

Adapting to non-uniform resource distributions in robotic swarm foraging through work-site relocation

Adam Lein and Richard T. Vaughan
Autonomy Lab, Simon Fraser University
Burnaby, British Columbia, V3K 3W5 Canada.
{alein,vaughan}@sfu.ca

Abstract—We describe a simple controller for swarms of foraging robots that reduces mutual spatial interference and adapts to non-uniform resource distributions. We discuss several sources of such non-uniformity, and show that some non-uniform distributions are not well-handled by previously described foraging schemes. Our new method adapts each robot’s work site size and location during run time, to reflect the distribution encountered. The controller is very scalable, using only local information and no explicit communication. Simulation studies demonstrate the method’s effectiveness.

I. INTRODUCTION

Central-place foraging is a well-studied task in robotics and ethology [20]. The task requires agents to locate and collect spatially distributed items and deliver them to a “home” location. The delivery rate can be increased by adding more foragers, but typically with diminishing returns due to mutual spatial interference.

This interference effect creates a problem for foraging robot controllers. To maximize performance, we want robots to work where there are items to collect (“pucks” hereafter). Yet we need to keep the robots spread out to minimize interference. When pucks are clustered together, these two goals are mutually exclusive. Figure 1 shows an example of this difficult domain that is the focus of this paper.

In animal foraging literature, resources are frequently described as occurring in “patches” or “food-source locations”, or otherwise referred to as clusters implicitly in discussions of isolated food sources distant from a nest site [15], [11], [3]. Similar language and environmental setups are used in ant-inspired artificial agent research, often in the context of pheromone trail-following [17], [4].

Research in robot foraging has addressed scenarios in which resources are uniformly distributed and where resources are tightly clustered. The dominant approach for the former is “bucket brigading”, which separates robots in space to minimize mutual interference [19], [14], [6]. The dominant approach for the latter is ant-inspired trail-following, which directs robots to resource patches. Trail-following robots are separated in time rather than space, and *requires* that robots are able to navigate close together without significant mutual interference.

In [13], we extended the bucket brigading method introduced in [19]. Our adaptive bucket-brigade foraging algorithm (described in Section V) allowed robots to vary the size




Fig. 1: Robots in simulation foraging around a cluster of pucks. Robots are indicated by boxes with “arms”, and pucks by solid black circles. Blue/dark robots are searching, green/light robots are homing, and red/medium robots are returning to their work areas. The larger circle in the northeast corner is the home zone.

of their work-site dynamically in response to interference with other robots. We showed that this could increase the performance and scalability of the foraging team.

Here we further extend adaptive bucket brigading to allow robots to relocate their work zones to more productive areas. This allows the distribution of robots to reflect the current distribution of pucks, while still separating them in space to maintain resistance to interference.

This paper narrows the gap between previous work in bucket brigading, characterized by its assumption of high risk of spatial interference, and uniformly distributed resources, and ant-like foraging, with its assumptions of relatively very low spatial interference and tightly clustered resources. We discuss some ways that clusters may form spontaneously in a foraging system, further motivating cluster-friendly methods.

We describe and examine our novel adaptive forager in a series of simulation experiments.

II. RELATED WORK

Foraging strategies that seek to reduce spatial interference have been studied extensively. Liu *et al.*, in [14], produced an adaptive controller to dynamically decide between “resting” and “foraging” operations, in order to maximize energy income. Here interference was controlled by adapting the size of the foraging population.

Fontán and Matarić [6] investigated foraging with small teams (2, 3, and 4 robots) of real robots which forage in separate territories to reduce spatial interference. Territories were assigned *a priori* and covered all available space without overlap. In contrast, in our work, work areas (the analogue of territories) are assigned dynamically and adaptively.

Shell and Matarić [19] pioneered the bucket-brigade approach to foraging in large-scale multi-robot systems. They discussed interference in multi-robot systems. Instead of allowing all foragers to explore the entire environment, they restricted each forager to straying no further from its starting location than a universal, preset value: the search radius. That study found a relationship between the number of workers and the performance of the team as a function of search radius. In the paper, we use a similar algorithm.

