

On foraging strategies for large-scale multi-robot systems

Dylan A. Shell and Maja J. Matarić
Interaction Laboratory
Center for Robotics and Embedded Systems
Department of Computer Science, University of Southern California
Los Angeles, CA USA 90089-0781
{shell|mataric}@usc.edu

Abstract—Physical interference limits the utility of large-scale multi-robot systems. We present an empirical study of the effects of such interference in systems with hundreds of minimalist robots. We consider the canonical multi-robot foraging task, and define a new parametrized controller. This controller allows for evaluation of spatial arbitration strategies along a continuum with the traditional homogeneous and bucket-brigading algorithms at each end. We present data from thousands of simulations which suggests that methods surprisingly close to homogeneous foraging, but augmented with limited arbitration, can improve both performance and reliability.

I. INTRODUCTION

Typical multi-robot systems currently consist of tens of robots. This research explores the performance of existing algorithms in large-scale systems with hundreds of robots. Such systems raise several challenges that must be overcome. One key challenge is the issue of resource competition: as a finite resource is shared among an increasing number of robots, techniques are needed to deal with contention. In particular, spatio-temporal interference imposes limits on the number of robots that are useful for task achievement. Since all robots have spatial extent, this is not a problem that can be easily bypassed. This paper presents an empirical study of the performance effects of interference in a multi-robot system, with hundreds of simple robots performing a collective foraging task.

Multi-robot foraging, as studied here, involves several robots collecting randomly distributed objects (pucks) and transporting them back to a single location (home region) within a planar arena. This idealization of collection and transport tasks has several applications including mine-clearing, hazardous waste clean-up, and search and rescue. Multi-robot foraging is the problem domain most widely used to study group size scalability and hence the effects of inter-robot interference; see for

example Arkin et al. [1], Goldberg [4], Lerman and Galstyan [5].

Traditional *homogeneous foraging* has each robot searching for pucks and independently transporting them to the home region. Once any of the robots searching for a puck finds one, that robot will deliver the puck to the home region; the robot then begins searching anew. As one might expect, such an approach produces spatio-temporal interference that is strongly concentrated around the home region; many robots will attempt to enter the same space and are forced to issue obstacle avoidance commands. Thus, rather than performing task-related actions, additional robots may hamper the collective effort. In order to increase the effectiveness of a system with a high density of robots, the *bucket brigading* strategy has been proposed [4, 8]. This strategy requires that each robot focus on a sub-region of the total work arena. Any robot finding a puck will transport that puck to the neighboring sub-region in the direction of the home region. Thus, pucks are passed from robot to robot along toward the home region and overcrowding is reduced.

By permitting sub-regions assigned to each robot to overlap, we are able to consider bucket brigading and homogeneous foraging as two ends of a single *strategy continuum*. We present data from simulations of foraging performance along this continuum. Our experiments compare a range of robot group sizes and differing puck densities. The data also consider a much wider range of system sizes than found elsewhere in the literature.

Analysis suggests that a balance exists between interference arbitration and greedy puck delivery. Our data indicate that even large overlapping sub-regions can have a positive effect on overall performance. Also the results point to a trade-off between robustness and adaptability across puck densities and maximal performance with small groups of robots (less than 150).

II. RELATED WORK

Foraging is a canonical task in distributed robotics and is one of the most widely studied problem domains. Arkin et al. [1] presented the first simulation results showing the effect of group size on robot foraging performance, and in particular, the role of robot density and critical group sizes. Matarić [7] performed early physical robot experiments and described the ethological inspirations for the task. Østergaard et al. [8] defined a taxonomical framework in which to organize variations on foraging, as well as related task domains of clustering and collection (for an example of such domains, see Beckers et al. [3]).

Parker [9] considered a hazardous waste cleanup application in order to demonstrate the robustness of her ALLIANCE software architecture. In that work, robots communicated explicitly with one another. Both implicit and explicit communication strategies have been studied in the context of foraging [1, 2, 11, 10]. The robots in this work do not have explicit communication capabilities and other robots are merely perceived as obstacles. This is in the spirit of minimalist robotics, and is in anticipation of the types of systems expected to scale to hundreds of robots with least effort.

