

Emergent Bucket Brigading

A simple mechanism for improving performance in multi-robot constrained-space foraging tasks

Esben H. Østergaard, Gaurav S. Sukhatme and Maja J. Mataric
Robotics Research Laboratories
University of Southern California
Los Angeles CA 90089-0781
[esben, gaurav, maja]@robotics.usc.edu

ABSTRACT

This paper is concerned with the multi-robot foraging problem, which has been described as one of the canonical problems for cooperative robotics. In particular, the multi-robot transportation task in constrained space environments is considered. We describe a simple algorithm which produces bucket brigade-like behavior, where the robots "hand off" resources to each other, using only local sensing. The algorithm is tested in five experimental conditions, to empirically evaluate the strengths/weaknesses of the approach, and is found to perform significantly better in constrained space environments than a homogeneous foraging approach. The key results of the described work are: 1) a framework for foraging tasks in general and 2) a simple mechanism that acts locally to produce bucket brigading behavior globally.

1. INTRODUCTION

Foraging is considered one of the "canonical test beds for cooperative robotics" [4]. In this paper, we focus on multi-robot foraging, specifically considering the special case of multiple robots transporting items from a single source to a single sink in a spatially constrained environment.

The multi-robot foraging task have been studied by many researchers. In several of the studies, interference between robots has been reported to be a key issue [8], [5] and [11] and several different mechanisms has been used to reduce negative effects caused by interference [13]. One goal of our work in progress is to find methods to reduce the negative effect of interference in the multi-robot domain, in order to allow for greater scalability with respect to group size.

Previous research has applied a variety of methods to solve different configurations of the foraging task. The goal of the




Figure 1: A two-source, two-sink open field environment with four robots

work we report here is to explore simple strategies that can add richness and flexibility to foraging. In this paper we focus on bucket brigading, a process in which robots hand off a resource from one robot to the next. We empirically study the value of using this strategy in different types of environment.

The rest of the paper is organized as follows. First we give an overview of different aspects of the foraging task and outline a non-hierarchical taxonomy of foraging as a context in which to place previous and current work. Second, we present a set of experiments with a simple local mechanism for achieving global bucket brigading through local drop off and sensing between robots.

2. THE FORAGING TASK

[8] provides a definition for the *collection* task:

The collection task is a two-step repetitive process in which:

1. N ($N \geq 1$) robots search designated region of space for certain objects, and
2. once found, these objects are brought to a goal region using some form of navigation.

In this paper the phrase *foraging* is used synonymously with

→ Increasing problem complexity →		
Single Robot	vs.	Multiple Robots
Single Sink	vs.	Multiple Sinks
Single Source	vs.	Multiple Sources
Open Field Environment	vs.	Constrained space environment
A single type of items to be collected	vs.	Multiple types of items to be collected
Items placed at fixed areas	vs.	Items sprinkled around the environment
Homogeneous robots	vs.	Heterogeneous robots
Without communication	vs.	With communication

Figure 3: Characteristics of foraging tasks

```

While (terminate_condition) not satisfied do {
  1: Go to source area
  2: Search for item
  3: Pick up item
  4: Go to sink area
  5: Drop item
}

```

Figure 2: The basic foraging behavior

collection, as defined above, except that we do not require navigation. In our work, foraging robot(s) transport items from sources area(s) to sinks area(s) in a given environment. An example of the task is shown in Figure 1, where the task is described as multi-robot foraging in an open field-like environment with two sources and two sinks.

The simplest manifestation of the foraging task is a single robot which collects items from a single source and transports them to a single sink in an open field-like environment. The basic foraging behavior is shown in Figure 2. This behavior is the underlying algorithm for much of the work done in the foraging task domain [6], [1], [2], [8], [5], [10] and [11].

In the context of foraging, the transportation problem can be regarded as a special case, namely the foraging task without the need to search for items, that is, without step 2 in the algorithm shown in Figure 2.

2.1 A taxonomy of foraging

Foraging tasks in general can have different manifestations. Most research has addressed the more complex manifestations of the foraging task, to make the control and analysis problems more interesting. The characteristics displayed in Figure 3 identify a set of quantities that, together, determine the nature of the problem to be solved.

The quantities displayed in the Figure are all independent. They are arranged so that characteristics in the left column tend to make the problem simpler, and characteristics in the right column tend to make the problem more complex. The simple case of foraging, where all the characteristics can be found in the left column, is not a very interesting control problem. However, finding efficient control strategies becomes much more complex when studying a case with one of the characteristics from the right column. In cases with multiple characteristics from the right column, the solution space of the problem increases significantly, and so does the

complexity of controller design.