Bucket-brigade foraging had previously been explored, and its efficacy demonstrated in real robots [9], [16]. Bucket-brigade foraging in ants has also been studied (and explicitly named) in [11]. Gordon *et al.* studied spatial division of labor in red harvester ants [10].

Notably, Schmickl and Crailsheim, in [18], simulated robots acquiring food from a localized source *and* investigated the possibility of *positive* effects of spatial interference. Their robots communicated using trophallaxis (transferring food from one robot to the next directly) and avoided non-trophallactic collisions with one another. The use of trophallactic contact induced a gradient of food concentrations, and robots could measure the local gradient and use the information to navigate uphill. However, their research did not explicitly look at the effect that using a localized source had on the performance of their algorithm. They did not address alternate distributions of the “dirt”, such as the uniform distribution.

Other authors are closing the natural/artificial gap between ants and robots. Kube and Bonabeau, in [12], modeled a group of ants cooperatively carrying a single food item with real robots in a cooperative box-pushing task. They produced a coordinate movement effect without direct communication between robots.

Bucket-brigade foraging has been studied in insect societies as a form of task partitioning. Anderson *et al.*, in [1], described a form of bucket-brigading seen in several species of ant and termite in which insects pass resources directly from one to the next until it reaches the nest. They gave situations in which bucket-brigading would be especially effective, such as situations in which material must be passed along a narrow passageway, or when many insects foraging around the same source create a bottleneck.

Trail-following behavior in robotics research is well-studied. Vaughan *et al.* in [21] implemented trail-following

in robots in a transportation task (similar to foraging) using low-bandwidth wireless communication. Real-world implementation of ant-like trail following shows that mutual spatial interference is a severe limitation. Later work by the same authors addressed this problem directly by quickly resolving conflicts over navigation space [23].

III. SIMULATED ROBOTIC FORAGING

Our experiments were performed in a dedicated multi-robot foraging simulator. The environment is a square region with a home location in the northeast corner. Space is approximately continuous, and there are no obstacles in the environment apart from the robots themselves; robots avoid collisions with the environment’s boundaries and with one another. Pucks are distributed throughout the environment, either uniformly or in a cluster depending on the experiment; pucks are modeled as points and do not “interfere” with one another. Robots are initially placed at random intervals throughout the environment.

Robots are equipped with 12 short-range proximity sensors capable of detecting walls and other robots up to 1m away; the sensors can differentiate between walls and robots. Each robot has a gripper which can lift, carry and drop a single puck. Another sensor detects when a puck is under the gripper and therefore grippable. This is the only way the robots can sense pucks. The robots measure $15\text{cm} \times 15\text{cm}$, approximately $1/28,000$ the area of the world.

Except for the ability to place clusters of pucks (see below), this is the same experimental setup as used in [13].

IV. CLUSTERING

Much of the previous work with foraging in robotic swarms assumes that the distribution of pucks is initially uniform. However, in real-world foraging scenarios, it is possible that the resources to be collected are distributed non-uniformly, especially in patches or clusters. In this work, we examine scenarios where there is initially a single cluster of pucks. Clusters are of interest for two reasons, first because some resources naturally or inherently are produced or deposited in clusters (apples on a tree; a bucket of tennis balls; a pile of rocks), and secondly because clusters may emerge as a side-effect of the actions of the robots and the peculiarities of their controllers.

The controller developed in [19] allowed robots to keep track of the distance and bearing to their starting location, through odometry. The inevitable error and unbounded drift in odometry caused the robots’ work areas to drift over time. The authors noted that this had the added effect of ensuring coverage of otherwise neglected areas of the environment.

In our original experiments with global, non-drifting localization, we noticed that some areas of the environment were indeed neglected (see Figure 2a), in that no robot’s work area covered them. Clusters of pucks spontaneously formed in these neglected areas. Even when total coverage is achieved, if some areas are more frequently visited for drop-off than pick-up then a clustering effect will exist.




Fig. 2: Clusters forming spontaneously (a) in neglected regions and (b) near the boundaries of foragers’ work areas. The large circles indicate the boundaries of robots’ work areas. The areas outside the clusters are blurred to highlight the clusters.

