Exploiting Embodiment in Multi-Robot Teams

Barry Brian Werger and Maja J Matarić
Interaction Laboratory, Robotics Research Labs
Department of Computer Science
University of Southern California
Los Angeles, CA 90089-0781
barry@usc.edu, mataric@usc.edu

Abstract

This paper describes multi-robot experiments and systems which exploit their embodied nature to reduce needs for sensory, effector, and computational resources. Many natural phenomena, such as territorial markings or ant pheromone trails, take great advantage of the ability mark the environment, and of the natural dissipative processes which cause these markings to decay. In this way, globally complex behavior can result from simple local rules. The information invariants ([Donald 1995], [Donald et al 1994]) literature raises the issue of robots similarly recording information, or even "programs," into the physical environment. This paper provides example systems that dynamically encode information and "programs" into the physical environment, and by so doing, increase their own robustness and reduce their resource requirements and computational complexity. The main experimental system that we present, a robot "chain" used for foraging, is modeled after the natural phenomenon of ant pheromone trail formation. "Minimal" agents with local sensing and action form a system that can perform position-dependent tasks. We discuss how this system can dynamically adapt to environmental changes, both by forming efficient paths to changing resource locations and by dynamically assuming roles. We also demonstrate a robot soccer system that exhibits such dynamic role assumption and flexible teamwork, subject to global constraints, using only limited local sensing and no explicit communication. We discuss how moving information and computation into the shared physical environment improves our ability to generate complex global behaviors from simple locally interacting agents.

1 Introduction and Motivation

While current trends in robotics towards situated, embodied, multiple agents have provided many systems that react effectively and robustly to their environments, they have dealt only obliquely with the possibility of deliberate manipulation of the environment by the agents. Systems that implement behaviors such as aggregation, dispersion, and flocking [Matarić 1995] involve agents which, through their physical presence, affect the behavior of other agents in a manner that is more than mere "interference." [Beckers et al 1994] describe a task where physical effects of task performance allow a simple, local control strategy to produce a consistent global behavior. Work in behavioral economics and "robot ecology" (e.g., [McFarland 1994], [Steels 1994]) has investigated the influence agents have on each other through the use and production of shared, limited resources.

We have been inspired by the elegant simplicity of natural forms of direct environmental modification such as territorial markings and pheromone trails. These phenomena effectively exploit the benefits of agents which deliberately encode information into the physical environment. As discussed in [Aron et al 1990], [Goss et al 1989], [Muller and Wehner 1988], and [Hölldobler 1990], the release of pheromones by ants leads to trails that can be differentiated by pheromone "strength," a function of frequency of use and decay. If pheromones are released only during certain phases of a task (e.g., while carrying some item back to the nest), then trails can begin to form *efficient* paths to useful locations, such as rich supply areas. This, combined with a very simple control strategy of probabilistically choosing the most frequently used path, leads to group behavior that adjusts to follow dynamically determined shortest paths to dynamically changing useful destinations.

The ability to take advantage of information "encoded" into the physical environment through task mechanics has recently been under investigation from the perspective of information invariants ([Donald 1995], [Donald et al 1994]), a line of research which seeks to examine the interaction between sensing, computation, communication, and task mechanics in the performance of distributed manipulation tasks. This approach has provided some theoretical basis for comparing sensori-computational systems, and some steps towards a methodology for design of efficient distributed manipulation systems. Specifically, a number of systems are demonstrated which take advantage of physical effects of task dynamics to dramatically reduce requirements for sensing, computation, and communication. A methodology for minimalizing such requirements is proposed. However, work on this approach "is still biased towards sensing, and it remains to develop a framework that treats action and sensing on an equal footing" [Donald et al 1994].

Two questions raised by this research include: 1) the ability of agents to externalize, or encode, "state" into the physical environment, and 2) the ability to do the same with "programs." We believe that the ant pheromone trails discussed above can be viewed as "state," and possibly even as "programs" physically encoded into the environment, and that a similar system can be employed by robots to create distributed physical representations - or even distributed physical "programs" - in their environment. In this paper, we present such a system of autonomous mobile robots that modifies its environment in way that allows dynamically changing, globally position-dependent tasks to be performed through local physical contact and very simple control rules, and a second system which uses local environmental interaction for globally-constrained role allocation.

This paper is organized as follows. Section 2 presents an overview of our robot chaining system for ant-like foraging. Section 3 discusses some of the resource requirements of previous approaches to foraging. Sections 4 presents the high-level robot chaining algorithm, with Section 5 filling in the implementation details. Section 6 presents techniques that allow the chain to adapt to changing environmental conditions by means analogous to those of ant trail formation. Section 7 discusses two other systems which take advantage of externalization - a sweeping chain which can make coverage guarantees using only local sensing, and a soccer-playing system which exhibits effective dynamic role allocation with interesting global properties. Sections 8 and 9 provide discussion and conclusions.

2 An Ant-like Robotic System

Previous research inspired by insect behavior has aimed to reproduce particular instances of *stigmergy* - "the production of a certain behavior in agents as a consequence of the effects produced in the local environment by previous behavior" [Beckers et al 1994] (see also [Deneubourg et al 1991], [Theraulaz et al 1991]). While purely stigmergic solutions have been found for tasks such as clustering items in the environment ([Holland and Melhuish 2000],[Beckers et al 1994]) and even sorting of scattered heterogeneous items into homogeneous clusters ([Holland and Melhuish 2000],[Deneubourg et al 1991]), tasks which require particular behaviors to take place at specific locations have generally relied upon some type of global position sensing, globally visible beacons, or random encounter of some locally-sensible position marker. One of our most direct motivations was a system described in [Deneubourg et al 1990] and [Goss and Deneubourg 1991], in which simulated robots form chains, based on beacons carried by each robot and perceivable within a limited range.

The robot chaining system for foraging that we present replaces the chemical pheromones of the ant trails with the physical bodies of robots. We demonstrate that a group of robots equipped with only contact sensors is able to form a physical pathway that some members of the group can use to randomly explore the environment and efficiently return home. Since the chain is formed of active robots, the effects of decay and accumulation of pheromone strength can be replaced by addition or subtraction of chain-forming robots and shifting of direction of the chain in response to activity of robots traveling along it. By these means, robot chains are able to form shortest paths to rich deposits much as pheromone trails do.

3 The Foraging Task

Variations of *foraging* - collecting items from the environment and depositing them at a specific location - are examples of a common class of robotic tasks that require some knowledge of global positioning for efficient performance.

