

The Impact of Diversity on Performance in Multi-robot Foraging*

Tucker Balch[†]

College of Computing
Georgia Institute of Technology
Atlanta, Georgia 30332-0280
trb@cs.cmu.edu

Abstract

Quantitative relationships between performance and behavioral diversity are investigated in a multirobot foraging task. The task, referred to as *multi-foraging*, requires robots to collect different types of object and deliver them to different locations according to type. Multi-foraging was selected for investigation because it offers even more opportunities for agent specialization than simpler foraging tasks. Three team foraging strategies are evaluated: *homogeneous*, where each agent is capable of delivering all types of object; *specialize-by-color*, where each robot specializes in collecting one type of object; and *territorial*, where most of the robots drop objects off near the delivery area, while the remaining agent completes the sorting and delivery. Each strategy is evaluated for diversity and performance using quantitative metrics. Data is gathered in thousands of simulation runs and the behaviors are also verified on mobile robots. In contrast to the results of a similar study in robotic soccer [12], the results of this research indicate homogeneous behavior is the best strategy for foraging robot teams.

1 Introduction and background

An important issue in multiagent robotics research is the question of similarity between the agents on a team. Most research in multirobot teams has centered on homogeneous systems, with work in heterogeneous systems focused primarily on mechanical and sensor differences (such as Parker’s work [18]). But teams of mechanically identical robots are also interesting because they may be homogeneous or heterogeneous depending only on agent behavior. Recent investigations indicate that *behaviorally* heterogeneous systems offer advantages in some tasks [12, 15]. A study of robotic soccer, for instance, shows that diversity is important and is strongly correlated with performance [12]. Does this hold for all multiagent tasks? To address this question we investigate the utility of behavioral diversity in foraging robot teams.

*Agents ’99, Seattle WA.

[†]Current address: Computer Science Department, Carnegie Mellon University, Pittsburgh, PA 15213-3891

Foraging has a strong biological basis. Many ant species, for instance, perform the forage task as they gather food. Foraging is also an important subject of research in the mobile robotics community; it relates to many real-world problems [3, 4, 10, 16, 15]. Among other things, foraging robots may find potential use in mining operations, explosive ordnance disposal, and waste or specimen collection in hazardous environments (the Mars Pathfinder rover for example).

At Georgia Tech, Arkin and Balch have investigated several behaviorally homogeneous strategies for robot foraging. [3, 4, 11]. Their work specifically investigates the impact of communication on performance in foraging teams. Although communication is not used here, the research is extended here to include a more complex foraging task and several new team strategies (including heterogeneous approaches).

Fontan and Mataric have investigated a territorial heterogeneous foraging strategy where the search area is equally divided between agents [15]. Robots hand off collected attractors from area to area, with the last agent completing delivery to the homebase. Their work indicates that performance degrades if the number of robots is increased beyond a certain maximum.

Drogoul evaluated several homogeneous foraging strategies in simulation [14]. His research investigates the utility of laying “crumbs” as path markers for other agents. The idea was inspired by the technique of laying chemical trails to food sources utilized by many ant species [17]. One factor impacting performance is the degree to which the robots interfere with one another. In the most efficient “crumb-laying” foraging strategy, performance is reduced when the number of agents exceeds a particular mark. To address this, a “docker” behavioral strategy is explored. The docker robots are able to pass attractors from one to another while remaining in a fixed position. In robot simulations using this behavior, spontaneous chains of agents arise. Instead of carrying attractors back to the base individually, they hand them from one to another in the chain. Performance in foraging is maximized in the docker strategy. The key drawback to this approach is the mechanical challenge of building agents able to accomplish such hand offs.

In the work most closely related to this research Goldberg and Mataric propose *interference* as a metric for evaluating a foraging robot team [16]. Interference refers to the situation where two robots attempt to occupy the same place at the same time; it is measured as the amount of time agents spend avoiding one another. Since interference may reduce the efficiency of a robot team, Goldberg suggests pack and caste arbitration as mechanisms for generating efficient be-


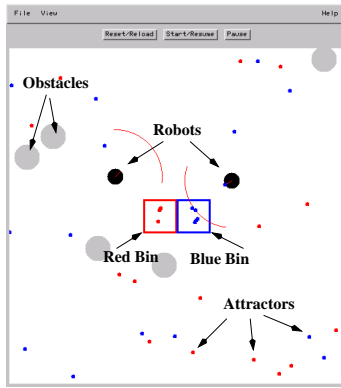


Figure 1: Multi-foraging robots in simulation (left) and in the laboratory (right). In the simulation robots are represented as black circles; arcs indicate the visual sensing range; obstacles are drawn as gray circles; the small discs are attractors. The robots deliver the attractors to the color-coded squares representing delivery areas.

havior and reducing interference. In the pack scheme, each agent is arbitrarily assigned a place in the “pack hierarchy.” Agents higher in the hierarchy deliver attractors before the others. In the caste approach, only one agent completes the final delivery; the other robots leave their attractors on the boundary of a designated “home zone.” Goldberg’s results indicate that interference per unit time is maximized in homogeneous foraging and minimized in pack foraging. In spite of the fact that interference is minimized in the heterogeneous pack systems, homogeneous systems perform best in terms of the number of attractors collected. The caste approach (referred to later as territorial foraging) is adopted as one of the team foraging strategies investigated here.

