

Adaptive multi-robot bucket brigade foraging

Adam Lein Richard T. Vaughan

Autonomy Lab, Simon Fraser University
Burnaby, British Columbia, Canada
{alein,vaughan}@sfu.ca

Abstract

Bucket brigade foraging improves upon homogeneous foraging by reducing spatial interference between robots, which occurs when robots are forced to work in the same space, and must spend time avoiding one another instead of carrying out useful work. Bucket brigade foraging algorithms restrict the motion of each robot to at most some fixed distance from its starting location. We reproduce the performance of known bucket brigade foragers, and then present a new controller in which robots adapt the size of their foraging area in response to interference with other robots, improving overall performance. This approach also has the potential to cope with nonuniform resource distributions.

Introduction

The *foraging problem* is a task in which one or more agents must find and collect target objects and deliver them to a “home area”. The simplest adaptation of this problem to multi-robot systems is for many robots to independently solve the foraging task, a solution known as *homogeneous foraging*.

Shell and Matarić (2006) explore foraging strategies in large-scale multi-robot systems. The key variable in their experiments is spatial interference, which refers to destructive interactions between robots forced to perform work in the same space. Multi-robot systems are advantageous only when what might be called the marginal performance – the benefit to performance of adding a single robot to the system – is positive. However, as those authors demonstrated, there comes a time when this is no longer so, when the loss of performance due to interference between robots outweighs the gain in work done by having more workers.

Goldberg and Matarić (2003) and Østergaard et al. (2001) describe *bucket brigading*, in which each robot is restricted to a finite search area, and instead rely on their coworkers to both deliver pucks into their search area, and remove pucks out of it. By restricting each robot to a finite area whose size is determined a priori, interference is ameliorated. Shell and Matarić (2006) empirically investigate the performance of large groups of robots as a function of varying search area sizes; homogeneous foraging corresponds to search areas

of infinite radius. Østergaard et al. (2001) describe the expected performance of multi-robot, space-constrained systems as a curve to which the $R = 40m$ curve in Figure 1 roughly corresponds; the curve has a local maximum, after which the marginal performance is negative.

In this paper, we investigate an approach to bucket brigading that does away with *a priori* search radii by allowing robots to adapt their search radii in response to interference with other robots, improving overall performance.

Related work

Foraging is a common analogy for a wide variety of robot tasks, including exploration and mapping, search, and actual objects retrieval.

Beckers et al. (1994) noted the use of *stigmergy* in insect swarms. Stigmergy is a process by which insects (in general, agents) communicate implicitly and indirectly by modifying the environment – their coworkers’ future behavior is affected by these changes. The authors built robots which utilized stigmergy to collect objects and gather them into a pile – essentially analogous to the foraging task discussed in this paper. They found that increasing the number of robots decreases the mean time required to complete the task, *but only up to a certain point* (three robots), after which the mean time increased, due to interference between robots.

Holland and Melhuish (1999) examined stigmergy and self-organization in physical robots: using very simple behavioral rules, the robots were able to cluster and sort Frisbees, despite possessing no memory or capacity for spatial orientation. The authors argued that their results hinged on the robots’ behavior taking advantage of real-world physics.

Wawerla and Vaughan (2007) applied the rate-maximizing foraging model to a single robot performing the task of foraging over a long period of time. The robot had a finite energy supply, and was required to travel to a charging station to recharge its batteries. While recharging, and while traveling between the work site and the charging site, the robot is not doing work. The authors presented a scalable, online, heuristic algorithm for the robot to recharge efficiently, maximizing the proportion of its time it spends

