the robot; the position of the maximum brightness (max-pyro) was returned whenever the sensor was probed. The behavior repertoire consisted of the following:

- *find-box*: move forward until one of the whiskers detects contact, then turn in that direction until both whiskers detect contact;

- *push-forward*: move forward while maintaining contact with both whiskers;
- *push-left*: turn to the left while pushing; accomplished by taking smaller steps with the left legs;
- *push-right*: turn to the right while pushing; accomplished by taking smaller steps with the right legs;
- *send-msg*: broadcast the max-pyro value over the radio;
- *stop*.

The reinforcement learning algorithm summed the reward and punishment values over time. The values were computed from the changes in the direction to the goal. Moving so that the sensors were facing the light or turning toward the more intense light generated internal reward, while turning away from the light generated internal punishment. Reinforcement had identical magnitude and opposite sign for the reward and punishment signals. Reinforcement was computed after the completion of a behavior, and was used to reinforce the most recent perception-behavior pair. The system maintained and updated a table of perception-behavior values, and used those to dynamically compute the policy.

The algorithm we employed used no discounting so the system was learning a purely reactive pushing policy. Upon completion of a behavior and computation of the reinforcement, another behavior was activated. The behavior selection algorithm chose the "best" behavior from the table 75% of the time and a random behavior 25% of the time.

Merely servoing to the light without a box is simple. Communication was the key to making this task learnable for the two robots having to do so with the large box and incomplete sensing: they could distinguish their direction relative to the light, but not the orientation of the box or the other agent. Through the use of communication, the robots could combine their sensory data and get information about two points along the box. As a result, the hidden state problem could in this case be solved by having the two agents pool their sensory resources. In general, pooling the sensory resources of all of the agents on a task is equivalent to a centralized control approach. Fortunately, most task do not require this level of communication as they are decomposable into concurrent subtasks. This two-robot experiment demonstrates two tightly-coupled sub-tasks that require communication to overcome sensory deficiencies. Alternative solutions could employ robots with more powerful sensors (Donald, Jennings & Rus 1993), or a larger group of robots with looser communication and information-sharing requirements, as demonstrated in our second robot experiment, in the next section.

Another challenge of this task is finding the joint robot actions that result in the appropriate coordinated state-action mapping for the task. The multi-agent credit assignment problem arises: if a box is getting closer to the goal, whose actions are responsible for the progress? Which of the agents should be rewarded? Is one agent undoing the work of the other? One way of solving this problem is again to use communication: each agent tells the other what action to perform, the two take coordinated actions, observe the outcomes, and *share* the subsequent reward or punishment. As a side-effect, communication also synchronizes their actions, ensuring coordinated movement. Note that this does not guarantee that the agents will always move in synchrony, nor that they will begin and end their behaviors at exactly the same time: such accuracy is beyond the capabilities of these robots' noisy communication and control mechanisms. The synchronization provided by the communication mechanism is only a coarse signal, but it provides sufficient information to initiate or terminate behaviors. The coarseness of the signaling also implies that the exact duration of behavior execution may vary, but this also had no effect on our system's performance. These points are important in that they illustrate how very sparse and minimal communication can be effectively employed as an augmentation of the basic sensing and behavior system.

In what we described so far, the robots were learning a policy between perceptual inputs and motor outputs. However, our learning system was more complex. In our experiments the robots learned both what behaviors to execute and what to communicate to each other. The pre-programmed *send-msg* behavior enabled them to communicate their sensory information (i.e., they were not learning language) in order to overcome the hidden state problem. In addition to this built-in capability, the robots also learned how to choose, based on the pooled sensory data, what action to take (described above) and what action to communicate to the other robot. Thus, each robot learned a functional mapping between the combined perceptual states and the best behaviors for itself (my-action) and the other agent (its-action). The resulting policy table consisted of the following: $S_{combined} \rightarrow A_{my-action} \times A_{its-action}$. In our case, $|S_{combined}| = 25$ and $|A_{my-action} \times A_{its-action}| = 25$. The desired policy was learned by both robots in over 85% of the trials; on the average it was learned in 7.3 minutes (i.e., about 40 contiguous turns). The details of the experimental trials are described in Simsarian & Matarić (1995).

