Robotic "Food" Chains: Externalization of State and Program for Minimal-Agent Foraging

Barry Brian Werger and Maja J Matarić
Interaction Laboratory
Department of Computer Science
Brandeis University
Waltham, MA 02254
barry@cs.brandeis.edu, maja@cs.brandeis.edu

Abstract

This paper describes experiments inspired by theoretical work on information invariants ([Donald 1995], [Donald et al 1994]), a means of comparison and a methodology for design of single- and multi-agent systems. Analysis reveals the environmental information that the systems assume and exploit, while the design methodology seeks to move information and processing into the "physical" environment and task mechanics. The approach raises the issue of agents actively recording information, or even "programs," into the physical environment. This paper provides an example system that dynamically encodes information and "programs" into its physical environment.

The second source of inspiration for this work is the natural phenomenon of ant pheromone trail formation, shown to involve agents with simple, local control that encode information into the environment to arrive at globally complex behavior. Analogously, our robotic system actively encodes information into its physical environment in order to reduce sensing, actuation, and computational requirements. Thus, "minimal" agents with local sensing and action form a system that dynamically and globally adapts to environmental changes. We discuss how moving information and "processing" 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 numerous systems that react effectively and robustly to their environments, they have dealt only obliquely with the 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, influence the behavior of other agents in a manner that is more than mere "interference"; [Beckers et al 1994] describes a task where physical effects of task performance allow a simple, local control strategy to produce a consistent global behavior; and 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 marking or pheromone trails. These phenomena exploit the benefits of having agents deliberately encode information into the physical environment. As discussed in [Aron et al in press], [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," which is a function of frequency of use and decay. If pheromones are released only during certain phases of tasks (such as 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 determined 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]), 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 the amount of sensing, computation, and communication which naively seems "necessary," and 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 are: 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. We discuss this system as both the object of analysis and inspiration for development of an extended information invariant-based approach we hope to develop in the future, including making steps towards extension of the methodology for the development of distributed manipulation protocols [Donald et al 1994].

2 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 requires some knowledge of global positioning for efficient performance. While purely stigmergic solutions have been found for tasks such as clustering items in the environment ([Beckers et al 1994]) and even sorting of scattered heterogeneous items into homogeneous clusters ([Deneubourg et al 1991]), tasks which require particular behaviors to take place at particular locations have so far relied upon some type of global position sensing, globally visible beacon, or random encounter of some locally-sensible position marker [1].

The following subsection gives 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.

[1] With the notable exception of one of our inspirations, the simulated beacon chain system described in [Deneubourg et al 1990] and [Goss and Deneubourg 1991].

2.1 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. Heading information can be determined through analysis of change in position. This is the basis of several successful foraging systems ([Fontán and Matarić 1996], [Goldberg Matarić 1996], [Matarić 1995]), but requires preparation of the environment (radio-sonar broadcasters at precise locations), complicated sensing equipment (radios, sonar detectors), 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 highly accurate and potentially complicated actuator motion sensing and calibration, and suffers from cumulative errors.
- *Taxis:* Following of some sort 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 map representation of the environment, and sensing the landmarks sufficiently accurately to local-

ize within that map.

2.2 Methods specific to multiple agents

- *Pheromones:* Requires the ability to emit and detect the presence of varying concentrations of pheromones. Work towards a robotic odor sensing/depositing system has been done by [Russell 1995], but has not yet been applied to tasks.
- *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 beacons) from the home location.
- *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. We propose an extreme case in which the robots use a small number of the simplest, most reliable sensors available - contact sensors - as discussed below.

3 Robot Chains

The system we present involves the formation of "hand-holding 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.

Earlier research on robot chains has been conducted through simulation in [Goss and Deneubourg 1991] and [Deneubourg et al 1990]. The chains were "line-of-sight" and required that each simulated robot be able to distinguish among beacons and locate those which communicated numbers representing distance (in beacon links) to *Home*. Unfortunately, this approach requires sophisticated sensors and transmitters, and given those, could still be sensitive to sensory errors and other noise in non-ideal environments.