This research is distinguished from other work in that it is the first to investigate *quantitative* relationships between behavioral diversity and performance in multirobot foraging. To investigate these relationships three multirobot foraging strategies are evaluated in thousands of simulation runs for one to eight robots. Quantitative investigation of behavioral diversity is enabled through the use of a new metric called social entropy [7]. The entropy (diversity) of a multiagent society is calculated based on the number and size of the groups making it up.

This paper provides a description of the multi-foraging task and the development of robotic behaviors for accomplishing it. Three foraging strategies are evaluated in thousands of simulation runs. The behaviors are also verified through implementation on mobile robots. The robots were evaluated in the Mobile Robot Laboratory at Georgia Tech and at the AAAI-97 Mobile Robot Competition where they won first place.

2 The multi-foraging task

Most robotic foraging tasks investigated to date involve the collection of attractors of a single type and their delivery to a single destination. This basic task is referred to as *simple foraging*. Simple foraging is an important robotic capability, but many practical industrial and military tasks call for more functionality. Consider, for example, a janitorial robot responsible for collecting and sorting recyclable trash objects into glass, aluminum and paper bins. Similarly, many assembly and construction tasks involve collecting parts or materials and placing them in a specific location. These more complex tasks are referred to as *multi-foraging* tasks. Multi-foraging was selected for this investigation because

it provides more opportunities for behavioral specialization than simple foraging.

In general, the multi-foraging task calls for several types of object to be collected and placed in specific locations according to type. Here *multi* refers to the multiple types of object to deliver, not the number of robots engaged in the task. Examples of simulated and real robots executing a multi-foraging task are presented in Figure 1.

Performance in the multi-foraging task is defined as the number of attractors collected and properly delivered in a fixed time. In terms of the multirobot task taxonomy introduced in [8] this task has the following characteristics: **TIME_LIM** (time-limited), because performance is measured over a fixed period; **RESOURCE_LIM** because as agents collect objects, the availability of attractors is reduced; **OBJECT_BASED** since performance is based on the location of objects, not agents; **COMP_INT** (internally competitive), because robots on the team compete for access to attractors among themselves; **SINGLE_AGENT** since an individual agent can perform positively, even though multiple agents may provide improved performance; and **SENSOR_LIM** since agents only have a limited view of the environment.

In order to compare simulations it’s important to control several factors which affect efficiency:

- **Number of attractors:** Since performance is measured as the number of attractors collected, more attractors available for collection may tend to provide for increased performance in a fixed time trial. In simulation runs there are 40 attractors, 20 of each type (red and blue).
- **Obstacle coverage:** Higher obstacle density can lead to degraded performance because the robots must slow down and/or take a longer route around hazards to deliver attractors. In simulation runs, each playing field includes five 1 m² obstacles (5% coverage). The AAAI Competition field included approximately 10 rock piles varying from about 0.5 m² to 1 m². In most laboratory runs, no obstacles other than another robot and the arena boundaries were present.
- **Playing field size:** Larger search areas may lead to a decrease in performance. In simulation, the field measures 10 by 10 meters. At the AAAI Competition, the field was a hexagon measuring approximately 8 by 8 meters. Runs in Georgia Tech’s Mobile Robot Lab were conducted in a 5 by 10 meter area.
- **Number of robots:** In most cases, increasing the number of robots on a team improves performance. There is some concern however, that as the number of robots increases,

interference between the agents will degrade performance [16]. In simulation experiments the number of agents is varied from one to eight. In laboratory runs one and two agents were used. At the AAAI Competition two robots were employed.

The next section explains the development of multirobot behaviors for the multi-foraging task.

3 Behaviors for foraging

A motor schema-based reactive control system is used for robot programming [2]. In this approach, the agent is provided several pre-programmed behavioral assemblages that correspond to steps in achieving the task (e.g. *wander*, *acquire*, *deliver*, and so on [1]). The behavioral assemblages are in turn composed of more primitive behaviors called motor schemas. Binary perceptual features (also referred to as perceptual triggers) are used to sequence the robot through steps in achieving the task. Three strategies for foraging are implemented and evaluated for performance and diversity:

- **Behaviorally homogeneous:** all the robots collect all the different types of attractor and deliver them to corresponding color-coded delivery areas.
- **Territorial:** In this scheme one robot, referred to as the *sorting* agent, is responsible for collecting attractors within a three meter circle centered on the homebase. The other agents search at a distance from the homebase. When these robots find an attractor, they drop it off at the boundary of the “home zone.” Final delivery is then completed by the sorting agent. This is a behaviorally heterogeneous approach.
- **Specialize-by-color:** half the robots specialize in collecting one type of attractor while the rest specialize in collecting the second type. Specialization by color is a heterogeneous strategy as well.