Next we present an overview of the most commonly used sensory modalities and strategies for performing variations of the foraging task, and some of their associated requirements and overhead.

3.1 Methods useful for single or multiple agents

The Omniscient Planner: The use of a planner that can "see" the whole environment and the forager's position within it, and plan accordingly. This is infeasible for non-trivial environments and group sizes.

Position/Orientation Sensing: The use of absolute global position information. There are various ways to perform position and orientation sensing that can be considered to be effectively equivalent. Popular approaches include:

> *Global Positioning System (GPS) and Compass:* requires environmental preparation (the GPS), and a potentially sophisticated local sensor (the compass) that is typically very sensitive to environmental noise.

Radio-Sonar Positioning System: triangulation based on time differences between arrival of sonar and radio signals provides position information. This is the basis of several successful foraging systems ([Fontán and Matarić 1996], [Goldberg and Matarić 1996], [Matarić 1995]), but requires preparation of the environment (radio-sonar broadcasters at precise locations), complicated sensing equipment, and triangulation computation.

Dead-Reckoning: Determination of robot position and orientation through careful monitoring of actuator motion, such as wheel rotation. Does not require modification of the environment, but does require that initial location be known. This approach necessitates accurate and potentially complicated actuator motion sensing and calibration, and suffers from cumulative error.

Taxis: Following of some type of beacon. This method involves some modification of the environment (the beacon), and is limited by the range of visibility.

Recognition of Unique Locations: The use of local environmental features, through such means as vision or sonar, to identify certain locations (landmarks) to which the agent can orient itself. While this approach has been used successfully in various experiments (e.g., [Matarić 1992b], [Horswill 1993], [Gomi 1995]) it relies on having or acquiring some topological map representation of the environment, and sensing the landmarks sufficiently accurately to localize within that map.

3.2 Methods specific to multiple agents

Chemical Deposit (Pheromones): Requires the ability to emit and detect the presence of varying concentrations of pheromones. Some work towards robotic odor sensing/depositing systems has been done by [Russell 1995] and others; [Grasso et. al. 1996] have demonstrated systems which are able to locate sources of odors.

Beacon Chains: The same as taxis, except that the beacon does not have to be globally perceivable; instead, robots are equipped with beacons that can be left within visible range of each other, together forming chains of indefinite length. The approach requires the ability to distinguish between beacons, which must broadcast information regarding their distance (in units of beacons) from the home location ([Deneubourg et al 1990], [Goss and Deneubourg 1991]).

Contact Chains: Only simple local sensors (such as infrared or contact) are used in the process of chain formation and following. Agents follow a chain composed of the "bodies" of other agents towards the home location. Below we propose an extreme case in which the robots use a small number of the simplest, most reliable sensors available - contact sensors.

Figure 1: A robot returns to the chain carrying a puck after a circular excursion.

4 Robot Chains

The system we present involves the formation of chains by a group of robots in order to provide local information sufficient for the performance of globally position-dependent tasks. The chain maintains contact with a starting point (*Home*). Robots that are not currently part of the chain are able to follow the chain both away from *Home* and back towards it. The chains can adjust to link *Home* with other points, such as rich supply areas, and re-form when the supply diminishes or new deposits are discovered, and, potentially, be put into motion to completely sweep an area. Simple communication can be sent up and down the chain, allowing a wide range of fairly complex behaviors to emerge.

Our approach to chaining uses only physical contact sensors (simple microswitches). Each member (link) of the chain maintains periodic contact with the links ahead and behind, by touch. Limited communication is implemented through the same mechanisms to allow for chain maintenance. Most communication between members of the chain is *phatic*, intended only to assert the existence of the line of communication (i.e., the integrity of the robot chain). This is implemented as a "double tap." One robot begins the communicative act (Figure 2) by moving enough to tap the robot ahead or behind twice and returning to its (approximate) initial position. The tapped robot answers by tapping back twice and returning. Two taps are used to distinguish communication from the many random taps of other robots in the environment.

More informative communication can be performed similarly, with contact held for a fixed period, or taps added, at points B, D, and F of Figure 2. Many interesting behaviors require no more than just this simple 1-bit phatic communication, but it is possible to pass more elaborate messages through combinations of "short" and "long" taps (as in Morse code) or through the use of an electrical signal transmitted through the contact point.

The basic behaviors involved in chain formation and maintenance are:

HomeLink: Remains still, except to maintain communication with the "next" (farther from *Home*) link in the chain.

MiddleLink: Maintains communication with the "next" and "previous" links by returning taps and passing messages.

EndOfChain: Maintains communication with the previous link, assists in positioning of new links by interacting in the alignment process, and establishes communication with newly aligned links to pass on the status of *EndOfChain* before becoming a *MiddleLink*. If not useful for a given period of time (see Section 6.2), it communicates its intention to the "previous" link to transfer *EndOfChain* status, and leaves the chain.

JoinChain: Allows a robot to attempt to append itself to the end of the chain. Interacts with the *EndOfChain* robot to align properly at the end of the chain, establish communication with the current *EndOfChain*, and become the new *EndOfChain*.

The behavior involved in chain following is:

FollowChain: Follow along the chain by "tacking" to the left (the tacking process is described in Section 5.1.2). When rounding the end of the chain, attempt to *JoinChain* only if not carrying an item. The end of the chain can be determined by a longer than usual time between contacts during tacking.

The foraging task our chain-making system performs involves the collection of metal pucks scattered either randomly or in clusters around the test environment. Searching is done by the following behavior:

ExcursionSearch: Follow the chain, occasionally taking roughly circular (clockwise) journeys of random radii into the area next to the chain (Figure 4).

5 Implementation

We have implemented a foraging system which moves metal pucks, distributed around an area, to the *Home* location using only physical contact-level sensing. The system is designed for the foraging team to begin in the *Home* area and commence chain construction from there. Infra-red emitter/detector sensors with an effective range of less than 1 inch, located on the underside of the robots, are used to determine when the robots are at *Home*, which is marked with a non-reflective black area on the floor. These extremely short-ranged sensors can be replaced with physical sensors capable of detecting some property of *Home*, or of the *HomeLink*. Currently the

Figure 2: Communication passed down the chain. A) The chain in resting state. B) Robot 3 taps robot 2 twice to initiate the communication act C) 3 returns to normal position. D) 2 taps 3 twice to acknowledge communication. E) 2 returns to normal position, terminating communication act. F) 2 taps 1, initiating next communication act in the passing of the message down the chain. In non-phatic communication, stages B, D, and E are modified.

robots are powered up sequentially at appropriate times. In the future this will be done either by fixed timing based on unique ID numbers or, ideally, through messages passed back through the chain to team members waiting at *Home*.