For example, when more complex capabilities are added, different issues arise, such as appropriate load balancing, exploration vs. exploitation, and deciding what to communicate, to whom, and how often.

Issues not captured by the characteristics listed in Figure 3 include the amount of position information is available to the robots, how they find their way to the source(s) and sink(s), and what type of of sensor/motor capabilities they have.

2.2 Related work

This subsection attempts to place many of the previously reported implementations of the foraging task in the context of the taxonomy described above.

In [12] Vaughan et al describe an algorithm for robust transportation of items from a single source to a single sink using multiple robots in a spatially constrained environment.

In [5] Fontan and Mataric manually split up the multi-robot foraging task into a territorial bucket-brigade system, where the sink area for one robot serves as the source area for the one next to it. This approach reduces the problem for the individual robot to the simple case of single-robot single-source single-sink environment, thereby significantly reducing interference. Individual robot failures induce a readjustment of the territories so as to cover the foraging area completely. [6] discusses the resulting critical mass results of the work, showing different levels of effectiveness for 3 different group sizes.

In [2] Balch investigates the “multi-foraging” task in which several different types of items are being collected. The experimental environment is a multi-robot open field. Items are sprinkled out over the source area, and each type of item has its own sink. A homogeneous approach and two different heterogeneous approaches are studied. In [3], Balch suggests a method for measuring diversity in heterogeneous robot groups; he provides empirical foraging results with items sprinkled around in a large single source area, using no communication and two sinks.

In [8] Goldberg and Mataric give a definition of the multi-robot foraging task which allow multiple source and sink areas in open field environments. The homogeneous approach is compared to two heterogeneous approaches, *caste*

and *pack*. Caste foraging involves the use of dominance hierarchies as a method of interference reduction. If more than one robot has a resource to take to the sink, only the most dominant is allowed to go, then the next dominant, and so on. Pack foraging involves a division of labor between the foragers, in terms of a spatial division: one collecting robot remains near the sink area and brings the delivered resource into it, while the other search freely for resource. This results in a significant reduction of overall interference.

With respect to adding communication in foraging, not much work has been done. In [1] Arkin et al describe a simulation domain in which multiple robots collect “goals”, corresponding to items. Multiple robots can help carry a single item, resulting in increased transportation speed. Communication is shown to increase the overall performance of the system, in terms of time to task completion.

Rongier and Liegos [10] describe a petri net model of the multi-robot multi-source single sink foraging task, with “through the environment” communication. By representing the robot behaviors with stochastic petri nets, convergence of specific robot and environment configurations is demonstrated.

In [9] Parker does a proof of concept test of the *ALLIANCE* architecture, in which three robots perform a “waste cleanup mission”, a two-source, single-sink, multi-robot foraging task in an open field environment, using communication to reduce interference among the robots.

In [15] Werger & Mataric describe a multi-robot single-sink multiple source foraging task with a group of mobile robots equipped only with contact sensors. The robots use tapping to form chains, then others use the chain structure to forage in circular excursions. The length of the chain is adjusted in response to the amount of resource found by the foragers, and chain members can become foragers depending on resource abundance.

None of the work surveyed above has addressed bucket brigade in space constrained environments. This article addresses this strategy, and compares measured performance with some of the other methods.

3. MULTI ROBOT PERFORMANCE

The experiments described in this article are concerned with developing general methods for dealing with a large array of different foraging tasks, while having good scalability characteristics with respect to the number of robots. Coming up with a general yet scalable solution to the foraging problem is difficult, due to the many dimensions of the task. We focus on multi-robot scalability in the homogeneous transportation task.

In the case of foraging behavior in a bounded environment, there is a finite number of any particular type of robots that is optimal for performing the task. This optimum lies somewhere between the two extreme cases; no robots in the environment, and an environment completely full of robots, in which the performance of the system will be zero. Figure 4 illustrates how we might expect system performance to vary over the number of robots; the x-axis shows the number of robots, the y-axis is the system performance. The notion

Expected performance model for multi-robot systems in constrained environments




Figure 4: Expected performance as a function of number of robots

of “maximum performance” corresponds to the *critical mass* number of robots as reported in [5] and [1]. Exactly at which number of robots we will find this point, and how the graph looks on either side of the point, depends on the particular task and the particular implementation. We know that by varying the implementation of a given task, we can move the point of “maximum performance” and we can change the shapes of the curve on either side of it, but we cannot change the general shape of the graph. However, a controller capable of changing strategy as a function of the number of robots in the environment, might produce a bi-modal performance curve. This will be examined further in future work.