To ensure a fair comparison between the various foraging strategies a fixed repertoire of behaviors is utilized across all implementations. A range of behaviors were developed to support several foraging strategies and to avoid bias towards any particular approach. The repertoire is suitable for building behaviorally homogeneous foraging teams as well as territorial and other heterogeneous strategies. Agents utilizing distinct strategies differ in the order they activate behaviors. The behaviors developed for foraging teams are summarized below:

- *wander*: move randomly about the environment in search of attractors. Upon encountering an attractor agents automatically transition to an appropriate *acquire* behavior.
- *stay_near_home*: similar to the *wander* assemblage but with an additional attractive force to keep the agent close to the homebase. This assemblage is utilized in the territorial strategy by a sorting agent.
- *acquire_red*: move towards the closest visible red attractor. When close enough to grasp the attractor, the agent closes its gripper and transitions to a *deliver* assemblage.
- *acquire_blue*: move towards the closest visible blue attractor.
- *deliver_red*: move towards the red delivery area. When close enough to deposit the attractor in the delivery area, the robot opens its gripper and transitions to one of the *wander* assemblages. Territorial agents are programmed to drop attractors on the boundary of the home zone.
- *deliver_blue*: move towards the blue delivery area.

The reader is referred to [8] for a more complete description of these behaviors.

Recall that behaviors are sequenced using perceptual features. A perceptual feature is a single, abstracted bit of

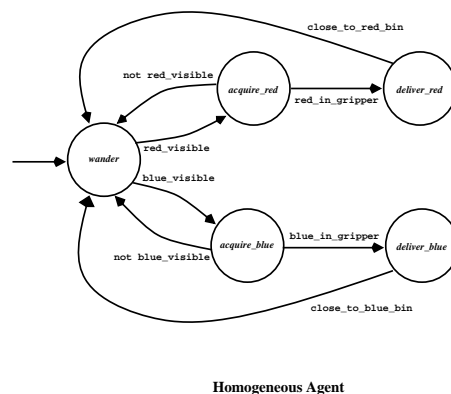


Figure 2: An FSA representing the homogeneous foraging strategy. This kind of agent can collect both types of attractor.

environmental or sensor state germane to the robot’s task (whether or not the robot is holding an attractor in its gripper for example). Robots must decide on the basis of these environmental cues which behavior to activate at each point in time. The robots are programmed as Finite State Automata (FSAs) that sequence from one state to another. Each state of the FSA determines which behavior is to be activated, with transitions between behavioral states triggered when particular perceptual features are activated. This approach is called *perceptual sequencing* [1].

A fixed set of perceptual features are utilized across all implementations to ensure a fair comparison between the various foraging systems. The perceptual features for foraging are cataloged in Table 1. In addition to the features advising the robot whether an attractor is visible, there are also features indicating whether attractors are visible outside the home zone. The visibility cues are used to allow territorial agents to search for attractors inside or outside the zone while ignoring the others. The *close_to_homezone* feature is used by territorial robots as a signal to drop an attractor on the boundary of the zone so that a sorting robot can complete the final delivery.

4 Behavioral sequencing strategies

An obvious approach to the design of a multirobot multi-foraging team is to program each agent to complete the entire task on its own. This strategy is referred to as *homogeneous* because all the agents are programmed with the same behavior. This approach was used by Georgia Tech’s robots in the AAAI-94 and AAAI-97 competitions and in research concerning multiagent communication [11, 10]. The homogeneous approach provides fault-tolerance because when one or more agents fail, the remaining robots can still accomplish the task.¹

An FSA representing the homogeneous strategy is shown in Figure 2. Each agent begins with the *wander* behavior activated, roaming about the environment in search of attractors. When a robot encounters an attractor, either the *red_visible* or *blue_visible* perceptual feature is triggered, causing the agent to transition to the corresponding *acquire_red* or *acquire_blue* behavior. Upon capturing an attractor, a robot returns back to homebase using one of the

¹In fact, in both competitions one of the robots failed. Fortunately, the remaining robots were able to complete the task.

perceptual feature	meaning
red_visible	a red attractor is visible.
blue_visible	a blue attractor is visible.
red_visible_outside_homezone	a red attractor is visible outside the three meter radius home zone.
blue_visible_outside_homezone	a blue attractor is visible outside the home zone.
red_in_gripper	a red attractor is in the gripper.
blue_in_gripper	a blue attractor is in the gripper.
close_to_homezone	the agent is within 3 meters of the homebase.
close_to_red_bin	close enough to the red delivery area to drop an attractor in it.
close_to_blue_bin	close enough to the blue delivery area to drop an attractor in it.

Table 1: Perceptual features available to the foraging robots. Each feature is one bit of environmental state; the entire perceptual state is a nine-bit value.

deliver behaviors. Finally, upon reaching the corresponding delivery area, the agent drops the attractor and transitions back to *wander*.