To avoid these problems, the approach to chaining that we present uses only sensors that operate within the range of physical contact - microswitches and break-

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

beams. The microswitches indicate physical contact with another object. Each member of the chain actively maintains contact with the link ahead and behind by touch, through microswitches.

Limited communication is implemented through the same mechanisms to allow for chain maintenance. The most common type of communication 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.

The basic behaviors involved in chain formation and maintenance are:

- *HomeLink:* Remains still, except to maintain communication with next link in the chain.
- *MiddleLink:* Maintains communication with "next" and "previous" links by returning taps and passing messages.
- *EndOfChain:* Maintains communication with link "ahead," 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, it communicates its intention to the "previous" link to transfer *EndOfChain* status and leaves the chain.

- *JoinChain:* Works with *EndOfChain* to align a robot properly at the end of the chain, establish communication with the current *EndOfChain*, and become the new *EndOfChain*.

Robots not part of the chain can determine which way along the chain *Home* is, and follow the chain towards or away from it, using a physical feature that allows simple sensors to determine a very rough estimate of heading relative to another robot.

Behaviors involved in chain following are:

- *GoHome:* Determine direction *Home* and follow chain in that direction.
- *GoOut:* Determine direction away from *Home* and follow chain in that direction.

The foraging task our chain-making system performs involves the collection of metal pucks scattered either randomly or in clusters around the test environment. Two types of searching are used to locate and retrieve pucks:

- *Random Search:* Robots search for pucks throughout the environment, then locate the chain randomly once carrying a puck. This is often performed when the end of the chain is reached.
- *ExcursionSearch:* Robots follow chain, occasionally taking roughly circular journeys into the area next to the chain (Figure 4.)

Modifications of the basic behaviors discussed above allow for dynamic adjustment of the chain to various environmental factors. As mentioned earlier, in certain tasks it is desirable for the chain to connect a rich supply directly to *Home*. One way for the chain to move towards such a configuration is for the links to collect statistics on the number of times they are tapped on each side, and gradually shift towards the side that sees the most "traffic." We see this as somewhat analogous to the gradual buildup of pheromones on paths frequently used by ants; it should eventually lead to the same type of convergence on a shortest path to a highly useful destination ([Aron et al in press], [Matarić 1990], [Goss et al 1989]).

4 Current Implementation

We have implemented a foraging system which gathers 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

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.

Home area. The system is functional with the following qualifications:

- *Sensing Home:* Infrared emitter/detectors with an effective range of less than 1 inch, located on the underside of the robots' fork arms, are used to determine when the robots are at *Home*, which is a nonreflective black area on the floor. This extremely short-ranged sensor can be replaced with a physical sensor of the same length capable of detecting some property of *Home*.
- *Initial Timing:* Currently the 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 the waiting team members.
- *Number of Robots:* As described below in 4.2, our herd of 20 robots is undergoing major renovations and modifications. The described experiments were performed with four robots fully capable of chain-building behavior, and two additional robots performing only behaviors based on following the chain. 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.

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

- *Environmental Assumption:* The current system assumes that the environment contains only robots, pucks, and *Home*.

4.1 Behaviors

4.1.1 Initial Chain Location - **Skirt**

Robots start gathered at *Home*. A behavior **Skirt** navigates to the edge of *Home*, then tacks along this edge until it encounters a physical obstacle projecting outside of the *Home* region. This obstacle is assumed to be the chain.

4.1.2 Chain Following - **Tack**

Chain following is performed through simple tacking. The following robot angles towards the chain until contact is made, backs off at a sharper angle, then angles back to make contact further down 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.

4.1.3 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 sends out a signal. This signal is used to deactivate **Tack** and activate **BackInto**.

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 front of 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 sends a signal which deactivates **BackInto** and activates **AlignBack**.