At the time of these experiments, our herd of 20 robots was undergoing renovations such that only 6 (described below in 5.2) were capable of the chain-based foraging task; of these, two were only capable of chain following. The length of our chains was thus limited to four, though we occasionally increased it by switching "dead" robots for chain links closer to *Home* (which are the least active) in order to re-use functional robots further down the chain. The current system assumes that the environment contains only robots, pucks, and *Home*.

A QuickTime movie of robots forming chains and performing the *ExcursionSearch* is available on the World Wide Web at *http://www-robotics.usc.edu/~barry/chaining.mov*.

5.1 Behaviors

5.1.1 Initial Chain Location - Skirt

Robots start gathered at *Home*. A behavior **Skirt** navigates to the edge of *Home*, then tacks (see below) along this edge until it encounters a physical obstacle project-

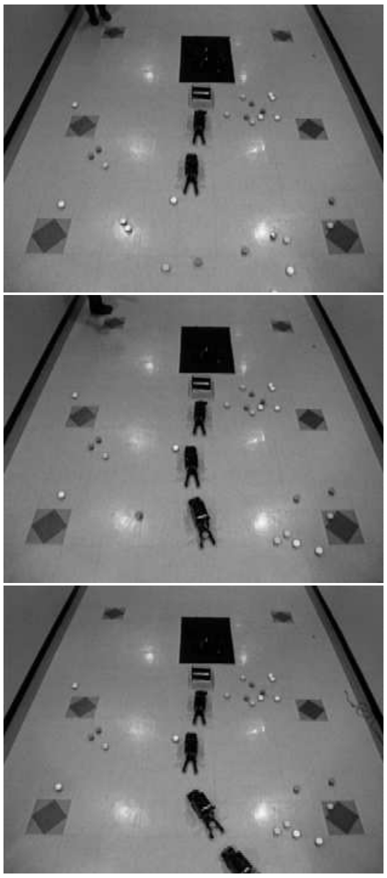

Figure 3: Four robots form a chain from *Home*

ing outside of the *Home* region. This obstacle is assumed to be the chain.

5.1.2 Chain Following - Tack

Chain following is performed through simple tacking. The following robot angles slightly towards the chain (that is, towards the robot's left) until contact is made, backs away from the chain at a sharper angle, then resumes forward motion angling slightly towards the chain in order to make contact again, further along the chain. This tacking allows a following robot to round the end of the chain and continue down the other side. In current experiments we enforce directionality on chain traffic. Tacking is always done with the chain on the following robots' left side; thus the right side of the chain (when viewed from *Home*) is for outbound traffic, and the left side for inbound traffic. This simplifies control and also minimizes potential inteference among the robots.

5.1.3 Extending the Chain - JoinChain

JoinChain is implemented as a combination of three behaviors: an extended **Tack**, **BackInto**, and **AlignBack**.

The extended **Tack** times the intervals between contacts with the chain. If the more than a given time passes (in our experiments, 10 seconds), it outputs a signal. This signal is used to deactivate **Tack** and activate **BackInto** when the robot is not carrying a puck.

BackInto reverses at a sharper angle than that used in forward tacking. When combined with an appropriate (empirically determined) time-out for **Tack**, this gives the robot a likelihood of contacting the *EndOfChain* robot with its back bumper. If contact is not made within a certain time period (30 seconds in our experiments), it is assumed that the end of the chain has been missed and the robot continues forward at the tacking angle until something is contacted. If contact is made, the robot withdraws enough to clear contact and outputs a signal which deactivates **BackInto** and activates **AlignBack**.

AlignBack pauses to avoid confusion between the first contact tap of **BackInto** and its own communicative taps, then taps the *EndOfChain* robot twice (within three seconds), adjusting its angle relative to the *EndOfChain* according to contact indication through left, right, or both rear contact switches. If the *EndOfChain* responds with two taps (within three seconds), the robot is fairly well aligned and considers itself to have joined the chain. In doing so, it deactivates, and takes the role of *MiddleOfChain*. If it does not receive two answering taps within a fixed interval (10 seconds), it circles left at the tacking angle until it hits something, then deactivates **AlignBack** and activates **Tack**.

5.1.4 MiddleOfChain - Link

MiddleOfChain has been implemented as the behavior **Link**, which detects and responds to double taps to its front and rear. This corresponds to A-F of Figure 2, and is also enough to satisfy the requirements of the *EndOfChain*. Time-outs on tap attempts to the next chain link, and single contacts with the previous link, allow recovery from most errors.

5.1.5 Foraging Excursions - ExcursionSearch

ExcursionSearch has been implemented through an extended **Tack** and a behavior **CircleRight**. The extended **Tack** makes a decision every time it contacts the chain as to whether it should make a circular excursion to the right to search for pucks. No excursions are made if the robot is already holding a puck, otherwise the choice is random (12.5 % chance in our experiments).

5.2 The Robot Herd

Our experiments are implemented and tested on the Nerd Herd, the Interaction Lab's group of 20 IS Robotics

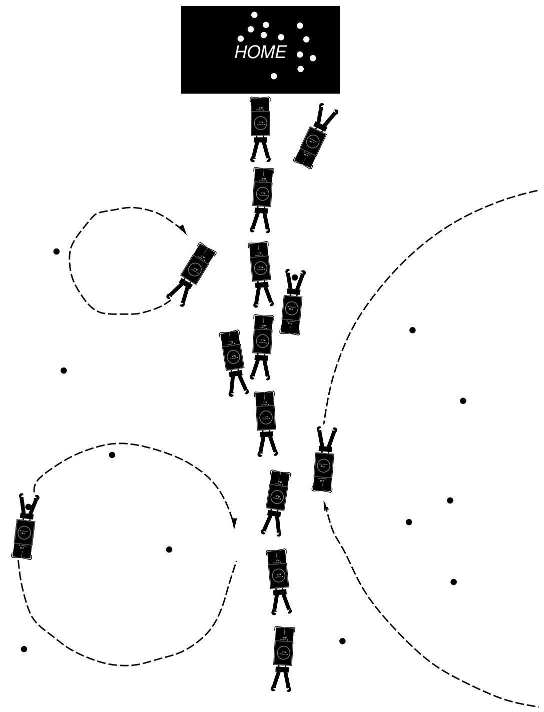

Figure 4: Excursion Search strategy: Robots search for pucks and return to the chain by making roughly circular "excursions" from the chain.