When several robots simultaneously attempt to deliver attractors to the same delivery area, they may interfere with one another and degrade performance (see Figure 6 for an example). One way to reduce interference and potentially improve performance is to partition the task so that responsibility for collecting red and blue attractors is divided among the robots. Half of the agents are responsible for collecting the red objects and the other half for blue. This way, the chance of a “traffic jam” at either delivery area is reduced.

FSA for these specialized agents are illustrated in Figure 3. All agents start with the *wander* behavior activated. They begin to search the environment for attractors. What follows depends on whether the agent is a red specialist or a blue specialist. Red specialists ignore blue attractors, but when they encounter a red attractor while in the *wander* phase, they immediately transition to the *acquire_red* behavior. Similarly, blue specialists ignore red attractors. After acquiring an attractor, the agents deliver it to the appropriate delivery area using one of the *deliver* behaviors, then they switch back to *wander* to search for new items. An additional transition is provided for situations where an agent loses sight of the attractor. In that event it transitions back to *wander*.

Another way to reduce interference near the delivery bins is to assign one robot to the sorting task while other robots collect the attractors and drop them nearby. This approach, initially investigated by Goldberg and Mataric is adopted for this investigation as well [16]. In this scheme, referred to as territorial foraging, one robot remains close to the homebase in the “home zone,” delivering attractors that other agents deposit on the zone’s boundary. Figure 4 shows the FSAs for robots in this system. The sequencing strategies for the agents are similar to the approach for homogeneous foragers (Figure 2). One significant difference is that the sorting agent utilizes a *stay_near_home* behavior rather than *wander* while searching for attractors. This results in the agent staying close to the delivery areas for sorting. The roaming agents are also similar to the homogeneous strategy, except they are triggered to drop attractors at the boundary of the home zone instead of depositing them in the delivery areas.

5 Performance in simulation

Now the performance of the foraging systems are examined experimentally. The JavaBots system was utilized for simulation and mobile robot experimentation [9, 5]. Behaviors coded in JavaBots may be run in simulation, and without modification, on Nomadic Technologies’ Nomad 150 mobile robots. Statistical results are gathered by running the robot behaviors in thousands of simulation trials.

In simulation, each robot is a kinematically holonomic vehicle (a simulated Nomad 150) which is controlled by one of the behavioral systems described above. Simulated motor and sensor capabilities are based on performance of the physical robots. The robots can detect hazards with sonar out to a range of nine meters. Attractors can be detected visually out to three meters across a 90 degree field of view.

Each type of control system under investigation was evaluated using one to eight simulated robots in five different randomly generated environments. The environments differ in the arrangement of hazards and attractors. The arena is 10 by 10 meters and includes five randomly placed 1 m² obstacles for 5% obstacle coverage. There are 20 each of red and blue attractors distributed about the environment for collection. 100 trials were run in each environment, or 4,000 runs for each control strategy, and 12,000 total.

Time is measured in seconds. Since reactive control systems are very fast, several thousand control cycles are completed each second. The simulation is allowed to proceed faster than real time with each control cycle fixed at 200 milliseconds (simulation time). Each trial runs for 10 simulated minutes or 30,000 control cycles.

Average performance for each of the three systems is plotted versus the number of robots per team in Figure 5. Performance is measured as the total number of attractors properly delivered by the team in 10 minutes; larger numbers are better, with 40 being the best possible. The plotted values are determined by computing the average performance of the teams in each of the five randomly generated environments over 100 trials.

In all cases, performance increases monotonically with the number of agents. The data also show that regardless of team size, the homogeneous strategy performs best, followed by the territorial method and finally the specialize-by-color approach. In the case of territorial systems with one robot, the single agent is programmed as a sorting robot. It is able to collect only the attractors placed nearby. Of the

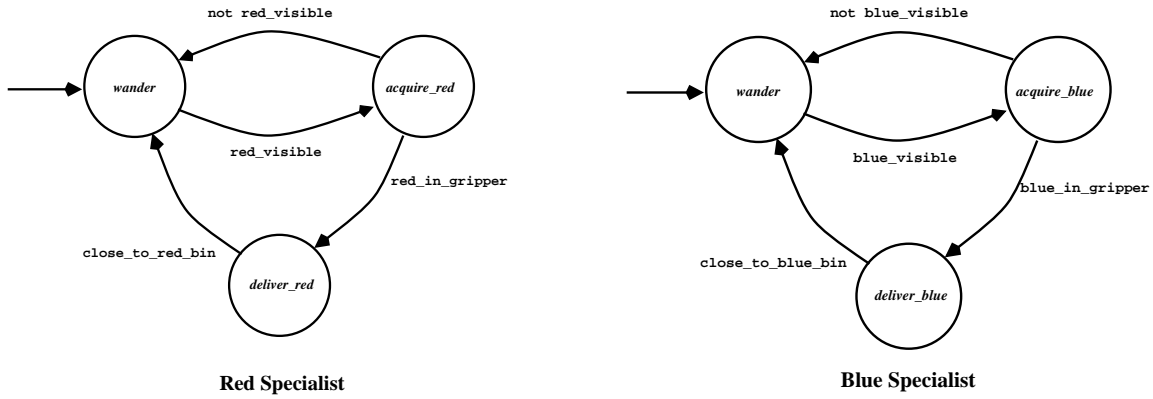


Figure 3: FSAs representing specialized behaviors for foraging. The FSA on the left shows the sequence behaviors are activated in for an agent specializing in collecting red attractors. The FSA on the right shows the sequence for blue specialists.

