

Exploring coordination properties within populations of distributed agents

Elizabeth Sklar², Martijn Schut¹, Konrad Diwold¹, Simon Parsons²

¹Artificial Intelligence Section, Dept of Computer Science
Faculty Of Sciences, Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV Amsterdam, NL
schut,kdiwold@few.vu.nl

²Dept of Computer and Information Science
Brooklyn College, City University of New York
2900 Bedford Ave, Brooklyn, NY, USA 11210
sklar,parsons@sci.brooklyn.cuny.edu

Abstract

This paper presents a discussion of coordination properties within populations of geospatially distributed embodied agents. We define two axes: *interaction mechanisms* and *population diversity*; and we present a new framework designed for exploring the relationship between values along these axes and the efficiency of solutions for a set of related tasks, e.g., foraging, resource allocation and area coverage.

Introduction

There are many issues that arise in coordinating populations of geospatially distributed agents. This paper presents an overview of properties of coordination in an embodied multi-agent system and a framework for examining them. We consider a class of tasks in which a physically dispersed, large group of heterogeneous agents, if well coordinated, can achieve better results than a single, multi-capable agent or a small group of agents possessing mixed capabilities. By “large”, we mean 50 agents or more. By “well coordinated”, we mean that the results achieved are better with a coordination mechanism than without one (i.e., random acts of agent kindness).

The term *coordination* is discussed from many perspectives within the embodied multiagent (or multi-robot) literature. Cao, Fukunaga, & Kahng (1997) offer a broad survey of coordination in multi-robot systems, provide historical perspectives and theoretical background, and categorize their findings along five “research axes”: group architecture, resource conflict, origin of cooperation, learning and geometric problems. Kalra, Stentz, & Ferguson (2004) outline a spectrum of tasks, from “tightly” to “moderately” to “loosely” coordinated, providing a dimension in which to consider the necessity for robots to coordinate in order for a task to be accomplished successfully. Farinelli, Iocchi, & Nardi (2004) propose a two-dimensional taxonomy and parameterization for multi-robot systems, based on coordination — cooperation (yes or no), knowledge (aware or unaware), coordination (weak to strong) and organization (centralized to distributed) — and system-level features (communication, team composition, architecture and size).

In the non-embodied multiagent systems (MAS) world, researchers are primarily concerned with understanding theoretically how systems might best coordinate, and they design simulated systems and mechanisms that approximate (or ignore) the physical world in order to demonstrate, test and evaluate theories. Conversely, in the multi-robot (MRS) arena, most researchers are more focused on building systems that operate robustly in the physical world; they are primarily concerned with “making it work” and frequently rely on highly engineered or emergent *ad hoc* methods. In the long term, we are interested in bridging the gulf that divides theoretical MAS and applied MRS approaches to coordination.

Here, we concentrate on two axes: *interaction mechanisms* and *population diversity*. It is our contention that there exists a dependent relationship between the diversity of a population and the level of interaction amongst population members for achieving coordinated solutions. Our theory is that heterogeneous populations which make proper use of interaction mechanisms to coordinate will perform as well or better than homogeneous populations of multi-capable agents. Indeed, there are many real-world applications where large populations of homogeneous agents are impractical due to financial or logistical constraints; thus, it is important to find efficient means for coordinating populations of heterogeneous agents.

Our longterm goal is to understand how the complexity along each of these axes affects efficiency within the solution space. In this paper, we outline some ideas in this area and describe a framework which we have constructed to explore them. The paper is organized as follows: first we detail several types of interaction mechanisms, of varying complexities; second, we explain factors relating to population diversity. Then, we describe our framework and newly constructed prototype simulation environment designed for experimentation. Finally, we discuss current and future work.

Interaction mechanisms

Coordination is enabled via interactions between agents, which can be either implicit or explicit; an interaction mechanism is used to facilitate any resulting coordinated activity on the part of the agents in a system. We use the term *coordinated plan formation* to refer to the three-part process that occurs between an agent receiving sensory information

and performing actions, as shown in Figure 1. Interactions are interpreted and considered before an agent commits to a plan of action. Implicit interaction refers to mechanisms where environmental cues serve to pass information from one agent to another. Explicit (or intentional) interaction refers to mechanisms where agents purposely and purposefully send messages to one another. Note that in dynamic environments, agents must be able to “walk and talk at the same time”, i.e., they must be able to sense, plan, interact and act in parallel.