As an example of this effect, we observed that clusters of pucks tended to spontaneously form near the edges of robots’ work zones, visible in the simulator screenshot depicted in Figure 2(b).

The existence of non-uniform resource distributions, typically clusters caused either by the asymmetry inherent in the initial supply of resources, or by asymmetrical pick-up and drop-off frequency, motivated us to study adaptive approaches to foraging that work well under these conditions.

In our experiments, we looked at two distributions of pucks:

- the “hat” distribution, in which pucks are uniformly distributed, but only appear within a fixed circular area (the “cluster”) of radius $\frac{1}{c} \times 25\text{m}$, and
- the uniform distribution (effectively $c = 0$).

Here c is the “clustering parameter”, which indicates how tightly clustered the pucks are.

V. FORAGING IN A BUCKET-BRIGADE

In the bucket brigade foraging controller of [19], each robot stays within the same fixed radius of its work-space location. This restriction means that robots approximately maintain their initial uniform distribution, keeping the robots spread out and limiting interference. Our adaptive method allows the robot’s distribution to change over time in response to the environment. We describe the original method and our extensions in this section.

A. Basic bucket-brigading algorithm

Robots can be in one of three states: searching, homing, or returning. The initial state is searching.

- A robot in the searching state searches its “work area”, the circle defined by the robot’s initial location and a fixed work-radius. The search method is implementation-specific; the details of how a robot goes about locating a puck are not relevant to the application of the bucket-brigade strategy. In our implementation, robots followed a random walk except when avoiding obstacles and steering to stay within the work area.

On detecting a puck, the robot picks it up and transitions from the searching state to the homing state.

- A robot in the homing state drives in the direction of the global home location (supplied to the robot) while avoiding collisions with other robots.

On reaching the home zone, or on leaving its work area, the robot drops its puck and transitions to the returning state.

- A robot in the returning state, drives towards the center of its work zone. On arriving within *half* their search radius from this location, the robot returns to the searching state.

In any state S , if a robot must avoid an obstacle it will transition to an avoiding state until obstacle avoidance is no longer necessary, at which point it will return to state S , as per a subsumption architecture [2].

The result of this strategy is that very few robots actually deliver pucks to the home zone. Assuming robots are uniformly distributed in space, and assuming a sufficiently large population, one or more robots will have a work area that overlaps the home zone. These robots will transport pucks just outside the home zone into the home zone. Other robots will have work areas that overlap these robots’ work areas, and this pattern of overlap will continue so that most of the environment is covered by a work area that is “connected” to the home zone, and pucks will travel toward the home zone in a manner suggested by the algorithm’s name—a bucket-brigade.

B. Adding adaptive work-zone size

Our first extension is to allow each robot to adapt its work-zone radius with experience. While the robot is avoiding collisions with other workers, it decreases the radius at a preset rate. Otherwise, the radius increases at a (possibly different) preset rate. Allowing different robots to use different work area radii gives the team more flexibility in adapting to the distribution of workers in the environment. We have previously shown that this simple extension improves performance and allows larger population sizes to work efficiently in uniform puck distributions [13].

C. Effects of clustering on the bucket-brigade controller

Given the controller described above, we found that if the initial placement of pucks is clustered instead of uniformly distributed, then performance decreases. Using the simulation environment shown in Figure 1 we ran 30 trials for a range of cluster parameters c , and for two robot populations. Figure 3 shows the results. The more tightly clustered the pucks (larger c), the fewer pucks are delivered over the length of the experiment. The success of the non-adaptive bucket brigade method depends on exploiting the uniform distribution of pucks.

D. Where to start searching

Recall that the bucket-brigade foraging algorithms require each robot to maintain an approximation of its start location. They begin their puck search near this point each time.




Fig. 3: Performance of the standard, non-adaptive bucket-brigade algorithm degrades significantly as the degree of clustering increases. As the “clustering parameter” increases, the cluster gets smaller but has the same number of pucks in it. Mean performance over thirty trials are plotted with 95% confidence interval.




Fig. 4: If D is the distance between the old starting location and the newly found puck, then ρD is the new distance; ρ (for $0 \leq \rho \leq 1$) is the relocation parameter.