Fixed range, density=0.781

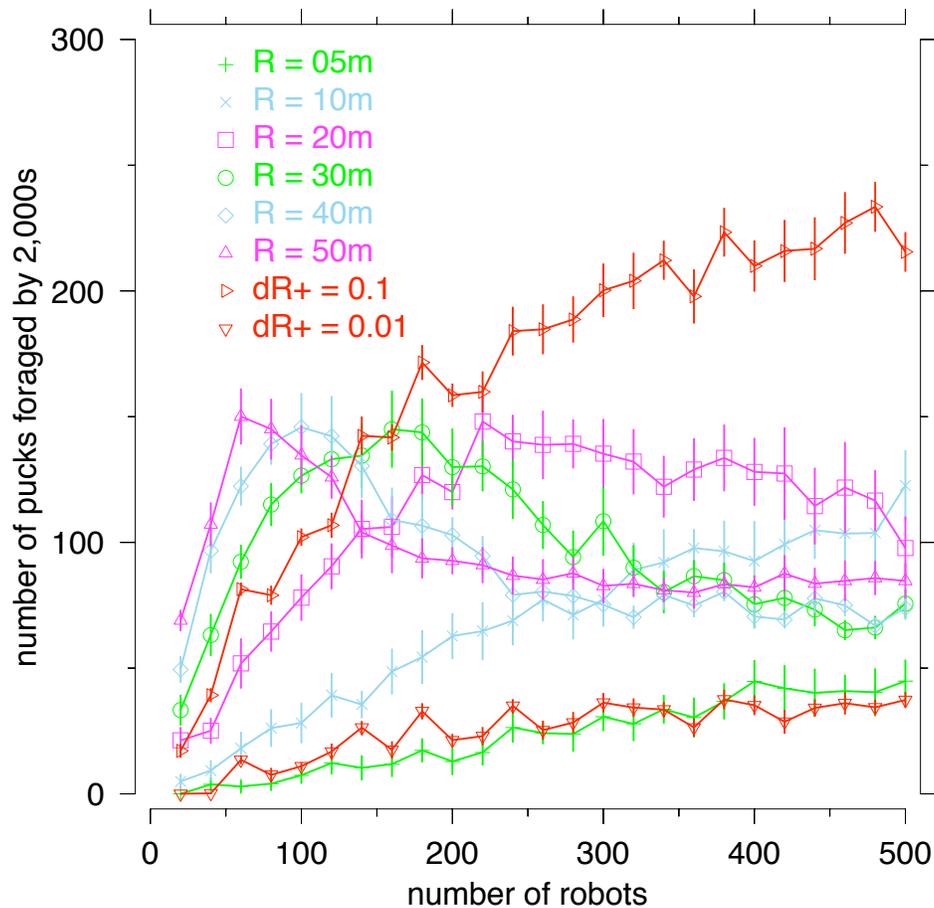


Figure 1: Performance with fixed and adaptive search radii at $0.781 \text{ pucks}/m^2$. R is the search radius of each robot in the trial. The curves labeled dR^+ show the results for adaptive range selection.

working.

In Zuluaga and Vaughan (2005), building on Brown et al. (2005), the problem of spatial interference in multi-robot systems was addressed through the use of aggressive display behaviors. Several robots were required to perform a transportation task (akin to our foraging task) in shared space. Robots selected an “aggression level” based on the amount of work they had invested up to that point. The discrepancy in aggression levels between interfering robots was used to break the symmetry that would otherwise have lead to deadlock. The authors showed their approach to be effective, both in simulation and in a real-world implementation.

Lerman and Galstyan (2002) formally modeled the effect of interference on the performance of a swarm of foraging robots. Their model formulated as a system of coupled first-order nonlinear differential equations. They found that group performance grows sublinearly with group size, so that individual performance actually decreases with increasing group size. Simulations verified the predictions of

their model.

Rybski et al. (2004) performed experiments in which real robots perform a foraging task using a variety of simple communication methods. Robots communicated by flashing a light bulb under various circumstances. The authors showed that communication can reduce the variance in the robots’ performance. In contrast, this research does not use explicit communication between robots; communication is implicit, however, in that robots must alter their behavior in the short term in response to the presence of other robots (collision avoidance), and in the long term by adapting a parameter of their behavior (discussed later).

Simulation

In order to reproduce the results of Shell and Matarić (2006), and compare them to the performance of the adaptive bucket brigaders, we developed a simulator similar to that in Shell and Matarić (2006). A description follows:

Robots are located in an arena, a 64 meter-square plane

scattered with pucks, in the northeast corner of which plane is a 3 m quarter circle “home area”. Robots are equipped with grippers which they can use to retrieve these pucks, and when a puck is dropped off in the home area, it is said to have been foraged.

Robots can move forward at a rate of 0.1 m/sec, and can turn to either side at a rate up to once every five seconds. It takes four seconds to retrieve a puck, but a carried puck can be dropped instantaneously.