In order to make the learning converge, the robots took turns in controlling the box. An alternative solution would establish a master-slave system where only one of the robots is learning and commanding the actions of the other. Given the intrinsic differences in the sensor and effector characteristics of the two robots, we opted for a scheme that did not favor either. The coarse turn-taking, easily accomplished through the use of the communication channel, effectively alleviated the credit assignment problem between the robots. After the behaviors (my and its) were executed, each robot could compute the reinforcement directly from the updated combined sensory data. Since the robots shared identical internal reinforcement functions, they received identical reinforcement, modulo sensory and communication noise. However, the strategies learned by the two robots were not identical; they depended on the side of the box each robot happened to be on. Each learned a "handed" strategy, becoming either a left-pusher or a right-pusher. (This specialization could be further generalized through the use of a single variable representing the box side.)

In this learning task, communication was effective in extending the individual agent's very local and incomplete view of the world at least in part because it comes from the other, near-by agent's sensors and actions involved in solving a shared task. Our next experiment, involving a larger group of robots, demonstrates that this idea of using communication to tie together spatially local agents also has important uses in learning in larger, more loosely-coupled groups with less frequent and completely asynchronous communication and behavior.

6 Communication for Shared Reward

This section describes the use of communication to handle hidden state and credit assignment in an experiment featuring four mobile robots concurrently learning four non-greedy social rules dealing with yielding to each other sharing information about the location of pucks they are all collecting. The robots are equipped with infra-red sensors for obstacle and puck detection, bump sensors, position sensors, and radios for communication (Figure 2). The sensors are noisy and error prone, and do not provide information about other robots' external state or behavior.

As in the previous experiment, the robots were pre-programmed with simple behaviors that enabled them to collect pucks (*pickup, drop, go-home, wander, follow, avoid*) and send messages (*send-msg*). Again as in the first experiment, the form of the messages was to broadcast small amounts of information, but in this case the communication was done asynchronously and the content of the messages was quite different, as will be described below.

The goal of the experiment was to have the robots improve individual puck collecting behavior by learning when to yield to one another and when to proceed, as well as when to send a message sharing information about the location of pucks and when to follow the received message to the location it communicated. As before, this experiment was set up within a reinforcement learning (RL) framework, so that each of the robots was learning a functional mapping between its sensory perceptions and its behavior repertoire. A table of situation-behavior pairs was maintained by each robot, and the behavior selection algorithm chose the "best" behavior 60% of the time, a random behavior 10% of the time, and *follow-msg* behavior 30% of the time. This behavior induced the robot to follow the most recently received message.

The sensors on this group of robots included infra-red, bump, position, and communication, and their inputs were clustered into a set of predicates: *have-puck, at-home, near-pucks, near-another, too-near-another*, and *got-msg* indicating whether the robot had a puck in its gripper, whether it was within the home region (the area where pucks were deposited), near a source of pucks, within a reasonable proximity or very near another robot. Not all behaviors were used in the learning; the following subset enabling the social rules was a part of the learning system: *stop*, *send-msg*, *proceed* and *follow-msg*. The perceptual conditions for the rest of the behaviors were fixed, based on previous learning experiments (Matarić 1996). Consequently, the robots did not need to learn how to collect pucks, but they did need to learn how to do so in a more social and efficient fashion.

Learning social rules in this scenario is difficult, due to delayed reinforcement and the credit assignment problem between the agents. In a reinforce-

Figure 2: Four IS Robotics R2e robots used to learn to yield (e.g., top two robots) and communicate the location of the pucks (e.g., bottom two robots), equipped with infra-red sensors for obstacle and puck detection, bump sensors, position sensors, and radios for communication.

ment learning paradigm, each agent tries to maximize its own payoff over time. Social behaviors like communicating puck location do not increase the agent's payoff and behaviors like yielding may decrease it by delaying or even preventing reward. Thus, it is difficult for individual agents to learn such social rules without having them centrally imposed or pre-programmed, as advocated in some game-theoretic work (Axelrod 1984).

Our experiments demonstrated that through the use of simple communication between spatially local agents, the described social rules can be learned. The communication channel was used to provide payoff, i.e., to generate reinforcement, for observing the behavior of the near-by agent and mimicking it. This way of relating two near-by agents' behavior was an effective means of learning social rules. Observing other agents' behavior is a difficult perceptual task, but it can be easily solved with communication; our robots broadcast, within a limited area (i.e., the perceptual range) the behavior they were performing, as well as the reinforcement they were receiving. Note that the former could potentially be sensed with a richer sensory apparatus, such as a camera and clever processing, while the latter typically could not. Thus one message of this work is that communication not only replaces but also augment sensors with information that in some cases cannot be sensed at all and yet could be trivially communicated.