4. APPROACH

Our goal is to develop mechanisms that scale well with respect to the number of foraging robots, while still performing reasonably with only a few. We focus on the transportation task, a special case of the foraging task, in which the source location is known in advance. This corresponds to the basic foraging task shown in Figure 2, but without step 2 (search for item). To study this problem in spatially constrained environments, we implemented the following simple behavior:

```
While (terminate_condition) not satisfied do {
  1: Go to source area
  2: Find & pick up item
  3: Go to sink area
  4: Drop item
}
```

This behavior was implemented using the following four threads:

```
Thread 1:
  If (gripper_is_empty)
  { Go towards source }
Thread 2:
  If (gripper_not_empty)
  { Go towards sink }
Thread 3:
  If ((can_see_item) and (gripper_is_empty))
```




Figure 5: A maze-like environment we used.

```

{ Trigger pickup-behavior }
Thread 4:
  If ((can_see_sink) and (gripper_not_empty))
    { Trigger drop-behavior }

```

This implementation will be referred to as the *naive* algorithm, since it represents the baseline transpiration task we used, without any improvements.

4.1 Experimental environment

The experiments were performed using the Stage simulator [14], a real-time simulation of multiple Pioneer robots, interfacing to the Pioneers using Player [7]. When interfacing the pioneers through Player, the difference between running on the real Pioneers or in the Stage simulator is transparent to the programmer. This enables the Stage / Player combination is used as a tool for implementation of robot controllers, making it easy to test the controllers in simulation, before doing experiments on the real robots. In the experiments reported here, the simulated pioneer robots were equipped with SICK laser range finders, and top mounted color cameras capable of detecting unique color marker tags on the ceiling. A probabilistic topological map of the environment was provided to the robots at the start of each experiment. During the experiments there was 100% exploitation of the map, no exploration was employed.

The naive algorithm discussed above was implemented and tested in the environment shown in Figure 5. The sizes of the robots in the environment were such that two robots were unable to pass each other in the corridor. We ran ten twenty-minute long simulations for group sizes of one to eight robots, inclusively. Figure 6 shows the results. Not surprisingly, the performance of the system peaks at a very low number of robots. With group sizes larger than one, the corridors quickly get blocked with robots unsuccessfully trying to pass each other. The naive approach has no mechanism for recovering from this situation, so the system performance declines as the group size is increased. This shows an extreme case of interference.

To address the interference problem, a simple mechanism was added to the naive approach, consisting of one additional sensor and one additional behavior. The sensor triggers when it detects robots not carrying items in a 180 de-




Figure 6: The performance of the naive algorithm, 1-8 robots, 10 repetitions of 20 minute runs, in the maze like environment.

gree cone in front of a robot. When the sensor triggers, one of two responses is given. If the robot gripper is empty, nothing happens. If the robot has an item in its gripper, it triggers the execution of the added behavior, which drops the carried item and turns the robot 180 degrees.

This addition to the individual robot controller causes the entire system of robots to perform a bucket brigading type of global behavior. Figure 8 shows the performance of the bucket brigading algorithm compared to the naive algorithm in the above described environment.

Comparison between the naive algorithm and the bucket brigade algorithm should be done with caution. Neither of the algorithms has been “tweaked” to increase performance, and performance of each is highly sensitive to the time involved in picking up items, the robots’ maneuvering capabilities, and so on. Rather than absolutes, the results should be viewed for trends.

The bucket brigading behavior the robots exhibit can be considered emergent in that there is no explicit representation of a bucket brigade in the control algorithm. Instead, the simple addition of a sensor and a rule to the naive approach results in a bucket brigade being formed in the environment. No explicit information is given about where the robots should position themselves in the brigade. The algorithm is very robust; robots can be added to or removed from the brigade at any time during experiments without disturbing the global behavior of the system.

We chose three different environmental conditions to test the bucket brigading algorithm. These correspond to varying degrees of spatial constraint in the environment. Figure 7 illustrates the difference between them. In **a**) the corridor is so narrow that it is impossible for two robots to pass each other. In **b**) two robots can pass each other but with some difficulty, and in **c**) robots can easily pass each other.

5. RESULTS

Two sets of experiments were performed, both aimed at producing conditions as those shown in Figure 7. In the first set of experiments, the environment was varied to correspond




Figure 7: Three types of environments with varying spatial constraints

to the three cases shown in Figure 7. In the second set, the environment was kept constant throughout all experiments, but the sizes of the robots were changed, producing the effect of differing spatial constraints of the environment.