Lerman and Galstyan [5] constructed a differential equation model of foraging in a swarm of minimalist robots. They demonstrated good qualitative agreement between the model and experiments (including interference effects) with up to 20 robots. This style of model has been extrapolated to large systems with hundreds of robots (e.g., [6, pp. 12]).

Perhaps the definitive study of interference in multi-robot foraging was done by Goldberg [4], in which four categories of multi-robot interactions are identified: SPST, DPST, SPDT, and DPDT (where S \equiv same, D \equiv different and P \equiv place, T \equiv time). Goldberg notes that interference around the foraging home region is physical interference, a characteristic SPST interaction. Such interference can be reduced by arbitrating resources so that either:

- 1) robots are forced to stay in a different places at all time (DPST); or
- 2) scheduling is done so as to ensure no location is used by two robots simultaneously (SPDT).

Bucket brigading is arbitration of the first type, while homogeneous foraging does not implement arbitration.

Of course, distributed resource conflicts arise in natural systems, too. Vaughan et al. [12] proposed a general methodology for dealing with such circumstances through a stylized aggression game. Their technique (and




Fig. 1. Trajectories of three robots using: (i) homogeneous foraging ($T \sim \text{Uniform}(10, 40)$), and (ii & iii) bucket brigading ($D = 5$ and $D = 15$) controllers.

the recent follow-up [13]) is applied to break symmetry in SPST circumstances.

We measure the significance of DP arbitration on foraging performance, by considering a progression of bucket brigading parameterizations. Treating arbitration as a binary property may be an over-simplification; the present research aims to address this issue.

III. METHOD

We simulated homogeneous groups of minimalist robots across a number of sizes. Detailed simulator and robot controller specifications are given next, followed by a description of the experiments themselves.

A. Simulator

We implemented a custom tool in order to efficiently simulate the kinematics of large numbers of planar robots. Simulated robots are rectangular with sides $35\text{cm} \times 10\text{cm}$, pucks are circular with the radius of 3cm . The simulated environment is an open $64\text{m} \times 64\text{m}$ arena with a perimeter wall, and a quarter-disk with radius 3m in the North-East corner of the area represents the foraging home region.

Each robot has forward-pointing scoop for collecting pucks (included in the 35cm length). At most one puck can be held at any given time. The simulator provides a velocity control interface to the controller software. Commanded linear (\dot{r}_c in m/sec) and angular ($\dot{\theta}_c$ in degree/sec) velocities are corrupted by noise. Table I contains the details. Robot positions are integrated in random order, reshuffled for each time step, with P_{update} giving the probability of a robot's position being propagated.




Fig. 2. Performance data for homogeneous foraging with different exiting times. The exit duration is drawn from a uniform distribution ($T \sim \text{Uniform}(10, 40)$). Left plot is for low puck density, the right plot is for a higher density.

Parameter	Value	Comment
$\dot{r} = (1 + \xi_v) \dot{r}_c + \eta_v + \zeta_v$		Linear velocity integrated per time step in meters per second
\dot{r}_c		Commanded linear velocity
ξ_v	$\sim N(0, 0.01^2)$	Multiplicative noise, drawn each time step
η_v	$\sim N(0, 0.03^2)$	Additive linear noise, drawn each time step
ζ_v	$\sim N(0, 0.002^2)$	Per robot linear bias, drawn each time step
$\dot{\theta} = (1 + \xi_r) \dot{\theta}_c + \eta_r + \zeta_r$		Angular velocity integrated per time step in degrees per second
$\dot{\theta}_c$		Commanded angular velocity
ξ_r	$\sim N(0, 0.01^2)$	Multiplicative noise, drawn each time step
η_r	$\sim N(0, 2.5^2)$	Additive rotational noise, drawn each time step
ζ_r	$\sim N(0, 0.15^2)$	Per robot rotational bias, drawn each time step
P_{update}	0.8	Probability of robot being propagated for a given time step

TABLE I

SIMULATION INTEGRATION AND CONTROL PARAMETERS

Each robot is provided with a distance sensor. Noisy values are returned from 12 radial rays, each giving a distance reading to obstacles and robots (but not pucks) up to the maximum distance of 0.5m. Each robot is equipped with a four-bit compass with noise added. Robots have a binary sensor that detects position within the home-region. False negatives are returned with probability P_{HR}^- and false positives with probability P_{HR}^+ . Each of the robots also has a binary sensor to detect the presence of a puck within the scoop, but which gives false negatives with probability P_{scoop}^- and false positives with probability P_{scoop}^+ . Robots are unable to detect pucks at a distance. See Table II for complete details of noise models.