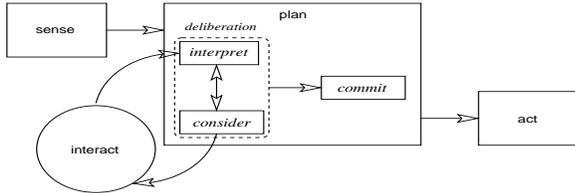


Figure 1: Coordinated plan formation

Coordinated plan formation can occur using either a centralized, distributed or hybrid approach. In a centralized approach, one “leader” agent determines a plan for all the other agents; thus the *consider* step in Figure 1 is effectively eliminated, as all non-leader agents automatically “commit” to the leader’s plan without deliberating. In a decentralized, or distributed, approach, each of the individual agents forms its own plan. Agents may receive input from other agents that might inform their local decision making, but the final determination of what plan to commit to is done at the level of the individual agent. Hybrid approaches, as the name suggests, combine the two. The most common types of hybrid approaches are market mechanisms, such as auctions, in which agents consider plans locally, make decentralized decisions about what they want to do, and then submit “bids” to a central decision-making facility which determines which bids to accept and sends assignments back to the bidding agents. If an agent’s bid for a plan is accepted, the agent is bound to commit to that plan. While some may argue that this, too, is a centralized approach, we contend that the distinction that divides plan consideration (bidding) from assignment of plans is important because there exists a large body of ongoing research examining one or the other aspect (or both, but considered separately). Parkes & Shneidman (2004) propose “distributed implementations” in which a market mechanism is used but there is no “auctioneer”, i.e., both plan consideration and assignment are decentralized. Work in this direction may be particularly fruitful for MRS due to inherent problems with communication.

Interaction mechanisms can be parameterized as follows:

- *medium* — the way a message is transmitted from one agent to another, either directly, through the environment or through a third, mediating agent; values can be: environmental, broadcast, brokered or peer-to-peer.
- *content* — the type of data a message contains; values can be: binary, scalar, vector or object.
- *decision mechanism* — who makes a decision about

Table 1: Taxonomy classes and parameter values for different interaction mechanisms, listed top-down from simplest to most complex

classification	interaction medium	content	decision mechanism	messaging	synchronization
tacit agreement	none	none	distributed	n/a	none
environmental cues	environmental	scalar	distributed	continual	none
signal broadcasting	broadcast	binary / scalar	distributed	continual	none
auctions	brokered	scalar or vector	hybrid	once (or in rounds)	turn-taking (with broker)

reaching an agreement; values can be: centralized, distributed or hybrid.

- *messaging* — how many times messages can be passed; messages are either sent once, in rounds until agreement (or stalemate) is reached, or continuously.
- *synchronization* — when messages can be passed, dictating whether agents wait for acknowledgement and/or responses before sending additional messages (i.e., turn-taking) or not.

Table 1 contains parameter values for four different interaction mechanisms, each of which is detailed in the paragraphs below. We note that our parameterization builds on that of Weiss (1999), where agent interactions are characterized according to frequency (how often a message is passed), persistence (how long a message survives), level (the complexity of the message content), pattern (how the data flows), variability (if the data flow is regular) and purpose (competitive or cooperative). Weiss’ definition of “level” is similar to our notion of “content”, though we outline specific content categories; our remaining parameters expand on his set.

Tacit agreements are the minimal form of interaction mechanism, because there is no explicit communication between agents, rather, social norms or pre-determined rules dictate what agents decide to do. Shoham & Tennenholtz (1992) explore the use of norms and social laws to coordinate the behavior of virtual agents in a MAS; conventions are either encoded directly into the system or emerge as the system runs. Minimal interaction mechanisms are most successful for coordinating fairly simple behaviors, primarily physical *formation control*. Kowalczyk (2001) presents mathematically three engineered solutions: leader-follower (where one robot is chosen to be the leader and the rest of the robots are designed to follow the leader’s trajectory), virtual structure (where the robots together form a virtual, rigid structure of fixed geometric shape), and behavior-based (where each robot is given a set of low-level individual behaviors, such as “avoid obstacles”, and a set of high-level group goals, such as “surround a target”). Fierro *et al.* (2002) describe another methodology for robot team formation in which robots’ control algorithms are designed so that individuals maintain certain distances from their neighbors while at the same time, they avoid obstacles. Balch & Arkin (1995) present work on formation control in groups of robots in which a hybrid mechanism is used for coordination; robots employ a two-step process in which they first