AlignBack delays 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* (there is currently no specialization required for the *EndOfChain*). 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**.

4.1.4 MiddleOfChain - **Link**

MiddleOfChain has been actualized 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 front and single contacts with the previous chain link allow recovery from most errors.

4.1.5 Excursion Search

Excursion Search has been implemented through an extended **Tack** and a behavior **CircleRight**. The extended **Tack** makes a decision every time it contacts the

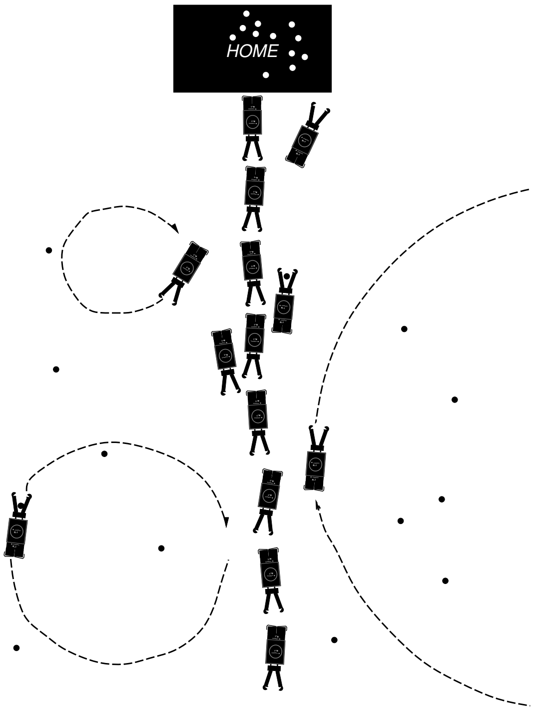

Figure 4: Excursion Search strategy: Robots search for pucks and return to the chain by making roughly circular "excursions" from the chain.

chain as to whether or not 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 (1/8 chance in our experiments).

4.2 The Robot Herd

Our experiments are implemented and tested on the Nerd Herd, the Interaction Lab's group of 20 IS Robotics R1 mobile robots. Each member of the Nerd Herd is a 12–inch 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 metal foam-filled disks, 1.5 inches 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 magnet. Each robot also has one piezo–electric bump sensor on each side of the chassis. Only the front contact, the stacking IRs,

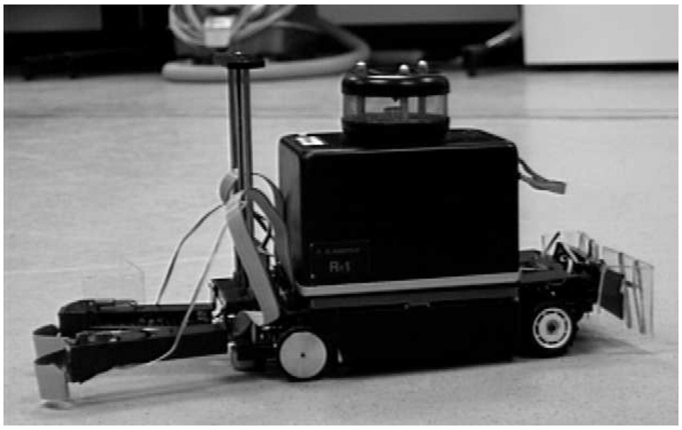

Figure 5: Each of the Nerd Herd robots is a 12"–long four–wheeled base equipped with a two–pronged forklift for picking up, carrying, and stacking pucks, and with a radio transmitter and receiver for inter–robot communication and data collection.

and rear contact sensors described in 4.2.1 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. 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].

4.2.1 Hardware Modifications

Originally equipped with piezo–electric bump sensors on the back of the chassis, the venerable robots are being modified to better suit the chaining task. The rear surfaces of some robots now have large bumpers that activate contact switches (see Figure 6). This is necessary due to the nature of the 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.

4.2.2 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 constrainted by various sensory, mechanical, and computational limitations.