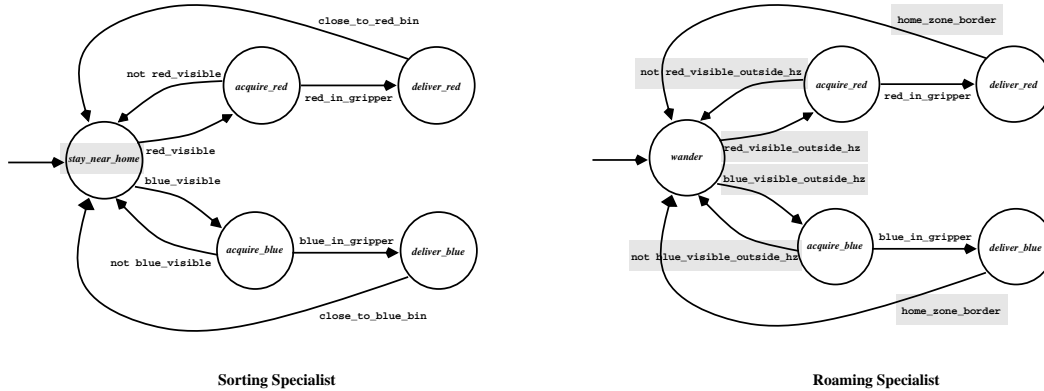


Figure 4: Territorial behavioral sequences for foraging. The FSA on the left shows the sequence of behaviors for an agent that remains close to the homebase, completing the delivery of the attractors. Agents using the strategy on the right search for attractors away from the home zone and deliver them to the home zone boundary. Differences from the homogeneous strategy are highlighted.

foraging strategies evaluated homogeneous systems perform best in this foraging task. These results confirm those of other researchers in simple foraging [15, 16].

As previously mentioned, inter-robot interference is a concern for homogeneous systems, while the heterogeneous strategies were specifically designed to reduce interference. Interference does occur during the delivery phase in both homogeneous and specialize-by-color strategies (Figure 6, left). This study does not include a quantitative measure of interference, but qualitative observations are consistent with results reported in [16]; namely that there is more frequent interference between agents in the homogeneous strategy than in other systems. Still, overall performance is best in homogeneous systems.

In the case of territorial foraging, robot-robot interference occurs much less frequently, but another factor limits performance. In most trials, the roaming agents quickly deliver a large number of the attractors to the boundary of the home zone, but the single sorting robot cannot always keep up. In simulations with seven and eight agents it is not uncommon for undelivered attractors to remain in a ring around the home zone at the end of the trial (Figure 6, center). Even though the number of sorters is constant (one), the territorial foraging experiments illustrate how the ratio of sorters to roamers impacts performance. The ratio varies from 1:0 to 1:7 as the number of agents goes from 1 to 8. In

the 1:0 case, the sorting agent “starves” because it quickly finds all the nearby attractors. Conversely, in the 1:7 case, the sorter is overworked; there are many more attractors for it to deliver than it is able to. Note that regardless of the ratio, territorial foragers never perform better than homogeneous strategy.

A different sort of problem crops up for the specialize-by-color teams. Towards the end of trials for these agents there are often uncollected red attractors on the right side of the field and uncollected blue attractors on the left (Figure 6, right). This occurs because the agents inadvertently segregate themselves geographically to the left or right depending on whether they collect red or blue attractors. After delivering a blue attractor, for instance, a blue-collecting agent is more likely to remain on the same (right) side of the field as the blue bin. Because of this there are no red-collecting agents on the right side of the field and vice-versa. In large robot teams the robot-robot repulsion employed as part of the *wander* behavior serves as an additional force preventing the agents from diffusing from one side to the other. In addition to segregation, the specialize-by-color agents occasionally interfere with one another when delivering attractors to the same delivery area simultaneously. These factors combine to drive performance down in the specialize-by-color teams.