Communicating external state (i.e., behavior) and internal state (i.e., reinforcement) between locally interacting agents has a dual effect. First, this information allows a robot to select what observed behavior to "imitate" based on the vicarious reinforcement, thus decreasing the need for experimentation. The key assumption is that positive reinforcement will be consistent across agents, and this assumption holds because of the spatial and temporal locality of the shared information. When two robots are at nearly the same place at the same time, they are likely to receive the same reinforcement for the same actions. Second, in addition to encouraging social behaviors, communicating and sharing reinforcement also provides payoff for social actions (such as yielding) that would not otherwise benefit the agent directly. In our implementation, when a robot yields to another, the "proceeder" reaches its goal and receives reinforcement directly, then *shares* it with the "yielder", thus reinforcing social behavior. The sharing is performed by communicating the reward. Any near-by robots that observed the yielding behavior and received payoff have the impetus to imitate it, and thus also learn it. Learning to send messages works in the identical fashion. A robot that finds pucks communicates its location, and the near-by agents that perceive the information move to that location, find the pucks, and receive and share the reward. Thus the mapping between finding pucks and sending a message is reinforced, as is the mapping between receiving a message and following it.

Locally sharing information was sufficient to enable the group to learn the following social behaviors: 1) when close to other robots they yielded, 2) when they found food they broadcast that information, 3) when they received a message with a location, and were not executing another behavior, they followed the message to the location, and 4) they proceeded when being yielded to. The desired social policy of yielding and sharing information was learned by all robots in over 90% of the trials; on the average it was learned in 20 to 25 minutes. Mutual yielding symmetries were broken by the noise inherent in the robots' sensing and actuation. Obeying the social rules resulted in better overall performance of the group compared to a non-social individually-greedy system. This result was based on the comparisons with the measured total collection time for a fixed number of pucks by each group in multiple trials.

As in the first experiment, in this learning task communication was used to extend each agent's very local and incomplete view of the world. Unlike the first experiment however, where sensory information was shared and reinforcement was individual, in this case sensory information was individual but reinforcement was shared. The described method is another demonstration of addressing the multi-agent credit assignment problem though the use of communication. Because the communication method is local, it is independent of group size and thus scales well. Broadcasting behavior and reinforcement were simple to implement and proved to be robust in the presence of noise given a minimal amount of radio filtering. Corrupted messages were largely detected and ignored. Lost messages were rare since each was broadcast a fixed number of times. If any were lost, however, the omission had the effect of slowing down the learning, as does any sensor noise. Given the confined work area, the robots encountered each other with high frequency, thus being exposed to social conditions that facilitated the learning.

7 Discussion

In the noisy, asynchronous, and dynamic worlds of multi-robot and other multi-agent systems, simple forms of communication can help address some of the important stumbling blocks in the way of system adaptability. The work we have described utilizes some simple built-in communication strategies to compensate for sensory limitations as well as exploit message-passing as a broader form of sensing and reinforcement.

Robot learning is known to be a very difficult problem (Mahadevan & Kaelbling 1996), and multi-robot learning an even more challenging instance of it that combines issues from learning, robots, and multi-agent systems. We applied some simplifying strategies to make the learning problem manageable, largely at the individual robot level. These include using behaviors as higher-level descriptions of action in both experiments, and the use of combinations of sensory values into higher level predicates in the second experiment. The combination of the two effectively collapses the potentially prohibitively large exponential search space of the learning system. The pros and cons of this approach are described at length in Matarić (1996), but are not relevant to the key message of this paper, which focuses on effective use of communication.

In the first experiment, what is communicated is pre-programmed but the resulting control policy is learned. Because of the sensory and effector noise inherent to physical systems and in particular to experimental robots, this strategy allows for the learning system, initially helped with the communication capability, to discover and adapt the control policy over time. While a controller for two coordinated box-pushing robots may not be impossible to program, the demonstrated learning time presents a viable and certainly a more adaptive alternative.

In the second experiment both what is communicated and the control policy are learned. This is a better argument for the scaling properties of local broadcast communication. As was argued earlier, learning social rules can be prohibitively time consuming, but may also be an unintuitive programming task in a distributed system. The presented approach biases the learning system toward a smaller state space and then allows it to learn the rules autonomously. It is likely that in system of sufficient sensing and control complexity, such as mobile robots, the built-in bias is crucial in order to enable the system to learn at a higher level, including group behavior and social rules. The few multi-robot learning demonstrations performed to date have tended in that direction, as discussed in the next section.

8 Related Work

The described work applies concepts from machine learning (ML) and distributed AI (DAI) to robotics. The problems of hidden state and credit assignment have been dealt with extensively in the ML literature (McCallum 1996, Sutton 1992). Whitehead (1991) analyzed cooperative RL mechanisms and Littman (1994) used the Markov games framework for simulated RL soccer-playing agents. Tan (1993) applied RL to a simulated hunter-prey scenario in which the agents used communication to exchange instantaneous information, much like in our experiments, but also used episodic and learned knowledge.