perceive where they are supposed to be in the formation, given their current location and those of the other robots, then they move into position to maintain the formation. Stone & Veloso (1998) propose the use of “locker-room agreements” for the RoboCup soccer simulation league as a way of implementing synchronized decision making in terms of roles and strategies for the team — the robots agree before the game to change their behavior based on some change in the state of the environment.

Environmental cues are mechanisms in which one agent signals to another by modifying the environment in some way, for example one agent may lay a trail that another can then follow. The classic example is ants depositing pheromone trails when they move towards a food source. Dorigo, Maniezzo, & Colomi (1996a) introduce the notion of an “ant system” and describe a series of experiments in which a population of simulated ants solves an instance of the classic travelling salesman problem. When an ant travels from one city to another, it leaves a trail which effectively increases the “weight” of a virtual edge between cities, symbolizing the deposit of pheromone by the travelling ant. The trail decays (or “evaporates”) over time, so that paths are learned not only because many agents traverse them but also because agents continue to traverse them as time passes. Kube & Zhang (1992) describe a common approach to using environmental cues for coordinating a team of robots wherein various “perceptual cues” can be observed that indicate changes in the state of the environment; recognition of certain states triggers behavioral actions which result in coordinated activity amongst the robots.

Signal broadcasting is an extension of environmental modification in which the messages are actively sent from one agent to another. A typical broadcast mechanism (like wireless radio communication) propagates fast; in most environments, all agents within range receive a broadcast at nearly the same time and make their decisions almost simultaneously. In contrast, when using environmental cues, the message is only readable when an agent is within sensor-range and so some agents may not read the message until some time after it has been left. Yanco & Stein (1993) present an experiment in which robots learned to coordinate their movements in the form of leader-follower behaviors by broadcasting signals. In ALLIANCE (Parker 1998), robots allocate tasks dynamically, communicating with each other by broadcasting messages indicating which task they are working on and the status of the task. Watson, Ficici, & Pollack (1999) present an on-line, evolutionary approach in which a population of perceptron-controlled robots learns to seek a light source by broadcasting weights at a frequency proportional to each robot’s success at the task. Werger & Mataric (2000) implement a “broadcast of local eligibility” architecture in which robots simply broadcast their fitness as they move around an environment, at the same time reading the fitness values of their neighbors, and acquiring more fit values as they are encountered. Vaughan *et al.* (2000) construct a team of robots that emulate honey bees by learning a path to a common resource and returning to a home position, performing a “waggle dance” in which they share information about the location of resources, broadcasting only

when successful paths are discovered. Saffiotti, Zumel, & Ruspini (2000) compare broadcast ranges in a multi-robot surveillance task; their results show that a local range works best in their environment, where timing and location are closely coupled. Kalra, Stentz, & Ferguson (2004) demonstrate signal broadcasting as the “passive coordination” approach taken in PC-MVERT, a system that was designed for experimentation with “tight”, “planned” and “computationally feasible” coordination; when agents communicate their current plan to their neighbors, results are improved over situations where there is no communication at all.