B. Controllers

Two controllers were implemented, both used the same low-level obstacle avoidance, odometry integration and sensor smoothing code.

1) *Homogeneous foraging*: The implementation comprises three states. The first, called searching, is active while the robot seeks a puck. When a puck is detected, this triggers homing which moves the robot toward the home region using the compass (*a priori* known to be North-East). This homing strategy works without requiring localization information on the robots. Once over the home region, the robot enters an exiting state in which it reverses briefly allowing the puck to leave the scoop. The robot performs a random turn, and navigates away from the home region (using the compass) for some random time (described with parameter T). The robot then transitions to the searching state. This simple controller means that if any robot fails, it is simply ignored by others.

2) *Parametrized bucket brigading*: The controller uses three states: searching, homing, and returning. The key difference is in the use of integrated odometric data. The controller is parametrized by D , the radius of each robot's sub-region within odometric space (in meters).

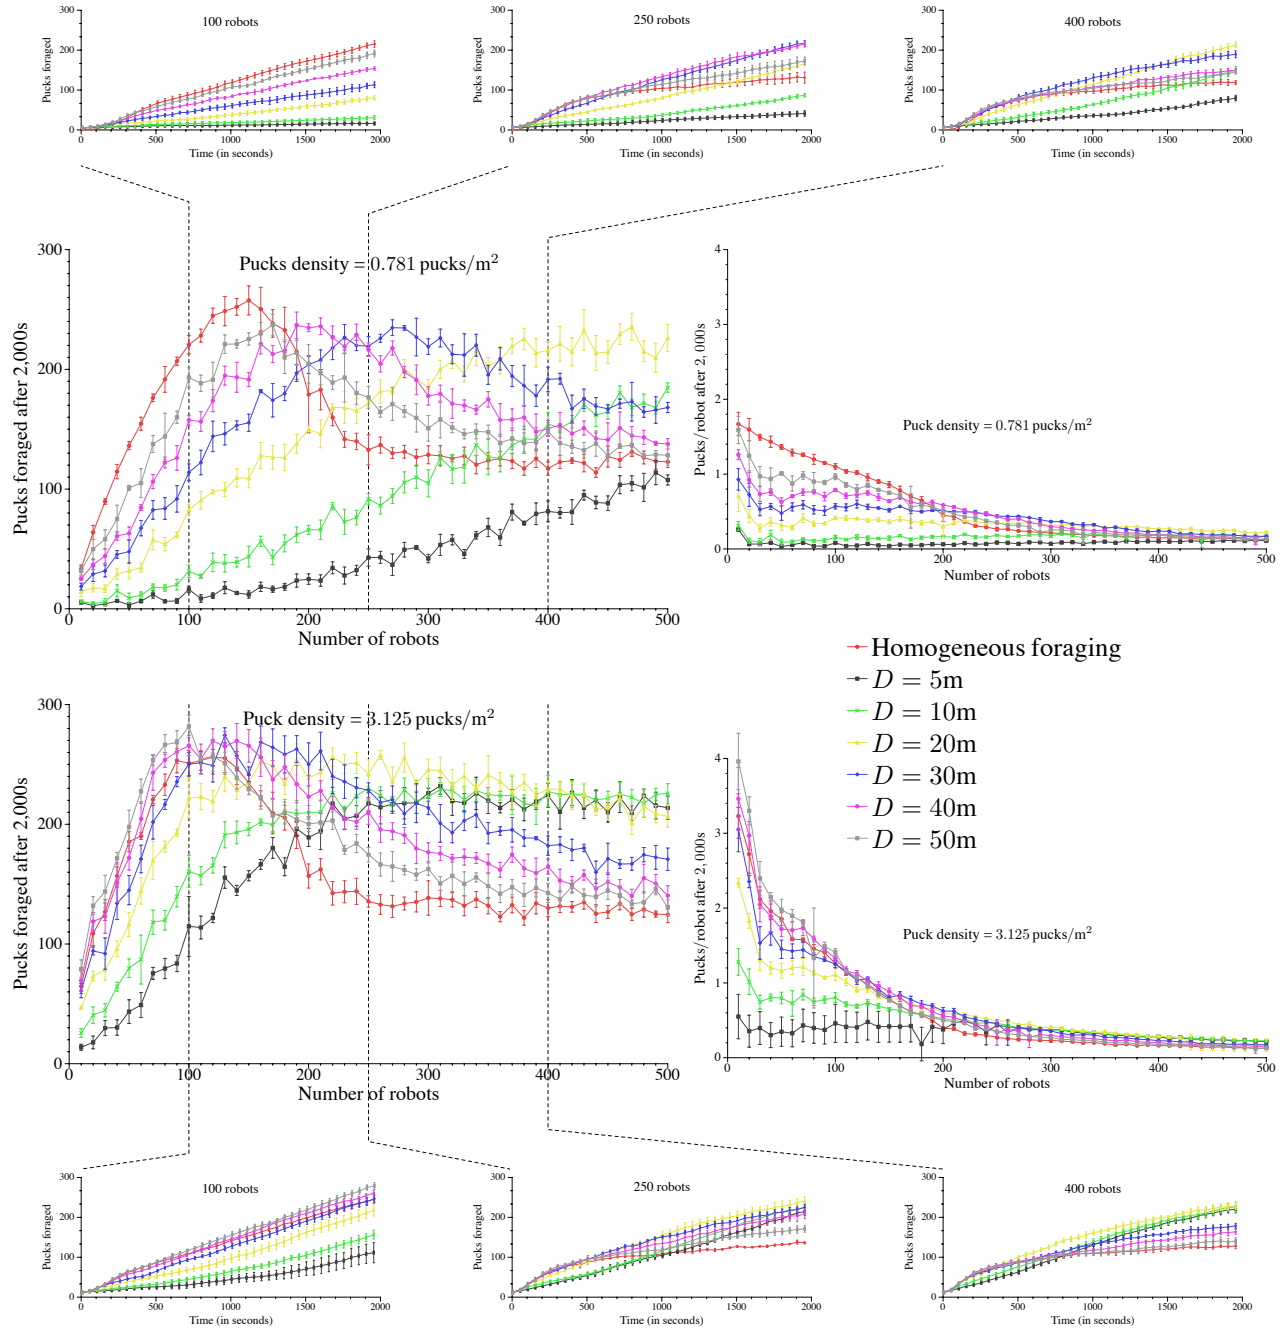


Fig. 3. Performance of parametrized bucket brigading controller across system sizes. Top half of figure (five plots) gives data on low puck density case (0.781 pucks/m²), bottom half the high density case (3.125 pucks/m²). Two large plots on the left give performance (number