5.1 First set of experiments

The first set of experiments consisted of three experiments, 1a, 1b and 1c. In two of the experiments, the maze-like environment shown in Figure 5 was used. In the first experiments, the walls of the maze were made so thick that two robots could not pass each other. In the second experiment, two robots could just barely pass each other. The third experiment was done in an open field-like environment (the walls were completely removed from the maze). These three experiments corresponds loosely to the three cases shown in Figure 7.

Each experiment was run 10 times, for 20 minutes each, with group sizes of one to eight robots. In each experiment, the naive algorithm was compared to the bucket brigading algorithm. The results from experiments are shown in Figures 8 to 10. The x-axis shows the number of robots, the y-axis the number of items collected during the experiment.

In Figure 8 we see that both algorithms perform equally well when run with a single robot, collecting on average about 14 items. This is as expected, since no interference problems arise with only one robot in the environment. In Figure 8 we also see that for more than 2 robots in the system, the naive algorithm performs badly, while the bucket brigade algorithm performs much better. At eight robots, the bucket brigading algorithm collects 133 items in 20 minutes in average, which is about 4 times as many as a single robot for the




Figure 8: Experiment 1a (2a), Narrow corridors

same amount of time. The naive algorithm collects about 3 items on average in 20 minutes.

Figure 9 shows the result from a run with corridors just wide enough for two robots to get past each other. The naive algorithm performs well up to 3-4 robots, but with more than four robots the amount of traffic in the corridors becomes too high, and performance drops. The bucket brigade algorithm performs less well than the naive algorithm for fewer than 4 robots, mostly because of unnecessary dropping and picking up objects.

Figure 10 shows results from the run with the open field environment. The naive algorithm performs better, however both algorithms seems to drop off after more than six robots. This effect is probably due to a combination of a slowdown of the simulator¹ and of robot interference near the source and sink areas. Additionally, in the case of the open field environment, the bucket brigade algorithm tends to litter somewhat; a few items remain on the ground over a 20 minute run. Littering occurs when one robot drops an item because it sees another robot, and the other robot does not see the item that is dropped. In the case of the constrained space environments, the litter is usually found by another robot at a later time, but in the open field environment, there is less chance of another robot taking a path near the same spot.

5.2 Second set of experiments

The second set of experiments, 2a-2c, was designed to provide for greater comparability between experimental data. In this set, instead of varying the environment, we manipulated the sizes of the robots in the environment, to produce the effect illustrated in Figure 7. The sizes and the sensor ranges of the robots were scaled down, but the speed of the robots and the structure of the environment were kept constant. Experiment 2a is identical to experiment 1a (the same data are used, shown in Figure 8), while experiments 2b and 2c correspond to robot sizes one half and one fifth of that in experiment 2a. Results from experiments 2b and 2c are shown in Figure 11 and 12.

¹An unfortunate artifact in which simulator becomes slow when simulating the SICK laser in open environments.




Figure 9: Experiment 1b, Less narrow corridors




Figure 10: Experiment 1c, Open field




Figure 11: Experiment 2b, Small robots




Figure 12: Experiment 2c, Tiny robots




Figure 13: Naive algorithm performance, 2a-2c

The data point to two conclusions. First, performance of the bucket brigade algorithm does not seem to vary as a function of the size of the robots. Second, performance of the naive algorithms varies a great deal as a function of the robot size.

When comparing tendencies for the bucket brigading algorithm in the first and second set of experiments, it is seen that the performance varies much less in the second set. This might suggest that the environment plays an important role in guiding bucket brigading robots to items dropped by other robots. Experiment 1c shows a loss of performance with more than 6 robots, while experiment 2c does not show this tendency. This indicates that the environment type is a larger factor in using the implemented bucket brigading algorithm than interference amongst the robots.

With respect to the naive algorithm, we would expect that in an environment with no robot interference, the performance would increase linearly with the number of robots. Figure 13 shows the performance of the naive algorithm from experiments 2a-2c, along with a line of “expected performance with no interference”. The series of experiments 2a to 2c seems to confirm our expectations, in that the performance curve of the naive algorithm seems to be converging toward the line of “expected performance with no interference”.

6. CONCLUSION

In this paper, we have first presented a set of characteristics that can be used toward a taxonomy for distinguishing different implementations of the foraging task. Second, we have described a simple mechanism that allows robot systems to do bucket brigading behavior, and we have done two sets of experiments with the algorithm, varying the spatial constraints of the environment. The experimental results lead us to conclude that the bucket brigading algorithm has a tendency to perform better than a naive approach in spatially constrained environments, and to perform worse than the naive approach in open field-like environments. That is, as the spatial constraints of the environment increase, the performance of the bucket brigade algorithm increases, while performance of the naive algorithm decreases.