Auctions, in economic terms, are market mechanisms — for selling and/or buying some commodity — in which messages that agents send are indications of how much they are willing to pay (or accept, in the case of a seller) and priority is given to high-price offers to buy and low-priced offers to sell (Friedman 1993). We distinguish “simple auctions”, where bids consist of a single number indicating an agent’s desired trading price, from “combinatorial auctions”, where a bid will consist of several values specifying a price for a particular combination of a set of goods. For agent interaction, the mechanisms are adopted with some scalar playing the role of currency, and the result typically determines not the allocation of commodities but what the bidding agents will do i.e., task allocation. The mechanism involves agents sending offers to an intermediary, an “auctioneer”, who determines a matching on the basis of these offers. One of the most widely used protocols in implemented multiagent systems today is the Contract Net (CNET) protocol (Smith 1977; Smith & Davis 1980), which is a simple yet robust methodology for task sharing. In the CNET protocol, one agent is designated as the “manager” who “announces” a task to a network of other agents by sending a message called a “task announcement”. Different levels of broadcasting this announcement can occur (to a whole network, portions of a network or individual agents) depending on what the manager knows about the capability (and availability) of the other agents in the network. Agents submit “bids” indicating their perceived ability to perform a task. The manager decides whose bid(s) to accept and sends “awards” to those agents. In typical MAS applications of CNET, bids are estimates of the cost of carrying out the task; and the manager accepts the lowest bid for each task — under such circumstances the CNET is a simple (first-price sealed-bid) auction. The effect of CNET is to distribute tasks to the agents that can best perform them. Sandholm (1993) developed an early application of CNET, used for vehicle routing, and later proved that exchanging sets of tasks amongst agents, instead of bidding for single tasks, can help to avoid local minima (Sandholm 1998). Dias & Stentz (2000) replicated a similar but simpler approach on a multi-robot system, experimenting with single-task contracts and contracts for swapping tasks, to solve a distributed traveling salesman problem. Golfarelli, Maio, & Rizzi (1997) described a task-swapping approach using the contract net protocol for groups of robots performing coordinated path-planning. Simmons *et al.* (2000) introduced a market-based approach to a coordinated map building problem wherein each robot constructs a “bid” based on its expected information gain

for going to a number of different cells; the bids are sent to a “central executive” who assigns tasks to each robot using the bids to determine a global maximization. Gerkey & Mataric (2002) implemented a market-based approach for task allocation in a team of robots working on both tightly and loosely coupled tasks.

Population diversity

There has been a considerable amount of related work addressing the notion of diversity within a population of agents, i.e., the degree to which a group of agents are heterogeneous (as opposed to homogeneous). Balch (1998) found that heterogeneity does not always pay off: for simple tasks, like foraging, a homogeneous population may do better, while for more complex tasks, like soccer, heterogeneity is better. However, several counter-examples have been found. Potter, Meeden, & Schultz (2001) reports that (Luke 1998) found that, for the RoboCup-97 competition, a homogeneous population was better off: under the time constraint of the competition, a heterogeneous population did not converge quickly enough to fit solutions (team compositions). In their own work, Potter, Meeden, & Schultz (2001) let a team of robots co-evolve controllers under different levels of task difficulty, and they examined the trade-offs between homogeneity versus heterogeneity. The robot controllers were neural networks; in heterogeneous populations, individuals each developed their own neural network; in homogeneous populations, individuals shared the same neural network. Their work demonstrated that simply increasing the difficulty of the task is not sufficient to justify heterogeneity.

Many models for division of labor use threshold (Beshler & Fewell 2001) or transition functions (Lerman & Galstyan 2003) to describe how to change from one role to another. A threshold model defines different limiting values for each role; when a specific internal variable in an agent falls within a certain limit range, the agent starts playing a role that corresponds to that range. This limit is usually based on some environmental cue. For example, in a honeybee colony, worker bees adopt a “patroller” role based on the number of perceived enemies. A transition function is based on some (event or observation) history and describes when to adapt another role. For example, Jones & Mataric (2003) developed a simulated robotic application with red and green puck collector roles where robots can change roles if necessary. A transition function is continuously evaluated that takes in (for a red puck collector) the number of observed red pucks and red puck collectors over some fixed time window and outputs the probability to change to a green puck collector.

Numerous models exist in biology for division of labor, which are extensively described by Beshler & Fewell (2001). They discuss six different models that describe the integration of individual worker behaviors into colony task organisation: response threshold models, integrated threshold-information transfer models, self-reinforcement models, foraging for work, social inhibition models and network models of task allocation. While it is beyond the scope of this paper to detail these models here, it is interesting to mention two observations made by the authors: first, the models are not

able to explain the phenomenon of division of labor, as they have merely been exploratory; and second, the evolution of division of labor is one of the necessary areas of study in order to assess how selection and self-organisation interact to produce complex behavioral properties in multiagent systems.