of pucks foraged after 2000s) for robot groups ranging from 10 to 500 robots. Medium sized plots on the right give performance per robot for each of the D parameters. The size small plots, three along the top and three along the bottom, give time series data for 100, 250 and 400 robots. All data are averages of 5 independent runs, error bars show one standard deviation. Homogeneous foraging has $T \sim \text{Uniform}(10, 40)$.

Parameter	Value	Comment
$\text{dist}_{\text{ret}} = (1 + \xi_d) \text{dist}_{\text{truth}} + \eta_d + \zeta_d$		Radial range reading returned after corrupted by noise
$\text{dist}_{\text{truth}}$		True distance between robot and obstacle
ξ_d	$\sim N(0, 0.02^2)$	Multiplicative noise, drawn each time step
η_d	$\sim N(0, 0.05^2)$	Noise (in meters) drawn for each sensor reading
ζ_d	$\sim N(0, 0.05^2)$	Per ray bias, drawn for each ray at initialization time
$\text{comp}\angle_{\text{ret}} = \text{comp}\angle_{\text{truth}} + \eta_c + \zeta_c$		The compass bearing returned after corrupted by noise in degrees
$\text{comp}\angle_{\text{truth}}$		True angle between robot and magnetic North
η_c	$\sim N(0, 15.0^2)$	Noise (in degrees) drawn for each sensor reading
ζ_c	$\sim N(0, 2.0^2)$	Per robot bias, drawn at initialization time
P_{HR}^-	0.15	Probability for a home region detector false negative
P_{HR}^+	0.08	Probability for a home region detector false positive
P_{scoop}^-	0.05	Probability for a puck scoop false negative
P_{scoop}^+	0.05	Probability for a puck scoop false positive