Our future work on the bucket brigading algorithm for foraging involves eliminating the need for the additional sensor. It might be possible to use interference among the robots to detect when it is a good time to drop the carried object.

Our future work in general is focused on finding mechanisms that improve performance of multi-robot systems over a large number of different foraging tasks. The beginnings of a taxonomy we presented in this paper provides us with several axes to pursue further. Particular interest was given in this paper to examining the role of spatial constraints of the environment. Future work will be concerned with heterogeneous robots using communication in more complex environments with multiple sources and sinks, which would introduce problems concerning load balancing, exploration vs. exploitation, and communication management. After having developed methods for handling many different environments, we will start investigating methods for dynamically changing robot strategies.

7. ACKNOWLEDGMENTS

This work is supported in part by DARPA grant DABT63-99-1-0015 under the Mobile Autonomous Robot Software (MARS) program, ONR grant N00014-00-1-0638 and NSF grants ANI-9979457 and ANI-0082498 under the Special Projects in Networking Program.

8. REFERENCES

- [1] R. C. Arkin, T. Balch, and E. Nitz. Communication of behavioral state in multi-agent retrieval tasks. In *Proc. IEEE Conf on Robotics and Automation*, Atlanta, GA, 1993. IEEE Press.
- [2] T. Balch. The impact of diversity on performance in multirobot foraging. In *Autonomous Agents 99*, Seattle, WA, 1999.
- [3] T. Balch. Hierarchic social entropy: An information theoretic measure of robot group diversity. In *Autonomous Robots, 8:3*, July, 2000.
- [4] Y. Cao, A. Fukunaga, A. Kahng, and F. Meng. Cooperative mobile robotics: Antecedents and directions. In *'IEEE/TSJ International Conference on Intelligent Robots and Systems'*, Yokohama, Japan., 1995.
- [5] M. S. Fontán and M. J. Matarić. A study of territoriality: The role of critical mass in adaptive task division. In P. Maes, M. Matarić, J.-A. Meyer, J. Pollack, and S. Wilson, editors, *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pages 553–561. MIT Press, 1996.
- [6] M. S. Fontán and M. J. Matarić. Territorial multi-robot task division. *IEEE&A*, 14(5), 1998.
- [7] B. Gerkey, K. Stoy, and R. T. Vaughan. Player robot server. Technical report, Institute for Robotics and Intelligent Systems Technical Report IRIS-00-391, University of Southern California, URL: <http://robotics.usc.edu/player/>, 2000.
- [8] D. Goldberg and M. J. Matarić. Robust behavior-based control for distributed multi-robot collection tasks. *USC Institute for Robotics and Intelligent Systems Technical Report IRIS-00-387*, 2000.
- [9] L. Parker. Alliance: An architecture for fault-tolerant multi-robot cooperation. In *IEEE Transactions on Robotics and Automation*, 14(2):220–240, 1998.
- [10] P. Rongier and A. Liegeois. Analysis and prediction of the behavior of one class of multiple foraging robots with the help of stochastic petri nets. In *Proc. IEEE Int. Conf. on Systems, Man, and Cybernetics (SMC'99)*, 12-15 October 1999, Tokyo, Japan, volume 5, pages 143–148, 1999.
- [11] K. Sugawara, M. Sano, and I. Yoshihara. Cooperative acceleration of task performance: Analysis of foraging behavior of multi-robot system. *Proc. IPSJ Int. Symp. on Information Systems and Technologies for Network Society*, pages 317–320, 1997.
- [12] R. Vaughan, K. Støy, G. Sukhatme, and M. Matarić. Blazing a trail: Insect-inspired resource transportation by a robot team. In *Proceedings, 5th International Symposium on Distributed Autonomous Robotic Systems (DARS)*, Knoxville, TN, Oct 4-6, 2000.
- [13] R. Vaughan, K. Støy, G. Sukhatme, and M. Matarić. make my day: Robot conflict resolution by aggressive competition. In *Proc. 6th Int. Conf. Simulation of Adaptive Behaviour*, 2000.
- [14] R. T. Vaughan. Stage: a multiple robot simulator. Technical report, Institute for Robotics and Intelligent Systems Technical Report IRIS-00-393, University of Southern California, 2000.
- [15] B. B. Werger and M. J. Matarić. Exploiting embodiment in multi-robot teams. Technical Report IRIS-99–378, University of Southern California, Institute for Robotics and Intelligent Systems, 1999.