Robots can sense walls and other robots within 0.6 m of their centers without error (compared with 0.5 m in Shell and Matarić (2006)), via the use of 12 radially oriented sensors. These sensors report the range in meters to the nearest obstacle (wall or robot). Specifically, each robot is idealized as a line segment, and if the center or either endpoint of that line segment is within sensor range, the sensor most closely oriented towards that point reports the distance to that point.

Robots cannot sense pucks unless the puck in question is located directly within the grip of their grippers, and this sensing is binary and also without error. Robots can determine the direction towards the home zone at any time; Shell and Matarić (2006) explain this as the use of a “four-bit compass”. Robots use the data from their proximity sensors to avoid walls and other robots.

Shell and Matarić (2006) add a variety of noise to each parameter in their simulation. For simplicity, we have dispensed with this noise, except in the case of odometry noise (described in greater detail below). While this does detract from the physical realism of the experiments, the conclusions in this work are drawn by comparing two controllers in identical worlds; we do not attempt to compare our new controller with one tested under different conditions.

Parametrized bucket brigading

The purpose of the overall robot system is to retrieve pucks and deliver them into the home area – to forage the pucks. In the bucket brigade approach to this problem, individual robots do not attempt to carry a puck all the way to the home zone themselves, but rather merely to shift the distribution of pucks towards the home area.

Each robot will attempt to stay within a fixed distance from its initial location. This zone is known as the robot’s “work area”. Through odometry, the robot can determine how far it is from the center of this zone, and can tell the direction towards the center of the zone. However, this odometry is noisy; as a result, the center of each robot’s work area drifts on a random walk at a rate of 0.01 m/sec.

The robot searches via a naïve algorithm within its work area. If it ever leaves the work area, it will drop off any puck it may be carrying, return towards its work area, and continue searching. If it ever discovers a puck, it will retrieve it and head towards the home area. The effect is that a “brigade” of similarly-behaving robots will slowly transport pucks from one robot (work area) to the next, gradually

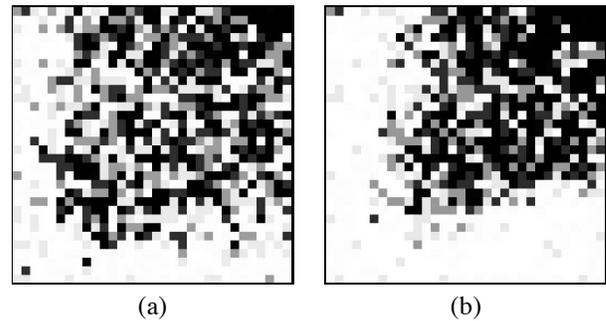


Figure 2: Typical density distribution of pucks after (a) 40 minutes and (b) 80 minutes. One square is 2 m × 2 m; darker squares indicate a higher concentration of pucks. The home area is in the northeast corner of the world. In this simulation, pucks were not added to the world after the start of the experiment.

bringing the puck closer to the home area. Of course, ultimate delivery of the puck requires a connected sequence of overlapping work areas ending in the home area, but this may be achieved over time (even if never simultaneously) due to the drift of the work areas.

Shell and Matarić’s robots all share the same work area radius, or “range”. In the following sections, we will explore other approaches to assigning these ranges to robots. In any given experiment, every robot uses the same approach to range selection.

Radial parametrization

Given that the emergent behavior of the robots is to shift the distribution of the pucks towards the home area, the first natural modification to the controller would be to allow the range parameter of robots to vary with the distance to the home area. Effectively, we are discarding the assumption that the distribution of pucks is uniform, but supposing that they may be more densely distributed near the home area, and therefore a different range parameter would be optimal. To demonstrate this, we simulated robots whose range parameter varied linearly with distance from the home area.

Let us consider the following idealized, one-dimensional situation: robots and pucks sit uniformly distributed on a line of length L , with the home zone at one end. Every robot collects a puck at the same time, drives a fixed distance D towards the “home” end, deposits its puck, and then returns to its starting location, at which point the process repeats. Let $p_n(r)$, where r is the distance from a given point to the home end, be the density of pucks at that point after n of these cycles has taken place. $p_0(r)$ is some constant (the initial puck density). Then

$$p_{n+1}(r) = \begin{cases} (1 - c)p_n(r) & \text{if } r > L - R, \\ (1 - c)p_n(r) + cp_n\left(\frac{Lr}{L-D}\right) & \text{otherwise.} \end{cases} \quad (1)$$

In our simulation, an analogous process occurs in two dimensions. Puck densities at two time-steps are displayed in Figure 2 during a typical run of an experiment in which pucks are not added to the system as fast as they are removed from the system by foraging (pucks in the other experiments in this work are added just as fast as they are foraged, to maintain constant average puck density).