The division of labor models from biology have inspired many roboticists. They exploit the parallel operation and transfer of information among the agents for designing flexible, robust and effective robotic systems (Krieger, Billeter, & Keller 2000). The well-known ant algorithms (Dorigo, Maniezzo, & Colomi 1996b), described earlier, are a prevalent example. Related research on ant-like task allocation (Krieger, Billeter, & Keller 2000) looks explicitly at the relationship between group size (clusters) of agents and the efficiency of the colony. The task allocation was based on variation in individual stimulus thresholds. Experiments show that groups where agents could recruit other agents were significantly more efficient at exploiting clustered resources than groups without information transfer.

Our approach

We have developed a framework which allows experimentation with varying values along the axes of levels of interaction and diversity. Our framework is a grid-based, toroidal artificial world in which agents exist and act in simulated real-time. The prototype version is implemented using NetLogo¹ (Wilensky 2002) and is pictured in Figure 2. The world is designed so that simple re-interpretation of parameters and slight modification of transition functions can easily transform the environment to represent several similar task domains, such as foraging, area coverage and constrained search. The discussion here is limited to a resource collection, e.g., foraging or de-mining task, based on earlier work (Eiben, Nitschke, & Schut 2005).

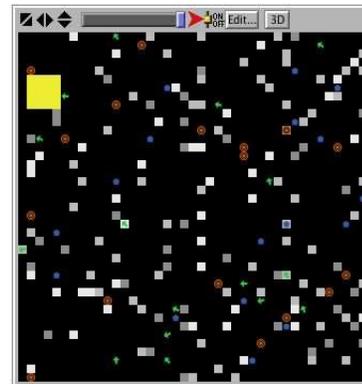


Figure 2: mindScape

The large square region near the upper left corner represents “home”. Small square regions, equal in size to one grid cell, contain resources. Agents are indicated by closed pentagons (“explorers”), open circles (“extractors”) and arrows (“transporters”).

¹<http://ccl.northwestern.edu/netlogo/>

The *mindScape* world is initialized with a number of resources and agents that are scattered randomly throughout the grid. The agents must find resources, gather them and return them to a fixed “home” location. There are three types of resources (labeled *A*, *B* and *C*), where each may require special agent capabilities to gather. Agents are rewarded for bringing resources to the home location, and for each resource, the reward is distributed amongst the three agents (explorer, extractor and transporter) who contributed along the way to the discovery and delivery of the resource. Agents are given three capabilities: *vision*, *extraction* and *transportation*. Agents who each possess only one capability inherently define three “roles”: *explorers*, *extractors* and *transporters*.

Explorers specialize in searching for and discovering resources. Explorers have 8 “vision” sensors, as shown in Figure 3. At each time step in the simulation, the sensors look in 8 directions and return the distance to the nearest cell containing at least one resource. The explorers then communicate to extractors the location of resources that are near them. The system currently can enable either environmental cues or signal broadcasting as a means for explorer-extractor and extractor-transporter interactions.

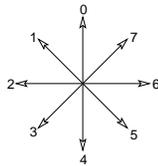


Figure 3: “vision” sensors

Extractors are good at gathering resources — the actual removal of resources from the location on the grid where they are discovered. At each time step, if an extractor is on a cell that contains at least one resource, then the agent can extract a resource from that cell, up to a fixed *capacity*. Each agent is endowed with an extraction capacity, represented by a vector indicating how many of each type of resource the agent can hold at a given time. If an extractor is holding at least one resource, then it cannot move; it has to wait until a transporter crosses its path, then it can transfer its resources to the transporter, move to a new location and start gathering again.

Transporters carry the collected resources to the home location. They are endowed with a transportation capacity, indicating how many resources they can carry at once. When transporters meet extractors, then they receive resources from the extractors and move to the home base (where they can unload the resources). They take the shortest path to the home. If they meet extractors along the way, then they can receive additional resources — up to their maximum capacity.