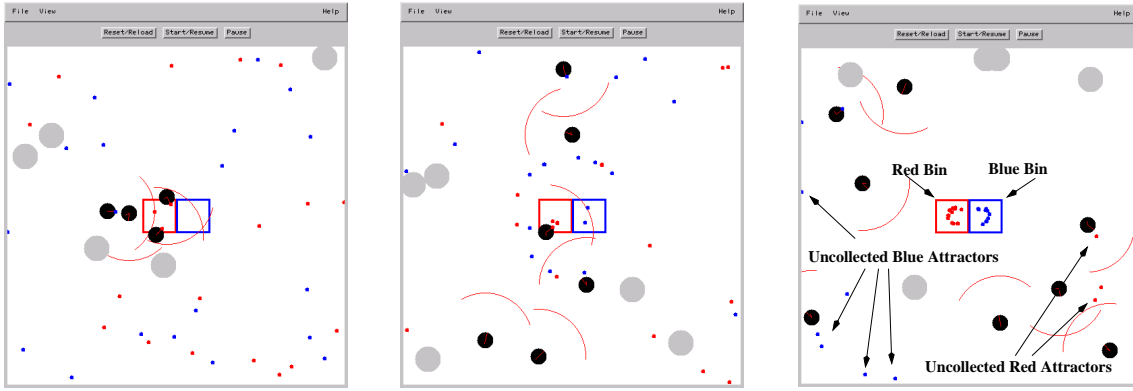


Figure 6: Simulations highlighting some of the factors that impact performance in foraging teams. From left to right: interference in a homogeneous team, an over-worked sorting robot, attractors left uncollected by agents specializing in one color of attractor.

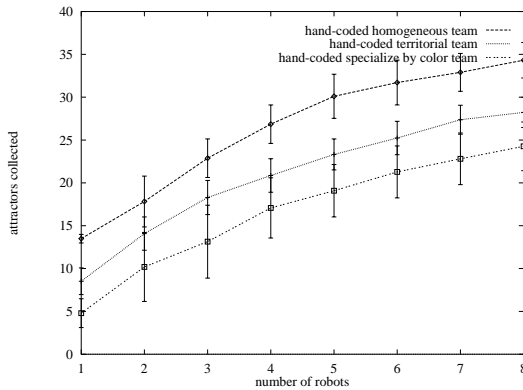


Figure 5: Performance of the foraging teams versus size of the team. Higher numbers indicate better performance; error bars indicate 95% confidence intervals. The homogeneous teams perform best in all cases.

6 Diversity

Previously, diversity in multirobot teams was evaluated on a bipolar scale with systems classified as either *heterogeneous* or *homogeneous*, depending on whether any of the agents differ [15, 16, 18]. Unfortunately, this labeling doesn't tell us much about the *extent* of diversity in heterogeneous teams.

Heterogeneity is better viewed on a sliding scale providing for quantitative comparisons. Such a metric enables the investigation of issues like the impact of diversity on performance, and conversely, the impact of other task factors on diversity. *Social entropy*, inspired by Shannon's information entropy [20], is introduced as a measure of diversity in robot teams. The metric captures important components of the meaning of diversity, including the number and size of groups in a society. Social entropy is briefly reviewed here. For more details, readers are referred to [13].

To evaluate the diversity of a multirobot system, the agents are first grouped according to behavior (e.g. all red-collecting agents are placed in one group). Next, the overall system diversity is computed based on the number and size of the groups. Social entropy for a multirobot system com-

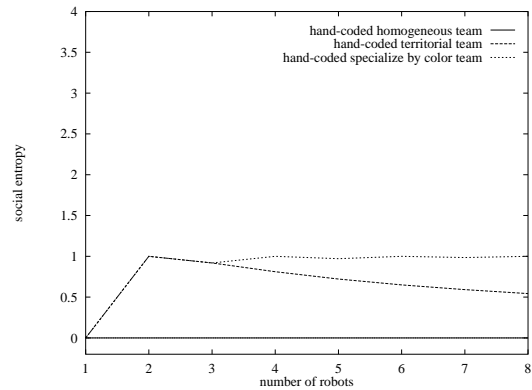


Figure 7: Diversity, as measured by social entropy, versus size of the team for foraging teams. Larger numbers indicate greater diversity.

posed of M groups is defined as:

$$H(X) = - \sum_{i=1}^M p_i \log_2(p_i) \quad (1)$$

where p_i represents the proportion of agents in group i . We will use this metric in the evaluation of the experimental foraging strategies.

The diversity of the three experimental systems is plotted versus the size of the team in Figure 7. Diversity is measured using the social entropy metric introduced in [7]. The entropy of a system is determined by grouping the agents according to behavior, then evaluating for diversity based on the number and size of the groups (a Java-based social entropy calculator is available online [6]). The entropy measure is analogous to the randomness of the system; greater entropy indicates greater diversity. The homogeneous teams always exhibit zero diversity, while the heterogeneous teams vary in entropy from 0.54 to 1.0.

The territorial system always contains one unique robot (the sorting agent that stays near homebase), while the rest are identical. In this system, entropy is 1.0 for two agents, then gradually declines as the number of agents increases. The value drops to 0.54 for eight robots.

For even numbers of robots the specialize-by-color team consists of half red-collecting robots and half blue-collecting

robots; this equates to an entropy of 1.0. For odd numbers of robots the entropy is slightly less than 1.0, but as the number of agents grows it approaches unity. This is illustrated in the graph (Figure 7).