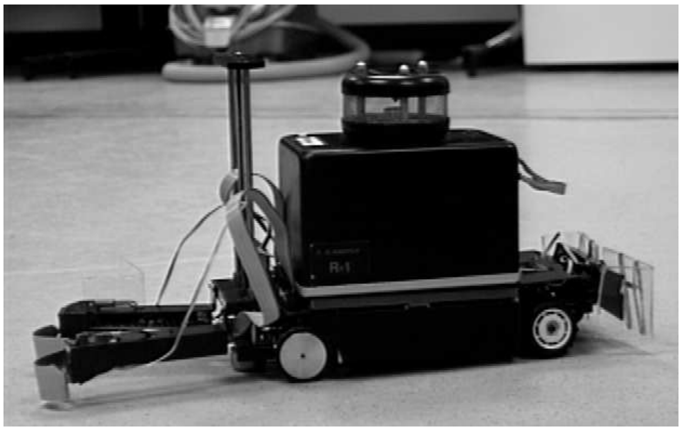

Figure 5: A Nerd Herd Robot: Each of the Nerd Herd robots is a 12"-long four-wheeled base equipped with a two-pronged forklift for picking up, carrying, and stacking pucks, and with a radio transmitter and receiver for inter-robot communication and data collection.

R1 mobile robots. Each member of the Nerd Herd is a 12-inch long four-wheeled vehicle, equipped with a two-pronged forklift for picking up, carrying, and stacking pucks (Figure 5). The forklift contains two contact switches, one on each tip of the fork, six infra-red sensors: two pointing forward and used for detecting objects and aligning onto pucks, two break-beam sensors for detecting a puck within the "jaw" and "throat" of the forklift, and two down-pointing sensors for aligning the fork over a stack of pucks for stacking (Figure 6). The pucks are special-purpose light ferrous foam-filled disks, 1.5 inches in diameter and between 1.5 and 2.0 inches in height. They are sized to fit into the unactuated fork and be held by the fork electro-magnet. Each robot also has one piezo-electric bump sensor on each side of the chassis. Only the front contact, the stacking IRs, and rear contact sensors described in Section 5.2.1 below are used in the described experiments.

The mechanical, communication, and sensory capabilities of the robots allow for exploration of the environment, robot detection, and finding, picking up, and carrying pucks. These basic abilities are used to construct various experiments in which the robots are run autonomously, with all of the processing and power on board. The processing is performed by a collection of four Motorola 68HC11 microprocessors on each robot. Two of the processors are dedicated to handling radio communication, one is used by the operating system, and one is used as the "brain" of the robot, for executing the down-loaded control system used in the experiments. The control systems are programmed in the Behavior Language, a parallel programming language based on the Subsumption Architecture [Brooks 1986, Brooks 1990].

5.2.1 Hardware Modifications

Originally equipped with piezo-electric bump sensors on the back of the chassis, some of the R1 robots were modified to better suit the chaining task. The modifications resulted in the rear surfaces of some robots having large bumpers that activated contact switches (see Figure 6). This is necessary due to the nature of the original bump sensors, which cannot indicate continuous contact, and to the fact that the width of the original rear surface is the same as the width of the opening of the fork - which leads to constant catching and damaging of the fork-mounted contact sensors in the alignment task.

5.2.2 Hardware Limitations

As discussed in Section 1, properties of physical hardware impose restrictions not only on the control strategies that can be applied, but also on the types of tasks and experiments that can be implemented. Robot hardware is constrained by various sensory, mechanical, and computational limitations.

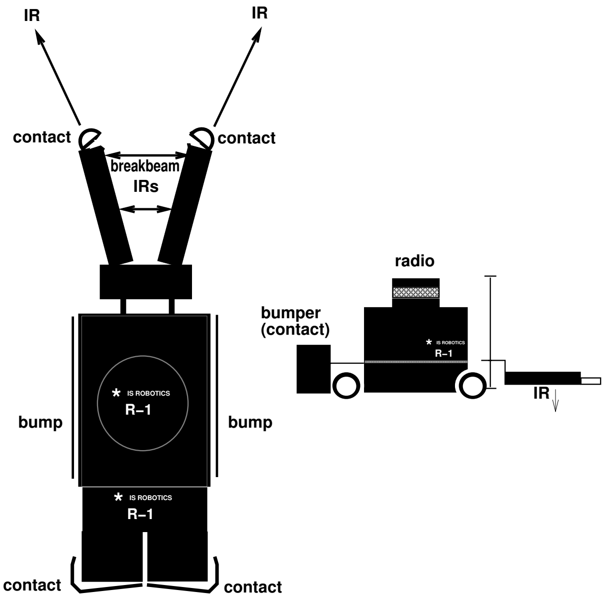

Figure 6: R1 Sensors: Each of the Nerd Herd R1 robots is equipped with contact sensors at the ends of the fork, piezo-electric bump sensors on each side and two on the rear of the chassis, and six infra-red sensors on the fork. Two forward-pointing IRs are located at the ends of the forks, two break-beam IRs in the jaw and throat of the fork, and two down-pointing IR for stacking pucks in the middle of each of the fork arms. The result of replacing the rear piezo-electric bump sensors with bumpers and contact sensors is shown.

The robots' mechanical steering system, when in perfect condition, is "accurate" to within 30 rotational degrees. At certain steering angles, a drive wheel is lifted off the ground, while at others, the steering wheels jam against metal parts of the chassis. During any type of physical interaction, parts tend to change alignment.

The uncertainty and variability inherent in any work with physical robots, and especially salient in the case of the R1s, although frustrating, is beneficial to experimental validity. Hardware variability between robots is necessarily reflected in their group behavior. Even when programmed with identical software, the robots behave differently due to their varied sensory and actuator properties. Small differences among individuals become amplified as many robots interact over extended time. As in nature, individual variability creates a demand for more robust and adaptive behavior. The variance in mechanics and the resulting behavior have provided stringent tests for our methodologies.

5.3 Performance

The foraging system that we tested with six working R1 robots demonstrates practicability of our robot chain concept. While some behaviors demonstrated a high failure rate, graceful recovery allowed multiple attempts, as detailed below. Ongoing software and hardware refinements are providing consistent increases in reliability, especially in regards to previously ubiquitous mechanical failures.