Each agent’s capabilities can be described by a trio of normalised values between 0 and 1: (v, x, t) , where v = vision capability, x = capacity to extract resources and t = ability to transport resources). We can imagine every agent as a point in this three-dimensional space. In a completely *ho-*

mogeneous population, every agent a would have the same (v, x, t) values as every other agent, e.g.:

$$\forall a_i, a_j, (v_i, x_i, t_i) = (v_j, x_j, t_j)$$

For *heterogeneous* populations, we can further analyze the diversity by attempting to cluster groups of agents: this refers to the individual’s degree of *speciation*. It will help us compute the level of specialisation and distribution of specialities within the population. We can define the fixed agent roles, mentioned above, as:

- *explorer* = (1, 0, 0)
- *extractor* = (0, 1, 0)
- *transporter* = (0, 0, 1)

Then, for every agent in the population, we compute its distance to each fixed agent role and “cluster” the agents around these pre-defined roles.

Because we are interested in the notion of specialisation within a population, we need a way to measure the similarity between any two agents. We also need a means for describing the level of “sameness” amongst all the agents in a population — this is a form of *parity*. We use the Euclidean distance to determine the similarity between any two agents.

$$d_{i,j} = \sqrt{(v_i - v_j)^2 + (x_i - x_j)^2 + (t_i - t_j)^2} \quad (1)$$

In order to examine the parity of an entire population, we perform pairwise comparisons for all the agents, filling in the upper triangular distance matrix:

$$\forall i, j : i > j, M_{i,j} = d_{i,j} \quad (2)$$

The larger the values in this matrix, the more disparity, or *heterogeneity*, there is in the population. We can use standard statistical techniques (e.g., mean and standard deviation) on the set of values in M in order to compare different populations.

Our experiments with the *mindScape* environment have involved crafting and evaluating various populations. We have been examining different *role ratios*, i.e., the relative number of explorers, extractors and transporters in the population. In addition, since the capacities v , x and t can be real-valued, we are considering populations where agents’ roles are less well defined and may not be singular; for example, an agent’s capacity values could be (0.99, 0.78, 0.12), indicating that the agent has strong abilities for exploration and extraction, but not for transportation. Initial work evaluated hand-coded populations, where the role ratios and capacity values were set manually. More recent work has been exploring the use of *evolutionary computation* to adapt the population in two ways. First, the role ratios within the population can change over time, allowing us to evolve populations with varying ratios of explorers, extractors and transporters. Second, the capacities of individuals can change through reproduction: if we allow an explorer and an extractor to “mate” (using an evolutionary algorithm), they may produce an offspring with hybrid capabilities. These features allow us to experiment with heterogeneity not only in terms of role allocation, but also in terms of evolved species. The evolutionary paradigm is a natural choice, since its rules for recombination provide an easy way to explore many closely related populations.

Discussion

The mindScape prototype is currently undergoing benchmark testing. We have compared the results for a set of hand-crafted role ratios, for example see Table 2. Preliminary results have revealed several issues.

Table 2: definition of sample benchmark experiments

experiment number	interaction mechanism	ratio of explorers	ratio of extractors	ratio of transporters
1	BROADCAST	0.33	0.33	0.34
2	TRAIL	0.33	0.33	0.34
3	BROADCAST	0.68	0.14	0.14
4	TRAIL	0.68	0.14	0.14
5	BROADCAST	0.14	0.68	0.14
6	TRAIL	0.14	0.68	0.14
7	BROADCAST	0.14	0.14	0.68
8	TRAIL	0.14	0.14	0.68
9	BROADCAST	0	0.5	0.5
10	TRAIL	0	0.5	0.5
11	BROADCAST	0	0.75	0.25
12	TRAIL	0	0.75	0.25
13	BROADCAST	0	0.25	0.75
14	TRAIL	0	0.25	0.75

First, the standard deviation of results is typically quite high. Current benchmark experiments count the number of resources delivered to home and the average energy level of agents at the end of a run. The average standard deviation value is approximately 20% relative to the mean.

Second, preliminary experiments, such as shown in Figure 4, result in very little difference between the two interaction mechanisms. As a result, we are tuning functions that determine decay of environmental cues and broadcast range in order to better understand these parameter spaces.