One goal of this research is to determine the relationship between diversity and performance in multiagent tasks; is the relative diversity of two multiagent teams a predictor of their relative performance? This question is addressed using *Spearman's Rank-Order Correlation Test* [19]. Spearman's test measures correlation between pairs of data points, where each pair reflects the ranking of each item according to separate metrics. In this case, we compare ranking in performance with ranking in diversity. The correlation value, r ranges from -1 (negatively correlated), to 0 (uncorrelated) to 1 (positively correlated). Statistical significance of the correlation is indicated by the probability that the same correlation could have occurred by chance.

Consider the plots of robot system performance and diversity in Figures 5 and 7. For each robot team size (> 1) the systems can be ranked by diversity and performance.² Spearman's test indicates that **diversity and performance are strongly negatively correlated in this foraging task**, with $r = -0.92$. Greater diversity is associated with lower performance. The probability of the null hypothesis, that the rankings occur by chance, is 0.000043 .

7 Implementation on mobile robots

To verify the simulation results, the foraging behaviors were ported to Nomad 150 mobile robots. Because the control systems are implemented in JavaBots, they can run in simulation and on hardware; the same behaviors and features can be utilized on mobile robots as in simulation. The homogeneous strategy was run on one and two mobile robots in the Mobile Robot Laboratory. The homogeneous and specialize-by-color strategies were also employed by Georgia Tech at the AAAI-97 competition.

Two robots executing the homogeneous foraging strategy are illustrated in Figure 1. In this set of experiments the robots utilize a passive gripper to collect attractors. The gripper is designed so that a captured object remains under the robot's control until the robot drops it by backing up.

The Mobile Robot Laboratory provides an arena measuring approximately 5 meters by 10 meters for the robot experiments. A total of 20 attractor objects, 10 of each type (red and blue), were distributed randomly about the lab for each trial. Both the size of the arena and the number of attractors available for collection are halved in comparison with the environment used in simulation experiments.

Five trials of 10 minutes were run for one and two robots. At the end of each trial, performance was evaluated as the total number of attractors properly delivered. Performance is summarized in Table 7. Qualitative behavior was essentially identical to that of homogeneous teams in simulation. As in simulation, the agents occasionally interfered with one another when they deliver attractors to the same bin. In these experiments with 20 attractors, each robot routinely collected and delivered 8 objects. As expected, two robots perform better than a single robot. Performance is slightly worse than the same strategy in simulation experiments. The decrease is most likely due to the reduced number of attractors available for collection (20 versus 40).

²Since diversity has no meaning for a single agent system, only teams of two or more agents are considered. Ties are declared in cases where values are exactly the same or when confidence intervals overlap.

The homogeneous and specialize-by-color strategies were also used on Georgia Tech's robots at the AAAI-97 Robot Competition. Rather than a delivery area for each type of attractor as in the task described above, in the AAAI task robots must deposit attractors in bins with doors. The doors of the bins are painted an identifying color to help the agents find them. The robots were equipped with active grippers to enable them to lift the objects and drop them in the bins. Behaviors were modified slightly to accommodate the differences in the task (full details are provided in [8]).

In the first round of the competition the robots had difficulty detecting rock hazards. The sonar sensors were not effective at detecting the hazards because they are mounted too high on the robot to detect the shorter rocks (this problem was corrected later by aiming the sensors downward). The robots encountered the hazards on several occasions. In one case, one of the robot's grippers was ripped off the vehicle (fortunately this occurred towards the end of the trial). Despite this setback, the agents were able to deliver a significant number of attractors and win the first round. One of the robots even captured a moving squiggle ball – this was a rare event at the competition.

A change in the task for the final round presented an additional challenge. The robots had to collect and deliver objects of six different colors instead of two as in the first round. This was a problem because the vision systems can only track *three* colors, and at least one of those tracking channels has to be dedicated to detection of the delivery door. This deficiency was addressed by exploiting a heterogeneous foraging strategy. Each robot was programmed to specialize in the collection of three of the six types of attractor.

In the final round the robots picked up 10 attractors and placed 9 of them in the correct delivery bin. Georgia Tech's robots won the competition. The success of a behaviorally heterogeneous team in this situation illustrates how the computational limits of individual agents can necessitate diversity in a multirobot solution. Each robot is potentially capable of detecting all six types of attractor, but computational limits of the embedded vision computers allow only three at one time — one robot cannot complete the entire task alone. In terms of the taxonomy presented in [8] this instance of the multi-foraging task is `MULTI_AGENT` instead of `SINGLE_AGENT`. Perhaps `MULTI_AGENT` tasks are more likely to require heterogeneous solutions.

8 Conclusion

This paper is the first to report a quantitative link between diversity and performance in foraging robot teams. In these multirobot foraging experiments diversity is negatively correlated with performance; homogeneous teams perform best. The result is in contrast with similar work conducted by the author in robot soccer where diversity is preferred [12]. In both studies conclusions are based on statistical analysis of thousands of simulation trials.

It is likely that differences in the soccer and foraging tasks contribute to the relative advantages of diversity. A key difference is that in soccer it is nearly impossible for a single agent to successfully compete against another team. Conversely, a single foraging robot using a homogeneous strategy could feasibly collect all attractors in a given scenario.