Our robots' mechanical steering system, when in perfect condition, is "accurate" to within 30 rotational degrees. At certain steering angles, the 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.

4.3 Performance of Current Implementation

The foraging system that we tested with six working 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. The effective length of a chain can be said to be approximately 1.5 times the length of the robots that form it.

Chain following suffered only rare mechanical failures. Some of those cases resulted in the following robot changing the position of one or more of the chain robots to such degree that chain integrity was broken. Detecting

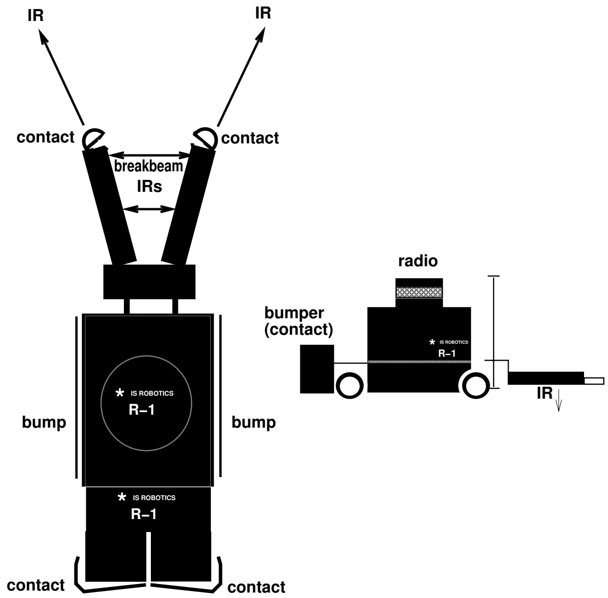

Figure 6: Each of the Nerd Herd robots is equipped with contact sensors at the ends of the fork, piezo–electric bump sensors on each side and two on the rear of the chassis, and six infra–red sensors on the fork. Two forward–pointing IRs are located at the ends of the forks, two break–beam IRs in the 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.

and recovering from such problems is a step on the way to dynamic chain readjustment, and is currently being developed.

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 prevents the front sensors from making any contact. This problem was resolved in some trials through a time-out which backs 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 will 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. This is exactly the type of trade-off we intend: that a large number of less capable, more robust, 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 outsurvives, 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.

5 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 raise 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. Future research will begin the process of extending the information invariants-based analysis and develop the existing design methodologies to encompass the notion of dynamic self-calibration.

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.

6 Conclusion

We have shown that chains of robots using only physical contact-range sensing can solve certain global position-dependent problems. This contradicts a heretofore assumed need for more complicated sensors, positioning systems, or 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.

Some robotics research has presented or reproduced 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, for example, [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 system is one example of a 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 reported here was done at the Interaction Laboratory at the Brandeis University Volen Center for Complex Systems and the Computer Science Depart-

ment. The work is supported by the Office of Naval Research Grant N00014-95-1-0759 and the National Science Foundation Infrastructure Grant CDA-9512448.

The authors thank Jaroslav Hook, Francisco Mello Jr., and Lester Lehon for their contributions to the functionality of the Nerd Herd, Dani Goldberg for emergency cool, Jordan Pollack for getting his hands dirty, and Pablo and Karina Funes for the *remedio*.

References

[Aron et al in press] "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*, W. Alt and G. Hoffman, eds.

[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 Matarić 1996] "Interference as a Guide for Designing Efficient Group Behaviors'," Dani Goldberg and Maja J Matarić. Brandeis University Computer Science Technical Report CS-96-186.

[Gomi 1995] "Non-Cartesian Robotics", Takashi Gomi, *Proceedings of the International Conference on Biorobotics: Human-Robot Symbiosis.*

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

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

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

[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.

[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.

[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 An-*

imats, D. Cliff, P. Husbands, J. Meyer, and S. Wilson, eds., Cambridge, MA, MIT Press.

[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 of Simulation of Adaptive Behavior: From Animals to Animats.*