The ability of robots to follow the formed chains was robust, and was lost only when mechanical failures led to following robots pushing chain robots so far as to open up wide gaps in the chain. Any gap wide enough to permit the front contact sensors of a following robot to cross the chain (about 1 robot length, 12 inches, depending on the turning circle of the particular robot involved) tended to result in unrecoverable errors. The average separation in well-formed chains was observed to be about six inches, and the nature of the communication along the chain tended to maintain this distance through minor (though not major) "pushing" by following robots. This can be improved with slightly more sophisticated chain maintenance behavior, and/or with the ability of chain link robots to anchor themselves against pushing. The effective length of a chain can be said to be approximately 1.5 times the length of the robots that form it (when using the R1s).

The only problems encountered during *Excursion Searching*, besides mechanical ones described above, occurred when a robot pushed more than two pucks at a time, which prevented the front sensors from making any contact. While ultimately best solved by better sensor placement, this problem was resolved in some trials through a time-out which makes the robot back up after a given period without contact (we actually used **Tack** and **BackInto** from *JoinChain* quite successfully). Our limited number of robots only allowed us to have one robot searching for pucks while the others formed the chain; in this case, the searcher brought an average of one puck *Home* each trip around a chain of four robots, which took about two minutes, depending on the number of circular excursions. With more functional robots we could examine the effects of interference along the chain and its influence on scalability.

The *JoinChain* process requires the most precision and was most prone to failure. Approximately fifty percent of attempts made a successful first contact (in **BackInto**), and of these approximately fifty percent exchanged taps and resulted in joining the chain. These rates could be improved by tuning the steering systems of the robots and/or tuning the timing of individual robots, but improvement would be only temporary since alignment changes rapidly. Though a raw success rate of twenty five percent does not seem impressive, graceful recovery and persistence of attempt allowed eventual joining in most cases.

As seen in Section 6.2, the probabilistic nature of chain joining is actually a benefit in the process of optimizing the length of the chain. This is exactly the type of trade-off we intend: that a large number of less capable, somewhat expendable agents can perform certain tasks at least as efficiently as a smaller number of more sophisticated agents. As in natural systems, such as ant pheromone trail formation, global behavior is a result of the cumulative effects of many actions. The key point we see in both natural (i.e., ant) and artificial (i.e., robotic) systems is that while individual successes benefit the system as a whole, individual failures do not accumulate. The most efficient ant paths are more frequently traveled than the longer ones, and are thus given a stronger marking that overpowers, and out-survives, the weaker ones. Analogously, in robot chains, only those robots that successfully join the chain have a lasting effect on the behavior of others. In both systems, success results in a persistent encoding of information in the environment, while failure does not.

6 Adaptive Chains

Once a chain has been formed, it can adapt to the environment in order to optimize system functioning.

Ants use pheromones to encode information into the environment about their current activities. When returning to the nest with useful material, an ant leaves a trail that can be followed back to the source of the material. As groups of ants forage, the trails accumulate and decay, so that little-used paths tend to fade while oft-used paths grow stronger. Since shorter paths lead to more frequent trips from a source to the nest, the shortest paths tend to be the strongest, and therefore most attractive, ones.

We can achieve a similar effect in our robot chains without the need to leave or sense markers, since the links of the chain are capable of computation and motion. Rather than depositing pheromones and having paths "emerge" through purely physical processes, the chain links can collect some statistics of the activity of the chain-following robots, and use them to adapt to the environment by physically modifying the chain. Two types of chain modification are sufficient for generating an optimal path to a rich source in a plane with no insurmountable obstacle: 1) shifting of chain direction, and 2) lengthening/shortening of the chain. Both are based on a simple report, analogous to a pheromone deposit: when a chain-following robot returns to the chain from an excursion having picked something up, it taps a message indicating this to the first chain link it encounters. We call this message a *SuccessReport*.

[Gordon 1999] describes findings about how ants change roles (e.g., from foragers to internal nest workers). This is found to happen in response to the number of encounters each ant has with ants fulfilling *other* roles - a nest worker that encounters a number of successful foragers in a given time period will decide to forage. As seen below, the process we describe for adjusting the length of the chain functions in a very similar manner.

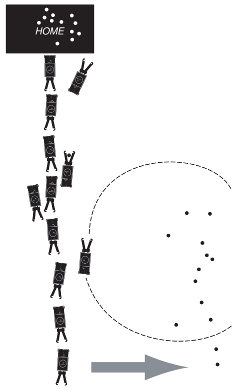

Figure 7: Moving Toward a Rich Deposit: After a robot picks up an item, it communicates the fact of the pickup to the first link robot it contacts when returning to the chain. The links collect statistics that allow them to determine whether most pickups are coming from the left or the right. Using this information, the links can begin shifting the chain towards rich deposits.

6.1 Chain Direction

In order for the chain to move to intercept a rich source, all that is necessary is for the chain links to monitor how many times they have had *SuccessReport*s on their right and left sides. If basic behaviors are in place that maintain chain integrity, individual robots can shift towards the direction of more *SuccessReport*s (within constraints of chain integrity) without need for explicit communication with neighboring links. In this way, the entire chain will slowly shift towards a rich source; indeed, over time, the chain should efficiently connect multiple sources, if they are present. Figure 7 illustrates a situation in which the chain shifts towards a source.

In order to more clearly replicate the ant systems, and eliminate the risk of the chain infinitely extending in a direction with no sources, it would be necessary to introduce random direction-shifting of chain links with some probability. Decay of trails could be replicated in two ways: either the links could factor recency into their statistics, or, more minimally, the links could merely react by shifting towards the direction of *every SuccessReport* , allowing such temporally-based statistics to to be computed "physically."

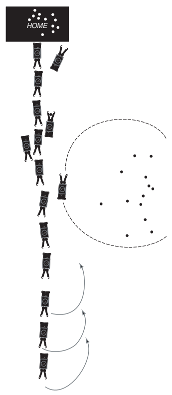

Figure 8: Optimizing Chain Length: Using statistics of where along the chain robots return after successful excursions, the chain length can be optimized so that foraging robots get to, and return from, rich deposits along a shortest path. When robots near the *End-Of-Chain* robot record significantly fewer successful returns than "earlier" links, the *End-Of-Chain* robot switches roles to become a chain-follower, decreasing the chain length by one. This is repeated until most successful returns are close to the end of the chain.

6.2 Chain Length

Ideally, once the chain has shifted to intersect a rich source, we would like it to end there - that is, we would like the *EndOfChain* to be near the center of the richest area, so that robots can return directly from the source to *Home*. Figure 8 demonstrates a situation where the chain should be shortened in order to both optimize the pathway and allow more robots to participate in transport of material.