Communication and multi-agent cooperation have been addressed in abstract and applied multi-agent work (Yanco & Stein 1993, Dudek, Jenkin, Milios & Wilkes 1993, Altenburg & Pavicic 1993). In DAI learning work, several related contributions have been made in the area of studying multiple simulated reinforcement learners (see collection by Weiss & Sen (1996)), but have not yet been applied to robotics, where multi-robot learning was identified as an area to be addressed (Mahadevan & Kaelbling 1996).

In robot learning, Mahadevan & Connell (1991) focused on state generalization and controller modularity applied to a sonar-based robot learning to push a box in an experiment complementary our our first one, which addresses the same problem with a similar approach, but introduces communication to extend it to the multi-robot domain. Work on cooperative robot pushing has been addressed in control; Donald et al. (1993) demonstrated a (non-learning) two-robot sofa-moving system using more powerful sensors to eliminate the need for communication. In contrast, our work uses communication to compensate for the robots' sensory limitations and to facilitate learning. Our previous work (Matarić 1994) treated learning in a multi-robot system, and introduced the use of behaviors and their associated conditions as a means for state space reduction which was applied in the work described here. The previous work focused on using internal progress estimators to address the temporal credit assignment problem on each agent individually, without the use of communication. In contrast, Parker (1993) addressed the issue of coordinating a loosely-coupled group of foraging robots, and described how communication can be used to combine the abilities of heterogeneous agents in order to complete a cooperative task more efficiently. Work by Asada, Uchibe & Hosoda (1995), Uchibe, Asada & Hosoda (1996) and related papers address the case of interacting and competing agents, and apply related behavior-based reinforcement learning ideas to learning aspects of soccer in a multi-robot environment. Murciano & del R. Millan (1997) address the homogeneous foraging problem in simulation, by learning to activate some behaviors through the use of limited communication between agents, and reinforcement based on time to complete the task by the entire group. This approach is similar to ours, but requires a collective reinforcement signal, which is not possible in all distributed multi-agent domains, and aggravates both the temporal and spatial credit assignment problem.

9 Conclusions

The goal of this paper was to draw from ideas in machine learning, robotics, and DAI in order to deal with two problems of learning in distributed multi-agent systems: hidden state, because agents cannot sense all of the relevant information, and credit assignment, because it is difficult to distribute reinforcement over multiple concurrently acting agents. We described how simple communication based on local broadcast can be used to address both problems and illustrated the ideas on two multi-robot demonstrations. In the first, two robots learn to cooperatively push a box by communicating sensory readings to each other; in the second, four robots learn social rules for yielding and sending information by using communication to share reinforcement. In both cases, communication is treated much the same as a sensory inputs, thus requiring no no additional processing overhead nor a departure from the reinforcement learning framework.

The principle effect of the described limited broadcast communication schemes was to temporarily bind together two or more spatially local agents. The idea utilizes simple communication to temporarily decrease the distributedness of the system just between the robots that are near-by at a given point in time. We demonstrated it in two very different domains: a tightly-coupled task that pooled sensory information and coarsely bound the agents with a coarse common action synchronizer, and a loosely-coupled task where local exchange of internal state information enabled a distributed group of robots to learn social rules in spite of the temporal and spatial credit assignment problem.

The simple idea of local broadcast scales well and is particularly suitable for noisy, dynamic domains in which agents interact based on incomplete sensory information. Such environments abound in robotics, where communication may compensate for the limitations of more direct sensory modalities, thus enabling learning in multi-robot systems.

9.0.1 Acknowledgments

The described work was supported by the Office of Naval Research Grant N00014-95-1-0759 and the NSF Infrastructure Grant CDA-9512448. The box-pushing learning work was performed at the Swedish Institute for Computer Science, in collaboration with Kristian Simsarian and Martin Nilsson. The author thanks François Michaud for the suggestion to write the paper, as well as for detailed and insightful discussions during its preparation, and the three anonymous reviewers for their constructive comments.

References

Altenburg, K. & Pavicic, M. (1993), Initial Results in the Use of Inter-Robot Communication for a Multiple, Mobile Robotic System, *in* 'Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots', Chambery, France, pp. 96–100.

Asada, M., Uchibe, E. & Hosoda, K. (1995), Agents That Learn from Other Competitive Agents, *in* 'Proceedings, Machine Learning Conference Workshop on Agents That Learn From Other Agents'.

Axelrod, R. (1984), *The Evolution of Cooperation*, Basic Books, NY.