In all experiments, pucks are initially distributed at random. However, it can clearly be seen that as soon as the robots interact with the pucks, the distribution becomes less random, biased toward the home area — the system has a form of entropy that decreases as a result of the work of the robots. Since the optimal search radius for a robot depends on the density of pucks, once the density of pucks changes, the robot's original choice for search radius may no longer be optimal. Ideally, robots would know and select the optimal choice for search radius on an ongoing basis, but in these experiments, robots are not given enough information to achieve this ideal.

In the next section, we propose a method for approximating this ideal.

Adaptive range selection

The aforementioned ideal puts an extra burden on the robot — it must be constantly aware of the distance between the center of its zone and the center of the home area. Alternatively, it must be able to measure the local puck density. In place of these assumptions we may also allow robots to adaptively select their range parameter using purely local information. In *adaptive range selection*, a robot will continuously increase its range parameter at rate dR^+ , except while it is avoiding a collision with another robot, to which situation the robot will react by shrinking its zone at rate dR^- , thus making it less likely to interfere with other robots in the future.

Consider the extreme example of a robot alone in the 64×64 meter-square arena. The robot's work area has some initial radius - say, 10m. The robot will remove pucks from his work area and carry them outside in the direction of the home area. In parametrized bucket-brigade foraging, once the robot has removed all pucks from his work area, there will be no work for him to do, but he will not know this, since he cannot sense pucks, or their absence, from afar — so he will continue searching. His restricted search radius doesn't improve efficiency since there are no robots to interfere with his navigation. Adding adaptive search radii into the picture, the robot's search radius grows at a rate of dR^+ and never shrinks (in practice, we limit the growth so that the radius of any search area is at most the diagonal of the arena).

Now consider the addition of a second robot into our example. As long as the two robots stay outside of each others' sensor ranges, their search areas will continue to grow as before; this scenario will be the same as the above-described

scenario. However, each robot is sensitive to other robots that come within a certain range, less than their sensor range. If such an encounter occurs, each robot will take action to avoid colliding with the other. At the same time (i.e. as long as the collision-avoidance behavior persists), each robot's search area will decrease at the rate of dR^- . Eventually, it may happen that one robot's search area shrinks so that the robot is no longer inside it; at this point, instead of continuing to search for pucks, that robot will try to return to his search area, thus lessening the chance that he will encounter and interfere with the other.

As mentioned above, no robot's search area will grow without bound — there is a maximum useful radius (the diagonal of the arena). In addition, no robot's search area is allowed to decrease below the space needed for the robot to drive in a full circle (at fixed forward/turn speeds).

Results for adaptive range selection are included in Figures 1 and 3.

Experimental design

Initially, we followed Shell and Matarić (2006) in experimental design. The following parameters were varied: puck density ($0.781/m^2$ and $3.125/m^2$), search area radius (5, 10, 20, 30, 40, or 50 m), and number of robots (20, 40, 60, 80, ..., 500). The task was simulated for each combination of parameters for 2000 simulated seconds, and the number of pucks foraged after that time was recorded. Twenty such trials were run, each with a different initial distribution of pucks; to control for robot position, robots were initially placed on a square lattice. The reported results, in Figures 1 and 3, are the averages of those twenty trials. Error bars indicate the standard deviations of the twenty-trial experiments.

Next, we tested our adaptive range selection controller using the same experimental setup. For these experiments, search radii were allowed to increase by $dR^+ = 0.1$ m/s and decrease by $dR^- = 0.05$ m/s, biasing the robots towards the limit of homogeneous foraging in the absence of significant interference. Each robot began with a small range of 5 m. For each parameter set, twenty trials were run, and mean performances were plotted with standard deviations shown.

Results

Our first step was to reproduce the results in Shell and Matarić (2006). Inspection of the data in Figures 1 and 3 indicates that this was accomplished, in that increasing the radius of robots' search areas in the fixed-radius regime led to an increase in the marginal benefit of adding robots (i.e., to the slope of the curves in those graphs), but only up to a point: eventually, adding more robots decreases the performance of the system as overcoming interference begins to dominate the robots' behavior. These critical points are clearly visible as the significant local maxima in the $R = 30$ m, $R = 40$ m, and $R = 50$ m curves.