There are two ways for this to happen; in either case, the chain will tend to shorten to the optimal length when there is a rich deposit, and naturally begin to grow again if this source begins to be exhausted. One way is for the chain links to collect *SuccessReport* statistics (most likely, the number of recent *SuccessReport*s at each link, for comparison) and pass them along the chain through some protocol, allowing the *EndOfChain* to decide when it should leave the chain and become a forager (by passing *EndOfChain* status to the preceding link).

Another minimal, environmentally-oriented way to adjust the chain length is to simply have the *EndOfChain* leave the chain after a period of time. If the chain extends past a rich source, there will be fewer robots attempting to append themselves to the end of the chain (since many will be carrying material and thus be ineligible); if the chain does not reach a source, few if any robots will be carrying and thus most will attempt to append themselves and lengthen the chain. This can be seen as dynamic role fulfillment such as [Gordon 1999] finds in ant colonies: when the *EndOfChain* encounters mostly successful foragers (which do not attempt to append themselves to the chain), it is likely to leave the chain and become a forager. When the foragers encounter mostly chain links without finding useful material, they tend to become chain links. The robots, like the ants, fulfill roles as determined by global constraints.

7 Other Tasks

Techniques of group coordination through the environment can be applied to many tasks performed by physical robots. Here we present two further examples: 1) a proposed extension to robot chaining for thorough coverage of an area, and 2) a successful robot soccer system that incorporates offensive and defensive formations which are entered into by an appropriate number of robots as dictated by circumstances, without explicit communication. Both of these systems, like the chaining system, are able to structure global system organization through local interaction. The area coverage system can make guarantees about coverage of an area without localization, and the soccer system has the effect of selecting and executing globally-appropriate roles without negotiation or explicit role selection.

7.1 Chains in Motion

Robot chains need not serve only as "passive" pathways for other robots; rather, they can structure motion of a group of robots. Specifically, a chain of robots that remain within perceptual range of each other can thoroughly sweep an area, for applications such as de-mining, planetary exploration, or search-and-rescue. As shown in Figure 9, a chain of robots "anchored" at one end can sweep a circular area with some guarantee of coverage: either the robots remain in sensor range, providing a maximum space between concentric circles of the sweep, or the system can report that the sweep was not complete.

The sweeping chain can be of great benefit in situations where coverage is essential but localization is difficult (such as surf-zone demining), or where hetereogeneous resources are combined. A single robot with a global positioning system can serve as the "anchor" for a sweeping chain, and allow for rough localization of any robot in the system (using the chain as a polar coordinate system relative to the ancor). Thus large areas can be thoroughly covered, and items of interest localized, with a single, non-moving GPS receiver. After a complete sweep, the anchor can move to a new location for a new sweep.

In situations where items being sought must be further investigated or manipulated (e.g., sensed in more detail or retrieved), it is possible to combine the concepts of

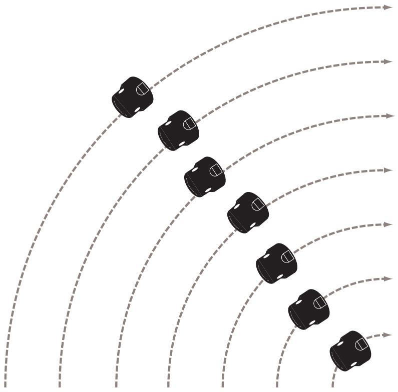

Figure 9: Sweeping an Area: Rather than growing outward from a *Home*, a robot chain can be "anchored" at a given point and sweep a circular area around it for thorough coverage. A chain of simple robots can be combined with a more capable anchor robot (e.g., equipped with GPS) for efficient and accurate patterned searches.

a sweeping chain and a chain pathway, as follows. A small number of specialized robots wait at the anchor point as the sweep progresses. If any robot in the chain encounters an object that requires further attention, it sends a report down the chain. The chain pauses and one of the specialized robots follows the chain to the notifying robot (Figure 10), at which point it can perform its specialized task. This allows heterogeneous systems to redirect resources effectively without need for either localization or global communication.

7.2 Dynamic Formations

As we have discussed in Section 6, ants are able to determine, through their local interactions with other ants, what roles would be globally optimal for them to assume (e.g., forager vs. domestic nest worker), and the chaining robots are able to use similar means of assuming optimal roles. Here we discuss a system we have implemented for a drastically different domain - robotic soccer - which is also able to use local interactions to determine globally optimal roles.

[Werger 1999] discusses at length our minimalist approach to team cooperation for a robot soccer team. Though individual players can perceive only the ball, the goals, and obstacles (which are not distinguished but may be walls, opponents, or teammates), and have no communication equipment, the team displays sophisticated cooperative behavior. The team falls into appropriate formations for offensive and defensive situations with the interesting property of formation size limitation.

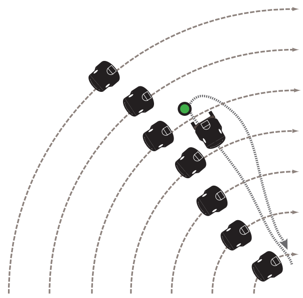

Figure 10: Following the Sweeping Chain: If objects encountered during a sweep require the attention of specialized equipment (e.g., more detailled sensing, or retrieval), the chain can pause the sweep upon encounter of such objects, allowing an appropriate robot to follow along the chain to the object of interest and back to the anchor area.

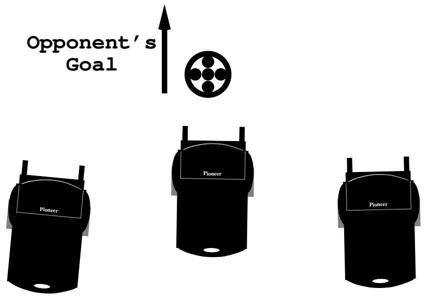

Figure 11: Offensive Formation in Robot Soccer: Interaction of *Push*, *Disperse*, and *Safety* behaviors cause the robots to fall into a V-formation when the ball is in motion roughly towards the opponent's goal. Perceptual properties limit the formation to three robots.

7.2.1 Offensive formation

The cooperative behaviors result from the interaction of simple behaviors. *Push* causes the robot to line up behind the ball and push it towards the opponent's goal. A second behavior, *Safety*, causes the robot to maintain the maximum safe velocity (as determined by sonar sensors). A third behavior, *Disperse*, causes the robot to rotate away from anything too close to its sides. Finally, a *Patrol* behavior causes the robot to patrol its half of the field defensively when it has not perceived the ball for a few seconds.

In an offensive situation, one robot serendipitously gets to the ball first and begins to *Push* it forward. Teamates also try to *Push*, but their *Disperse* and *Safety* behaviors slow them down and steer them away when they get very close to the *Push*ing robot, and thus tend to fall into the formation shown in Figure 11.