When puck density is uniform, one location is as good as another for puck searching, and the start location serves only to spread robots out. In non-uniform puck distributions, however, placing the center of the work zone in a puck-rich neighborhood is likely to reduce search time and improve performance. Our second adaptive modification attempts to improve the position of the work zone center over time.

As in Shell and Mataric’s original method [19], our robots’ estimates of their work areas drift on a slow random walk. We improve on this by adding non-random, deliberate work-zone *relocation* as follows.

Upon finding a puck, a robot will relocate its estimate of its starting location a certain portion of the way along the line from its current position to the location of the puck. This proportion is an adjustable *relocation parameter* ρ ; a value of $\rho = 0$ (no relocation) corresponds to non-relocating bucket-brigade foraging, in which the robot’s estimate of its starting location remains fixed (upon finding pucks; it is still subject to drift). A value of $\rho = 1$ (complete relocation) indicates that robots will return to the last place they found a puck before they begin searching again. All robots use the same value of ρ (though there is no a priori reason requiring this). This process is illustrated in Figure 4 and described algorithmically in Algorithm 1.

This tends to cause robot work zones to move towards areas of high puck density, reducing search time and poten-

Algorithm 1 Procedure for robots in the searching state.

Require: $0 \leq \rho \leq 1$ is the *relocation parameter*

Require: $r_Z \geq 0$ is the *range to zone center* from odometry

if I found a puck **then**

 Pick up the puck.

$r_Z \leftarrow (1 - \rho)r_Z$

 Switch to the homing state.

end if

tially improving performance.

VI. EXPERIMENTS

A. Simulated environment

We performed a series of simulation experiments to examine the new method. Our hypothesis is that relocation can improve performance by adapting to the current puck distribution.

The simulator was a modified version of the one used in [13]. The foraging environment was a $25\text{m} \times 25\text{m}$ enclosure with a quarter-circular home area, of radius 1m, in the northeast corner. Robots were initially positioned at random locations (different in each trial).

Pucks are initially placed at random locations throughout the environment (different in each trial), either in a uniform distribution, or according to the clustering model described in Section IV. In either case, on average there were two pucks per square meter (for a total of 1,250 pucks available to be foraged). Upon delivery to the home zone, a puck is replaced at random according to the same initial distribution.

B. Parameter space

Experiments were run with all possible combinations of the following parameter values:

- Clustering parameter: we tested several small values ($c = 1, \dots, 5$), and a few larger values ($c = 10, c = 20$).
- Robot population: could be either a “small swarm” of 36 (0.25 robots per square meter) or a “large swarm” of 100 (0.69 robots per square meter).
- Relocation parameter: could be 0.00, 0.25, 0.50, 0.75, or 1.00.

C. Performance metric

We make two claims; first, that clustering of pucks hurts the performance of existing algorithms, such as bucket-brigading, which ignore the distribution of the pucks. This is supported by the data discussed in Section V-C and displayed in Figure 3. The second claim is that we can adapt the foraging algorithm to work well in the presence of non-uniformity by using the relocation system described in Section V-D and illustrated in Figure 4. To support the second claim, we compare the performance of foragers in clustered-resource situations with and without relocation.

In each experiment, robots foraged for one hour of simulated time. At the end of the hour, the total number of pucks delivered to the home zone was counted.

#robots	Clustering	ρ	Performance		p
			μ	σ	
36	0	0	221.000	33.837	< 0.0001
36	0	1	130.300	19.701	
36	5	0	120.633	15.940	< 0.0001
36	5	1	179.800	14.372	
36	10	0	76.900	11.973	< 0.0001
36	10	1	146.133	12.260	
36	15	0	56.100	9.386	< 0.0001
36	15	1	124.367	17.604	
100	0	0	313.833	36.545	0.0022
100	0	1	347.433	53.187	
100	5	0	335.667	36.456	0.0002
100	5	1	367.500	33.470	
100	10	0	208.033	17.629	< 0.0001
100	10	1	328.067	19.311	
100	15	0	150.033	15.751	< 0.0001
100	15	1	282.833	24.552	

TABLE I: Analysis of performance for extreme values of the relocation parameter ρ . p -values for the significance of the difference between the $\rho = 0$ and $\rho = 1$ cases' performances are shown (mean over thirty trials).