Homogeneous foragers tend to interfere with one another as they deliver attractors to the delivery areas (interference in homogeneous foraging was also noted by Gold-

configuration	attractors collected per trial
1 robot, homogeneous strategy average	10, 8, 11, 12, 8 9.8 +/-2.2
2 robots, homogeneous strategy average	15, 16, 17, 12, 15 15.0 +/-2.3

Table 2: Summary of performance in foraging robot trials. 95% confidence intervals are indicated.

berg in [16]). To address this, territorial and specialize-by-color strategies were designed with the goal of reducing inter-agent interference. Even though interference is reduced in these heterogeneous teams, performance is *worse*. In fact, diversity is negatively correlated with performance in these foraging teams: Spearman's $r = -0.92$ and $\text{prob} = 0.000043$.

In related work Fontan and Mataric investigated a similar territorial heterogeneous foraging strategy [15]. Their work indicates that performance degrades if the number of robots is increased beyond a certain maximum. In contrast, the results in this paper indicate monotonically increasing, but leveling off of performance as the number of foraging agents increases. This difference may be due to: 1) the foraging strategy used by Fontan and Mataric was not reproduced exactly in this work; and 2) performance in the strategies introduced here may actually degrade as the number of agents in this work increases significantly beyond eight robots.

Finally, the behaviors, perceptual features and behavioral sequences used in simulation were also verified on mobile robots. Qualitatively, mobile robot behavior matches that predicted in simulation, including inter-agent interference and overall performance. Robots using these strategies won the AAAI-97 Robot Competition's "Find Life on Mars" event.

9 Acknowledgements

The author thanks Ron Arkin and Chris Atkeson for their many helpful conversations regarding this work. The robots and laboratory space for this research were provided by Georgia Tech's Mobile Robot Laboratory.

References

- [1] R. Arkin and D. MacKenzie. Temporal coordination of perceptual algorithms for mobile robot navigation. *IEEE Transactions on Robotics and Automation*, 10(3):276–286, June 1994.
- [2] R.C. Arkin. Motor schema-based mobile robot navigation. *International Journal of Robotics Research*, 8(4):92–112, 1989.
- [3] R.C. Arkin. Cooperation without communication: Multi-agent schema based robot navigation. *Journal of Robotic Systems*, 9(3):351–364, 1992.
- [4] R.C. Arkin, T. Balch, and E. Nitz. Communication of behavioral state in multi-agent retrieval tasks. In *Proceedings 1993 IEEE Conference on Robotics and Automation*, Atlanta, GA, 1993.
- [5] T. Balch. Javabots. www.cc.gatech.edu/~tucker/JavaBots.
- [6] T. Balch. Online diversity calculator. www.cs.cmu.edu/~trb/java/Dcalc.
- [7] T. Balch. Social entropy: a new metric for learning multi-robot teams. In *Proc. 10th International FLAIRS Conference (FLAIRS-97)*, May 1997.
- [8] T. Balch. *Behavioral Diversity in Learning Robot Teams*. PhD thesis, College of Computing, Georgia Institute of Technology, 1998.
- [9] T. Balch. Integrating robotics research with javabots. In *AAAI-98 Spring Symposium*, March 1998.
- [10] T. Balch and R.C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1), 1995.
- [11] T. Balch, G. Boone, T. Collins, H. Forbes, D. MacKenzie, and J. Santamaría. Io, Ganymede and Callisto - a multiagent robot trash-collecting team. *AI Magazine*, 16(2):39–51, 1995.
- [12] Tucker Balch. Learning roles: Behavioral diversity in robot teams. In *AAAI-97 Workshop on Multiagent Learning*, Providence, R.I., 1997. AAAI.
- [13] Tucker Balch. *Behavioral Diversity in Learning Robot Teams*. PhD thesis, College of Computing, Georgia Institute of Technology, 1998.
- [14] A. Drogoul and J. Ferber. From tom thumb to the dockers: Some experiments with foraging robots. In *From Animals to Animats: Proc. 2nd International Conference on the Simulation of Adaptive Behavior*, pages 451–459, Honolulu, HI, 1994. MIT Press/Bradford Books.
- [15] M. Fontan and M. Mataric. A study of territoriality: The role of critical mass in adaptive task division. In *From Animals to Animats 4: Proceedings of the Fourth International Conference of Simulation of Adaptive Behavior*, pages 553–561. MIT Press, 1997.
- [16] D. Goldberg and M. Mataric. Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings, AAAI-97*, pages 637–642, July 1997.
- [17] B. Holldobler and E. Wilson. *The Ants*. Belknap Press, 1990.
- [18] Lynne E. Parker. *Heterogeneous Multi-Robot Cooperation*. PhD thesis, M.I.T., 1994.
- [19] W. Press, S. Teukolsky, W. Vetterling, and B. Flannery. *Numerical Recipes in C*. Cambridge University Press, 1988.
- [20] C. E. Shannon. *The Mathematical Theory of Communication*. University of Illinois Press, 1949.