Fixed range, density=3.125

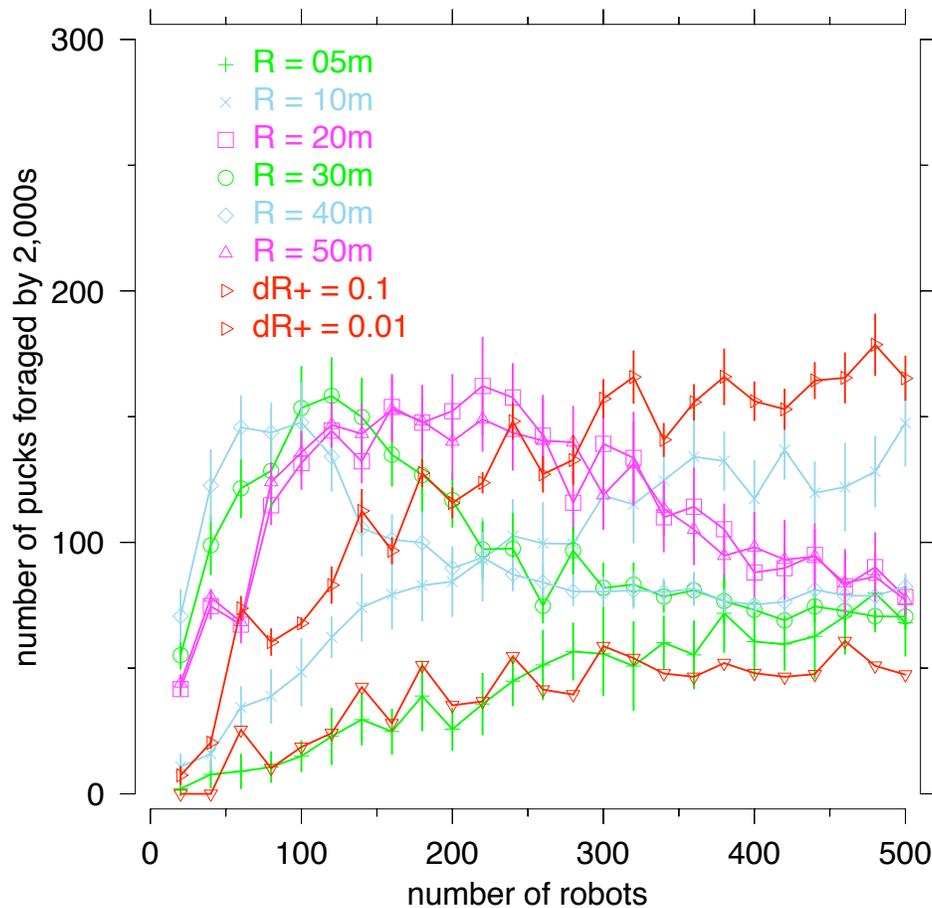


Figure 3: Performance with fixed and adaptive search radii at 3.125 pucks/ m^2 . R is the search radius of each robot in the trial. The curves labeled dR^+ show the results for adaptive range selection.

There is a noteworthy distinction in that robots in our experiments only foraged approximately half the pucks that those of Shell and Mataric (2006) did; this indicates more a quantitative difference in the efficacy of the robots' controller programs than a qualitative failure of our simulations to agree at the heart of the matter: that interference affects a growing population later when the individuals' foraging spaces are larger, which is indicated by the relative shapes of the curves.

Our adaptive range selection algorithm performed at least as well as the fixed ranged controllers in simulation, and scaled better. Fixed range algorithms suffered from one of two problems: robots with small search areas did not gather many pucks, and robots with large search areas interfered too much and the critical point at which the marginal benefit of increasing the number of robots was reached when the number of robots was still small. While the robots using adaptive range selection did not gather as many pucks when the number of robots was small as did robots with large,

fixed search radii, increasing the number of robots always increased the performance of the group.

Also noteworthy is that the adaptive controllers performed more consistently, as indicated by tighter error bars on those curves than on the fixed-radius performance curves. Since adaptation to interference is a form of implicit communication, this is in agreement with the findings of Rybski et al. (2004).