Donald, B. R., Jennings, J. & Rus, D. (1993), Information Invariants for Cooperating Autonomous Mobile Robots, *in* 'Proc. International Symposium on Robotics Research', Hidden Valley, PA.

Dudek, G., Jenkin, M., Milios, E. & Wilkes, D. (1993), On the utility of multi-agent autonomous robot systems, *in* 'Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots', Chambery, France, pp. 101–108.

Horswill, I. D. (1993), Specialization of Perceptual Processes, PhD thesis, MIT.

Kaelbling, L. P. (1990), Learning in Embedded Systems, PhD thesis, Stanford University.

Littman, M. L. (1994), Markov games as a framework for multi-agent reinforcement learning, *in* W. W. Cohen & H. Hirsh, eds, 'Proceedings of the Eleventh International Conference on Machine Learning (ML-94)', Morgan Kauffman Publishers, Inc., New Brunswick, NJ, pp. 157–163.

Mahadevan, S. & Connell, J. (1991), Automatic Programming of Behavior-Based Robots Using Reinforcement Learning, *in* 'Proceedings, AAAI-91', Pittsburgh, PA, pp. 8–14.

Mahadevan, S. & Kaelbling, L. P. (1996), 'The National Science Foundation Workshop on Reinforcement Learning', *AI Magazine* **17**(4), 89–97.

Matarić, M. J. (1994), Reward Functions for Accelerated Learning, *in* W. W. Cohen & H. Hirsh, eds, 'Proceedings of the Eleventh International Conference on Machine Learning (ML-94)', Morgan Kauffman Publishers, Inc., New Brunswick, NJ, pp. 181–189.

Matarić, M. J. (1995), 'Issues and Approaches In the Design of Collective Autonomous Agents', *Robotics and Autonomous Systems* **16**(2–4), 321–331.

Matarić, M. J. (1996), 'Reinforcement Learning in the Multi-Robot Domain', *Autonomous Robots* **4**(1), 73–83.

McCallum, A. R. (1996), Learning to use selective attention and short-term memory in sequential tasks, *in* P. Maes, M. Matarić, J.-A. Meyer, J. Pollack & S. Wilson, eds, 'From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior', The MIT Press, pp. 315–324.

Murciano, A. & del R. Millan, J. (1997), 'Learning Signaling Behaviors and Specialization in Cooperative Agents', *Adaptive Behavior* **5**(1), 5–28.

Parker, L. E. (1993), Learning in Cooperative Robot Teams, *in* 'Proceedings, IJCAI-93 Workshop on Dynamically Interacting Robots', Chambery, France, pp. 12–23.

Seeley, T. D. (1989), 'The Honey Bee Colony as a Superorganism', *American Scientist* **77**, 546–553.

Simsarian, K. T. & Matarić, M. J. (1995), Learning to Cooperate Using Two Six–Legged Mobile Robots, *in* 'Proceedings, Third European Workshop of Learning Robots', Heraklion, Crete, Greece.

Sutton, R. S. (1992), Machine Learning, Special Issue on Reinforcement Learning, Kluwer Academic Publishers, Boston.

Tan, M. (1993), Multi-Agent Reinforcement Learning: Independent vs. Cooperative Agents, *in* 'Proceedings, Tenth International Conference on Machine Learning', Amherst, MA, pp. 330–337.

Uchibe, E., Asada, M. & Hosoda, K. (1996), Strategy Classification in the Multi-Agent Enviornment – Applying Reinforcement Learing to Soccer Agents, *in* 'Proceedings, ICMASS Workshop on RoboCup: Soccer as a Problem for Mulit-Agent Systems'.

Wehner, R. (1987), 'Matched Filters – Neural Models of the External World', *Journal of Computational Physiology* **A**(161), 511–531.

Weiss, G. & Sen, S. (1996), Adaptation and Learning in Multi–Agent Systems, Lecture Notes in Artificial Intelligence, Vol. 1042, Springer-Verlag.

Whitehead, S. D. (1991), A complexity analysis of cooperative mechanisms in reinforcement learning, *in* 'Proceedings, AAAI-91', Pittsburgh, PA.

Whitehead, S. D. & Ballard, D. H. (1990), Active Perception and Reinforcement Learning, *in* 'Proceedings, Seventh International Conference on Machine Learning', Austin, Texas.

Yanco, H. & Stein, L. A. (1993), An Adaptive Communication Protocol for Cooperating Mobile Robots, *in* J.-A. Meyer, H. Roitblat & S. Wilson, eds, 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', The MIT Press, pp. 478–485.