This formation provides an effective "fumble protection" that is essential in the robotic soccer domain. Robots often accidentally knock the ball off course while dribbling it forward; this formation provides backup and recovery. With this formation it is not uncommon for possession of the ball to transfer between the robots of an advancing group without loss of possession by the team. The formation also provides for a very quick defense if the ball is stolen (see below).

The size of the offensive formation is limited by the interaction between the four behaviors above and the physical bodies of the robots. Once there are three robots in the formation, any other robot trying to *Push* the ball will have its view of the ball occluded by the bodies of the first three robots in the formation. When this occlusion lasts for more than a second or two, the *Patrol* behavior gains control of the robot and it gives up on following the ball. In this way, necessary roles are filled (attacker, supporters, and defense) without negotiation, explicit definition or assignment of roles, or even any representation of teammates.

7.2.2 Defensive group formation

In a defensive situation the ball is not advancing toward the opponent's goal. The same behaviors described in Section 7.2.1 cause the robots to fall into a semi-circular arrangement around the ball rather than the V-formation of the advance (see Figure 12), since the robots on the sides are no longer kept behind by lower speed. This formation very effectively prevents the opponent from continuing to move the ball up the field, and places players in a good position to gain possession of the ball. An emergent "batting behavior" (another result of the interaction between the four behaviors listed above, described in [Werger 1999]) makes it likely that the *Push*ing robot will jostle the ball towards one of its teammates, which can smoothly begin an advance from the side.

7.2.3 Transition between formations

Transition between offensive and defensive formations is determined by motion of the ball, and is not even perceived by the robots; there is no concept of "offensive" or "defensive" (or even of "formation") anywhere in the behavior structure. Simple sensing of the local environment leads to flexible, dynamic team behavior that many researchers claim requires higher deliberation and explicit communication ([Balch and Arkin 1998, Balch and Arkin 1995, Jennings 1995, Onn 1997, Stone and Veloso 1998, Tambe 1997]; see [Kraus 1997] for a detailed discussion).

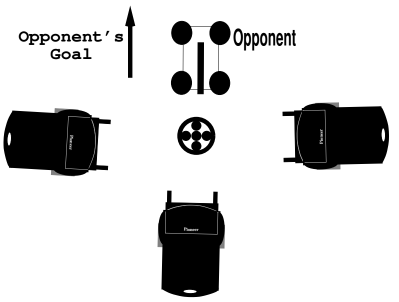

Figure 12: Defensive Formation in Robot Soccer: When the ball is not moving roughly towards the opponent's goal, the robots cluster around it to form an effective barrier and be in good positions for recovery.

8 Discussion

[Donald et al 1994] demonstrated the utility of their theoretical framework of information invariants in analyzing tradeoffs and equivalences between sensor systems. Specifically, they showed the reducibility of one system that used explicit communication between two robots to one that did not (i.e., which communicated solely through task dynamics). They also raised the following questions: 1) "can robots *externalize*, or record, state in the world?" and 2) "can we record *programs* in the world in the same way we may externalize state?" Our research addresses these questions with a system of robots that form distributed physical representations of spatial information. Where [Donald 1995] discusses "calibrations" of sensor systems which fix certain spatial relationships (effectively encoding spatial information) in the system, we present a system that continuously calibrates itself to encode changing information into a distributed representation of spatial relationships, or, in other words, to continuously re-engineer the environment so as to influence the behavior of individual agents. Since these physical representations direct the behavior of agents within the system, they may be seen as "programs" that the system as a whole encodes into the environment for "execution" by its parts.

Practically, we can see that such externalization of state and control allows a wider range of robots - particularly, much simpler robots - to perform various classes of tasks. Through collective behavior, local (at the extreme, physical contact) sensors can suffice for tasks that require global position information. Such sensors can also suffice for making coverage guarantees, and for adaptive task division and "role assumption" subject to global constraints.

More philosophically, in externalizing more and more of the cognition required to perform any task, we shift our focus farther from intra-agent processing and further towards interaction between agents. The extreme simplification of control within an agent allows us to locate interesting behavior at this level of interaction. Since these interactions are all physical and observable, our vantage point for observation of "emergent" behavior is substantially improved.

9 Conclusion

We have shown that chains of robots using only physical contact sensing can solve certain global position-dependent problems. This contradicts a heretofore assumed need for more complicated sensors, positioning systems, and processing. Many environments and applications (especially a number of those proposed for development of "nanorobot swarms", undersea exploration, and space exploration), due to size and/or ambient noise factors, impose exactly these types of restrictions on position-dependent tasks. Systems similar to that described here should drop the lower bound on hardware (and therefore cost) requirements for a wide range of position-dependent tasks, and extend the range of environments in which they are possible.

We have also shown that, in conjunction with a single GPS-equipped "base station," such a system can suffice to provide guaranteed coverage of an arbitrarily large region and to provide good approximations of the locations of encountered objects.

Finally, we have shown that, using only local sensing, both chaining and soccer-playing robots can divide into dynamically-assembled subgroups whose sizes are determined by global optimality constraints.

Some robotics research has presented or reproduced particular instances of stigmergy ([Holland and Melhuish 2000, Beckers et al 1994, Deneubourg et al 1991, Theraulaz et al 1991]), but analysis has remained at the level of claims of greater robustness or ease of scalability than an often undescribed "centralized" system. Many proposed robotic applications are poised to take advantage of these properties of stigmergy, but must wait for a better understanding of what the systems *can* do, and likely the ability to make some guarantees about what the systems *will* do. The robot chaining, sweeping, and soccer-playing systems are examples of deliberate and useful exploitation of stigmergic effects that we hope will serve as inspiration and object of analysis for development of methodologies for externalization.

Acknowledgements

The research described here was supported by the Office of Naval Research Grants N00014-95-1-0759 and N0014-99-1-0162, and the National Science Foundation Infrastructure Grant CDA-9512448.

References

[Aron et al 1990] "Functional Self-organization Illustrated by Inter-nest Traffic in Ants: The Case of the Argentinian Ant," S. Aron, J.L. Deneubourg, S. Goss, J.M. Pasteels, *Biological Motion , Lecture Notes in BioMathematics* Vol. 89, W. Alt and G. Hoffman, eds. Springer-Verlag, Berlin.

[Balch and Arkin 1998] T. Balch and R. Arkin, Behavior-based Formation Control for Multi-robot Teams. *IEEE Transactions on Robotics and Automation*, December 1998.