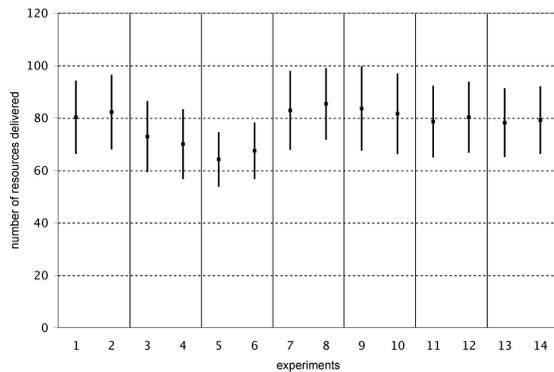


Figure 4: sample benchmark tests

Third, the definition of rules for energy consumption and renewal within the system is crucial. Inappropriate settings produce unexpected results. Transporters are the agents that collect the energy reward when delivering resources. The question is: how should those rewards be distributed to the other agents that helped discover and extract the resources? Essentially, this is a credit assignment problem. Within the evolutionary computing community, there exists a large body of literature examining various approaches, e.g., (Wilson 1994). In addition to evaluating the application of these

approaches to our scenario, we are also exploring a market-based approach to evolving appropriate functions for setting energy levels. This work involves comparing (1) a two-tiered market in which explorers interact with extractors, and extractors interact with transporters; versus (2) a single-tiered market in which explorers and extractors both interact with transporters.

Once benchmark testing is complete, a full set of experiments will be conducted along the two axes (interaction mechanisms and population diversity). Analysis will examine relationships between complexity levels of interaction and speciation vs diversity levels of heterogeneity within evolved populations. In future work, we aim to move the results from the simulated domain onto real robots. Plans are underway for a series of experiments using 50 LEGO Mind-storm robots, following the rules and environmental structure defined in mindScape.

Acknowledgments

This work was partially supported by NSF #IIS-03-29037.

References

- Balch, T., and Arkin, R. 1995. Motor-schema based formation control for multiagent robot teams. In *Proceedings on the First International Conference on Multiagent systems*.
- Balch, T. 1998. *Behavioral Diversity in Learning Robot Teams*. Ph.D. Dissertation, Georgia Institute of Technology.
- Besher, S., and Fewell, J. 2001. Models of division of labor in social insects. *Annual Review of Entomology* 46:413–440.
- Cao, Y. U.; Fukunaga, A. S.; and Kahng, A. B. 1997. Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots* 4:1–23.
- Dias, M. B., and Stentz, A. 2000. A free market architecture for distributed control of a multirobot system. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems*.
- Dorigo, M.; Maniezzo, V.; and Coloni, A. 1996a. The Ant System: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics Part B: Cybernetics* 26(1):29–41.
- Dorigo, M.; Maniezzo, V.; and Coloni, A. 1996b. The Ant System: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics Part B: Cybernetics* 26(1):29–41.
- Eiben, A.; Nitschke, G.; and Schut, M. 2005. Evolving an agent collective for cooperative mine sweeping. In *Proceedings of the IEEE Conference on Evolutionary Computation*.
- Farinelli, A.; Iocchi, L.; and Nardi, D. 2004. Multirobot systems: A classification focused on coordination. *IEEE Transactions on Systems, Man and Cybernetics*.
- Fierro, R.; Song, P.; Das, A.; and Kumar, V. 2002. *Cooperative Control of Robot Formations*, volume 66. Kluwer Academic Press. 73–93.