Performance is reported as the average number of pucks collected over thirty trials for each possible choice of parameters as given in Section VI-B. A 95% confidence interval was determined, assuming that real performance is normally distributed.

VII. RESULTS AND DISCUSSION

We found that relocation can have either a beneficial or an adverse effect, depending on the setting. The results support our hypothesis that relocation would allow robots to adapt to a tightly clustered distribution of pucks: swarms employing complete relocation outperformed those that did not relocate their starting locations, nearly doubling performance in the small-swarm (36 robots) scenario. These results are statistically significant; see Table I.

The results are illustrated graphically in Figure 5.

We found that the performance did increase when relocation was used, and we suggest that the observed improvements were due to the robots' search areas adapting to the puck distribution. See Figure 6 for an illustration.

VIII. FUTURE WORK

In focusing on adaptive work site relocation, we discovered a wealth of complexity arising out of a non-uniform resource distribution. In [13] we provided a simple, one-dimensional mathematical model of such distributions that form as a result of the action of robots using the bucket-brigade algorithm for transporting pucks. We would like to develop a more general model that includes the variety of sources of clustering. We would also like such a model to incorporate the energy requirements of the workers, as discussed in [14].

Furthermore, we would expect that relocation of work sites would arrange the robots approximately in an *ideal free distribution* as described by Fretwell and Lucas in [7]. That is, in an environment with multiple clusters, the robots will distribute their starting locations so that the number of




Fig. 5: Performance of relocating compared to non-relocating foragers, averaged over thirty trials, with 95% confidence interval, over a range of relocation parameters. In tightly clustered puck distributions, the graph shows an upward trend in performance as the relocation parameter is increased.

robots working at each cluster is proportional to the number of pucks in the cluster. Further investigation is needed to confirm this hypothesis.

The problem of spatial interference motivated the bucket-brigade foraging algorithm as a means to *spatially separate* the workers[19]; temporal separation has been studied in both ecology[5] and robots[8]. We would pursue a descriptive (or even better-generative) model that can describe or produce spatial and temporal foraging in a unified framework.

Wawerla and Vaughan modeled the time-value of labor in [22]. Fundamentally, all robotics enterprises must be judged in the context of a real-world economy; future work in foraging swarms should take this dimension into account. This would provide greater weight to pucks foraged early on, which highlights the value of early recruitment of workers into active foraging.

Also, since we have shown that relocation can be beneficial or detrimental to performance in different environments, we would like, as we did in [13], to develop an adaptive mechanism for selection of the relocation parameter.




Fig. 6: Distribution of pucks (dark circles) and robots' starting locations (squares) initially and after 20 and 40 minutes. As pucks are spread out via bucket-brigading, the foragers automatically relocate to the "trail" of pucks from the cluster to the home zone. For the clarity of this illustration, robots were initially placed on a grid.

IX. ACKNOWLEDGMENTS

Thanks to Jens Wawerla and the rest of the Autonomy Lab for providing valuable feedback in the development of this paper. This work was supported by NSERC and DRDC in Canada.

REFERENCES

- [1] C. Anderson, J. J. Broomma, and III J. J. Bartholdi. Task partitioning in insect societies: bucket brigades. *Insected Sociaux*, 49:171–180, 2002.
- [2] Rodney A. Brooks. Intelligence without representation. *Artificial Intelligence*, 47:139–159, 1991.