[Balch and Arkin 1995] T. Balch and R. Arkin, Motor Schema-based Formation Control for Multi-agent Robot Teams. *Proceedings of the 1995 International Conference on Multiagent Systems*, San Francisco, CA pp. 10-16.

[Beckers et al 1994] "From Local Actions to Global Tasks: Stigmergy and Collective Robotics" R. Beckers, O.E. Holland and J.L. Deneubourg, *Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, 181-189.

[Brooks 1986] "A Robust Layered Control System for a Mobile Robot", Rodney A. Brooks, *IEEE Journal of Robotics and Automation*, RA-2, April, 14-23.

[Brooks 1990] "The Behavior Language; User's Guide", Rodney A. Brooks, *MIT A.I. Lab Memo 1227*, April.

[Deneubourg et al 1991] "The Dynamics of Collective Sorting: Robot-Like Ants and Ant-Like Robots", J.L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain and L. Crétien, *Proceedings of the First International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, MIT Press, 356-363.

[Deneubourg et al 1990] "Self-organizing collection and transport of objects in unpredictable environments," J. Deneubourg, S. Goss, G. Sandini, F. Ferrari, and P. Dario, *USA-Japan Symposium on Flexible Automation*, Kyoto, Japan, July.

[Donald et al 1994] "Information Invariants for Distributed Manipulation," B. Donald, J. Jennings, and D. Rus, *The First Workshop on the Algorithmic Foundations of Robotics*, R. Wilson and J.-C. Latombe, eds. A.K. Peters.

[Donald 1995] "Information Invariants in Robotics," B. R. Donald, *Artificial Intelligence* 72.

[Fontán and Matarić 1996] , "A Study of Territoriality: the Role of Critical Mass in Adaptive Task Division," Miguel Schneider Fontán and Maja J. Matarić. *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, Cape Cod, September.

[Goldberg and Matarić 1996] "Interference as a Tool for Designing and Evaluating Multi-Robot Controllers" Dani Goldberg and Maja J Matarić. *Proceedings of AAAI-97*, Providence, RI, July, pp. 637-642.

[Grasso et. al. 1996] "Locating Odor Sources in Turbulence with a Lobster-Inspired Robot," Frank Grasso, Thomas Consi, David Mountain, and Jelle Atema. *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior: From Animals to Animats 4*. The MIT Press, pp. 104-112.

[Gomi 1995] "Non-Cartesian Robotics", Takashi Gomi, *Robotics and Autonomous Systems* 18:1-2, July 1996. pp. 169-184.

[Gordon 1999] Ants at Work, Deborah M. Gordon. The Free Press, New York.

[Goss et al 1989] "Self-organized Shortcuts in the Argentine Ant," S. Goss, S. Aron, J.L. Deneubourg, and J.M. Pasteels, *Naturwissenschaften* 76, pp. 579-581.

[Goss and Deneubourg 1991] "Harvesting By A Group Of Robots," S. Goss, J. L. Deneubourg, *Proceedings of the First European Conference on Artificial Life*, MIT Press.

[Hölldobler 1990] The Ants. Bert Hölldobler and Edward O. Wilson, Cambridge, Massachusetts: The Belknap Press of Harvard University Press.

[Holland and Melhuish 2000] "Stigmergy, self-organisation, and sorting in collective robotics," Owen Holland and Chris Melhuish. *Artificial Life* 5:2.

[Horswill 1993] "Specialization of Perceptual Processes", Ian D. Horswill, MIT PhD Thesis, May.

[Jennings 1995] N. Jennings, Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence* 75.

[Kraus 1997] S. Kraus, Negotiation and cooperation in multi-agent environments. *Artificial Intelligence* 94.

[Matarić 1992b] "Integration of Representation Into Goal-Driven Behavior-Based Robots," Maja J Matarić, *IEEE Transactions on Robotics and Automation*, 8:3, June.

[Matarić 1990] "Navigating with a Rat Brain: A Neurobiologically-Inspired Model for Robot Spatial Representation", Maja J Matarić, *Proceedings of the First International Conference on Simulation of Adaptive Behavior: From Animals to Animats*, MIT Press.

[Matarić 1995] "Designing and Understanding Adaptive Group Behavior," Maja J Matarić, *Adaptive Behavior* 4:1, December, pp. 51-80.

[McFarland 1994] "Towards Robot Cooperation", David McFarland, in *Proceedings of the Third Annual Conference on Simulation of Adaptive Behavior: From Animals to Animats*, D. Cliff, P. Husbands, J. Meyer, and S. Wilson, eds., Cambridge, MA, MIT Press.

[Onn 1997] S. Onn and M. Tennenholtz, Determination of social laws for multi-agent mobilization. *Artificial Intelligence* 95 155-167.

[Russell 1995] "Laying and Sensing Odor Markings as a Strategy for Assisting Mobile Robot Navigation Tasks," R. Andrew Russell. *IEEE Robotics and Automation Magazine*, September.

[Steels 1994] "A Case Study in the Behavior Oriented Design of Autonomous Agents,"Luc Steels, in *Proceedings of the Third Annual Conference on Simulation of Adaptive Behavior: From Animals to Animats*, D. Cliff, P. Husbands, J. Meyer, and S. Wilson, eds., Cambridge, MA, MIT Press.

[Stone and Veloso 1998] P. Stone and M. Veloso, Task Decomposition and Dynamic Role Assignment for Real Time Strategic Teamwork. *Proceedings of the Third International Conference on Multi-Agent Systems.*

[Tambe 1997] M. Tambe, Towards flexible teamwork. *Journal of Artificial Intelligence Research* 7.

[Muller and Wehner 1988] "Path Integration in Desert Ants: Cataglyphis Fortis," Martin Muller and Rudiger Wehner. *Proceedings of the National Academy of Sciences.*

[Theraulaz et al 1991] "Task differentiation in Polistes wasp colonies: a model for self-organizing groups of robots," Guy Theraulaz, Simon Goss, Jacques Gervet, and Jean-Louis Deneubourg. *Proceedings of the First International Conference on Simulation of Adaptive Behavior: From Animals to Animats.* The MIT Press.

[Werger and Matarić1996] "Robotic Food Chains: Externalization of State and Program for Minimal-Agent Foraging," Barry Werger and Maja J

Matarić. *Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior: From Animals to Animats 4*. The MIT Press, pp. 625-634.

[Werger 1999] "Cooperation Without Deliberation: A minimal behavior-based approach to multi-robot teams," Barry Werger. *Artificial Intelligence* 110, pp. 293-320.