- Friedman, D. 1993. The double auction institution: A survey. In Friedman, D., and Rust, J., eds., *The Double Auction Market: Institutions, Theories and Evidence*, Santa Fe Institute Studies in the Sciences of Complexity. Cambridge, MA: Perseus Publishing. chapter 1, 3–25.
- Gerkey, B., and Mataric, M. 2002. Sold!: Auction methods for multirobot coordination. *IEEE Transactions on Robotics and Automation* 18(5).
- Golfarelli, M.; Maio, D.; and Rizzi, S. 1997. A task-swap negotiation protocol based on the contract net paradigm. Technical report, CSITE, Research Center for Informatics and Telecommunication Systems, Universita di Bologna, Italy.
- Jones, C., and Mataric, M. 2003. Adaptive division of labor in large-scale multi-robot systems. In *Proceeding of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS-03)*, 1969–1974.
- Kalra, N.; Stentz, A.; and Ferguson, D. 2004. Hoplitest: A market framework for complex tight coordination in multi-agent teams. Technical Report CMU-RI-TR-04-41, Robotics Institute, Carnegie-Mellon University.
- Kowalczyk, W. 2001. Multi-robot coordination. In *Second Workshop on Robot Motion and Control*.
- Krieger, M.; Billeter, J.-B.; and Keller, L. 2000. Ant-like task allocation and recruitment in cooperative robots. *Nature* 406(6799):992–995.
- Kube, C. R., and Zhang, H. 1992. Collective robotic intelligence. In *1992 International Conference on Simulation of Adaptive Behaviour*, 460–468.
- Lerman, K., and Galstyan, A. 2003. Marcospocia analysis of adaptive task allocation in robots. In *IEEE International Conference on Intelligent Robots and Systems*.
- Luke, S. 1998. Genetic programming produced competitive soccer softbot teams for robocup97. In Koza, J. R.; Banzhaf, W.; Chellapilla, K.; Deb, K.; Dorigo, M.; Fogel, D. B.; Garzon, M. H.; Goldberg, D. E.; Iba, H.; and Riolo, R., eds., *Genetic Programming 1998: Proceedings of the Third Annual Conference*, 214–222. University of Wisconsin, Madison, Wisconsin, USA: Morgan Kaufmann.
- Parker, L. E. 1998. Alliance: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*.
- Parkes, D. C., and Shneidman, J. 2004. Distributed implementations of vickrey-clarke-groves mechanisms. In *Proceedings of the Third International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS)*, 261–268.
- Potter, M.; Meeden, L.; and Schultz, A. 2001. Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists. In *Proceedings of The Seventeenth International Conference on Artificial Intelligence*, 1337–1343. Morgan Kaufmann.
- Saffiotti, A.; Zumel, N. B.; and Ruspini, E. H. 2000. Multi-robot team coordination using desirabilities. In *Proceedings of the 6th International Conference on Intelligent Autonomous Systems (IAS)*, 107–114.
- Sandholm, T. 1993. An implementation of the contract net protocol based on marginal cost calculations. In *Proceedings of the 12th International Workshop on Distributed Artificial Intelligence*.
- Sandholm, T. 1998. Contract types for satisficing task allocation: II experimental results. In *AAAI Spring Symposium Series: Satisficing Models*.
- Shoham, Y., and Tennenholtz, M. 1992. On the synthesis of useful social laws for artificial agent societies. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92)*.
- Simmons, R.; Apfelbaum, D.; Burgard, W.; Fox, D.; Moors, M.; Thrun, S.; and Younes, H. 2000. Coordination for multi-robot exploration and planning. In *Proceedings of the 17th National Conference on Artificial Intelligence*.
- Smith, R. G., and Davis, R. 1980. Frameworks for cooperation in distributed problem solving. *IEEE Transactions on Systems, Man and Cybernetics* 11(1).
- Smith, R. G. 1977. The contract net: a formalism for the control of distributed problem solving. In *Proceedings of the Fifth International Joint Conference on Artificial Intelligence (IJCAI-77)*.
- Stone, P., and Veloso, M. 1998. Communication in domains with unreliable, single-channel, low-bandwidth communication. In *International Joint Conference on Artificial Intelligence*.
- Vaughan, R.; Stoey, K.; Sukhatme, G.; and Mataric, M. 2000. Blazing a trail: insect-inspired resource transportation by a robot team. In *Proceedings, 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, 111–120.
- Watson, R. A.; Ficici, S. G.; and Pollack, J. B. 1999. Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In *1999 Congress on Evolutionary Computation*, 335–342. IEEE Press.
- Weiss, G. 1999. *Multiagent Systems: A Modern Approach to Distributed Artificial Intelligence*. MIT Press.
- Werger, B. B., and Mataric, M. 2000. Broadcast of local eligibility for multi-target observation. *Distributed Autonomous Robotic Systems (DARS)* 347–356.
- Wilensky, U. 2002. Modeling nature’s emergent patterns with multi-agent languages. In *Proceedings of EuroLogo 2002*.
- Wilson, S. W. 1994. ZCS: A zeroth level classifier system. *Evolutionary Computation* 2(1):1–18.
- Yanco, H., and Stein, L. 1993. An adaptive communication protocol for cooperating mobile robots. In *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, 478–485. MIT Press.