- [3] S. Holmes D. M. Gordon, R. J. Schafer. Forager activation and food availability in harvester ants. *Animal Behaviour*, 71(4):815–822, April 2006.
- [4] J. L. Deneubourg, S. Goss, N. Franks, A. Sendova-Franks, C. Detrain, and L. Chrétien. The dynamics of collective sorting: Robot-like ants and ant-like robots. In *Proceedings of the first international conference on simulation of adaptive behavior: From animals to animats*, pages 356–363, Cambridge, MA, USA, 1990. MIT Press.
- [5] Joan H. Fellers. Interference and exploitation in a guild of woodland ants. *Ecology*, 68(5):1466–1478, October 1987.
- [6] Miguel Schneider Fontán and Maja J Matarić. A study of territoriality: The role of critical mass in adaptive task division. In *In From Animals to Animats 4: Proceedings of the Fourth International Conference of Simulation of Adaptive Behavior*, pages 553–561. MIT Press, 1996.
- [7] Stephen Dewitt Fretwell and Henry L. Lucas. On territorial behavior and other factors influencing habitat distribution in birds. *Acta Biotheoretica*, 19(1):16–36, 1969.
- [8] D. Goldberg. *Exploring the Dynamics of Agent-Environment Interaction*. PhD thesis, University of Southern California, May 2001.
- [9] D. Goldberg and M.J. Matarić. Maximizing Reward in a Non-Stationary Mobile Robot Environment. *Autonomous Agents and Multi-Agent Systems*, 6(3):287–316, 2003.
- [10] D. M. Gordon, J. Chu, A. Lillie, M. Tissot, and N. Pinter. Variation in the transition from inside to outside work in the red harvester ant, *pogonomyrmex barbatus*. *Insectes Sociaux*, 52:212–217, 2005.
- [11] Stephen P. Hubbell, Leslie K. Johnson, Eileen Stanislav, Berry Wilson, and Harry Fowler. Foraging by bucket-brigade in leaf-cutter ants. *Biotropica*, 12(3):210–213, 1980.
- [12] C. Ronald Kube and Eric Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30:85–101, 2000.
- [13] Adam Lein and Richard Vaughan. Adaptive multi-robot bucket-brigade foraging. In *Proceedings of the Eleventh International Conference on the Simulation and Synthesis of Living Systems, Artificial Life*, pages 337–342. MIT Press, August 2008.
- [14] Wenguo Liu, Alan F. T. Winfield, Jin Sa, Jie Chen, and Lihua Dou. Towards energy optimization: Emergent task allocation in a swarm of foraging robots. *Adaptive Behavior - Animals, Animats, Software Agents, Robots, Adaptive Systems*, 15(3):289–305, 2007.
- [15] Peter Nonacs and Joanne L. Soriano. Patch sampling behaviour and future foraging expectations in argentine ants, *linepithema humile*. *Animal Behavior*, 55(3):519–527, March 1998.
- [16] E.H. Østergaard, G.S. Sukhatme, and M.J. Matarić. Emergent bucket brigading: a simple mechanisms for improving performance in multi-robot constrained-space foraging tasks. In *Proceedings of the fifth international conference on Autonomous agents*, pages 29–30. ACM Press New York, NY, USA, 2001.
- [17] Liviu Panait and Sean Luke. A pheromone-based utility model for collaborative foraging. In *Proceedings of the 2004 Conference on Autonomous Agents and Multiagent Systems*, 2004.
- [18] T. Schmickl and K. Crailsheim. Trophallaxis among swarm-robots: A biologically inspired strategy for swarm robotics. In *The First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechanics, 2006. BioRob 2006.*, pages 377–382, 2006.
- [19] Dylan A. Shell and Maja J. Matarić. On foraging strategies for large-scale multi-robot systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2717–2723, Beijing, China., Oct 2006.
- [20] David W. Stephens, Joel S. Brown, and Ronald C. Ydenberg. *Foraging*. University of Chicago Press, Chicago, 2007.
- [21] Richard T. Vaughan, Kasper Stóy, Gaurav S. Sukhtame, and Maja Matarić. Whistling in the dark: Cooperative trail following in uncertain localization space. In *Agents*, pages 187–194, Barcelona, Spain, 2000.
- [22] Jens Wawerla and Richard T. Vaughan. Optimal robot recharging strategies for time discounted labour. In *Proceedings of the Eleventh International Conference on Simulation and Synthesis of Living Systems, Artificial Life*, August 2008.
- [23] Mauricio Zuluaga and Richard Vaughan. Reducing spatial interference in robot teams by local-investment aggression. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2798–2805, August 2005.