Adaptive range selection was sensitive to variations in dR^+ and dR^- . If dR^+ was too small, adaptive selection underperformed the fixed range foragers. Figure 1 shows results when $dR^+ = 0.01$, a fifth of dR^- . In that case, the adaptive controller performs no better than the worst-performing fixed-range controller we tested. This is not altogether surprising, since the search radius in the worst fixed-range controller and the initial search radius in the $dR^+ = 0.01$ m adaptive controller were both 5 m.

Figure 3 shows qualitatively similar improvements; however, in this scenario (where puck density is 3.125 pucks per

square meter, the improvement is only slight, even for large group sizes).

Future work

In this paper, we have explored only some of the ways — and naïve ones at that — of improving on the bucket brigading algorithm. Future work may explore generalization of the problem to discard hidden assumptions, or, increasing the adaptability of the system. For instance, a more plausible biological analog might be ants foraging for food. In such a scenario, the ants would not initially be uniformly distributed through the field, nor would the food.

The performance of the adaptive-parameter bucket brigade forager needs to be more deeply tested. In none of our experiments did we note negative marginal performance; the limits of the algorithm in terms of scalability remain to be seen. The parameter space for the adaptive controller — the possible values of dR^+ and dR^- — needs to be explored in greater depth. We should also test the controller in more topologically complex spaces, such as corridors, as did Østergaard et al. (2001), and investigate why the adaptive controller does not outperform as significantly in denser environments.

Briefly mentioned above is the point that the adaptive approach causes a net increase in the number of a priori parameters: new parameters added are the rates at which zones increase and decrease in size. In this work, those parameters were set experimentally. Future work may provide motivation behind values of these parameters, investigate a relationship between optimal values for dR^+ and dR^- , or do away with these parameters entirely.

Future work may also formalize a closed-form relationship among optimal search radius, puck density, and robot density. Additionally, we have briefly touched on a notion of entropy in robots' work space — a quantitative function of the environment that decreases through *useful* interactions between agents and the environment — this notion should be more formally analyzed.

Conclusions

To summarize the contributions of this paper: we replicated the results of Shell and Matarić (2006), confirming their findings with an independent implementation. Further, we propose a simple modification of their foraging scheme in which each robot's foraging area is adapted in response to interference. The new method was shown to improve performance, particularly in large population sizes.

Acknowledgments

Special thanks to Jens Wawerla, for saving me from learning the hard way, and to Yaroslav Litus, for braving the cluster with me.

References

- Beckers, R., Holland, O., and Deneubourg, J.-L. (1994). From local actions to global tasks: Stigmergy and collective robotics. In Brooks, R. and Maes, P., editors, *Artificial Life IV*, pages 181–189. Cambridge, MA: MIT Press.
- Brown, S., Zuluaga, M., Zhang, Y., and Vaughan, R. (2005). Rational aggressive behaviour reduces interference in a mobile robot team. In *Proceedings of the International Conference on Advanced Robotics (ICAR)*, Seattle, Washington.
- Goldberg, D. and Matarić, M. (2003). Maximizing Reward in a Non-Stationary Mobile Robot Environment. *Autonomous Agents and Multi-Agent Systems*, 6(3):287–316.
- Holland, O. and Melhuish, C. (1999). Stigmergy, Self-Organization, and Sorting in Collective Robotics. *Artificial Life*, 5(2):173–202.
- Lerman, K. and Galstyan, A. (2002). Mathematical Model of Foraging in a Group of Robots: Effect of Interference. *Autonomous Robots*, 13(2):127–141.
- Østergaard, E., Sukhatme, G., and Matarić, M. (2001). 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.
- Rybski, P., Larson, A., Veeraraghavan, H., LaPoint, M., and Gini, M. (2004). Communication strategies in Multi-Robot Search and Retrieval: Experiences with Min-DART. In *Proceedings of the Seventh International Symposium on Distributed Autonomous Robotic Systems (DARS-04)*, pages 301–310.
- Shell, D. A. and Matarić, M. J. (2006). On foraging strategies for large-scale multi-robot systems. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 2717–2723, Beijing, China.
- Wawerla, J. and Vaughan, R. T. (2007). Near-optimal mobile robot recharging with the rate-maximizing forager. In *Proceedings of the European Conference on Artificial Life*.
- Zuluaga, M. and Vaughan, R. (2005). Reducing spatial interference in robot teams by local-investment aggression. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, Edmonton, Alberta.