TABLE II

SIMULATED SENSOR PARAMETERS

The searching state randomly explores the arena, but once outside the disk of radius D (calculated from odometric data), the controller transitions to returning state. If a puck is found within the bounds of the sub-region, the homing state is triggered. The returning state navigates the robot back to a position within the sub-region. Our implementation is conservative, ensuring that the robot is within the disk of radius $D/2$ before transitioning to searching. The homing state simply moves the robot toward the home region. If the robot detects that it is over the home region, the puck is deposited identically as in homogeneous foraging; if the robot exits the D -disk, the puck is deposited and returning commences. Thus, bucket brigading with $D = \infty$ corresponds to the homogeneous controller.

Robots have only \dot{r}_c and $\dot{\theta}_c$ values (and neither \dot{r} nor $\dot{\theta}$) so the odometric estimates of each robot's sub-region drifts over time. Figure 1 shows the effect of this drift (for a single robot) and how this compares with homogeneous foraging. Since all robots' estimates are drifting, the regions themselves perform (slow) independent random walks over the arena. This has the added effect of ensuring coverage of otherwise neglected areas (and adds robustness in the case of robot failures).

C. Experimental Design

1) *Experiment 1: Effect of return time:* An initial experiment was conducted to compare the effect of different return time on performance of the homogeneous controller. We simulated two parameter values:

$T \sim \text{Uniform}(10, 40)$ and $T \sim \text{Uniform}(20, 80)$, and considered puck densities of 0.781 pucks/m^2 and 3.125 pucks/m^2 (corresponding to 3,200 and 12,800 total pucks, respectively). Experiments considered robot group sizes from 10 to 500 robots in increments of 10. Five trials were run of each case.

2) *Experiment 2: Effect of parameter D :* We simulated bucket brigading with D values $\{5\text{m}, 10\text{m}, 20\text{m}, 30\text{m}, 40\text{m}, 50\text{m}\}$, which is a wide range of values considering that the arena was 64m across (these values must always be interpreted relative to the size of the environment). Again we considered group sizes from 10 to 500 robots in increments of 10, and low and high puck density conditions.

IV. RESULTS

A. Experiment 1

Figure 2 shows the measured performance for Experiment 1. Each data point is the mean of five trial runs; the error bars show a single standard deviation. As might be expected, the larger mean T values result in decreased performance. Mean T is doubled, but the performance effect is limited. We conclude that the T parameter has a limited effect on the qualitative characteristics of the performance curve.

B. Experiment 2

Results from Experiment 2 are displayed in Figure 3. The plots show the collective foraging performance for a range of D values; homogeneous foraging is also

shown for reference purposes. The two large plots give collective performance across the range of robot system sizes. The six surrounding smaller plots show time-series data for systems with 100 robots (left two), 250 robots (middle two), and 400 robots (right two). The upper plots are for a low puck density; the lower plots have a higher (by a factor of four times) puck density.

These graphs display several results. Bucket brigading is shown to be more sensitive to puck densities than homogeneous foraging. The robots must have sufficient probability of finding another puck in order to drop one at a neighbor’s sub-region. This effect depends on puck density, and the minimalist nature of the simulated robots, specifically the lack of puck sensing at a distance, requires higher densities than found elsewhere (e.g., Goldberg [4]). Bucket brigading (especially with small D) has gradual performance increases, and is inefficient with few robots, where for example disks of radius D may be insufficient to cover the arena.

With increasing D values the bucket brigading controller approaches the performance of the homogeneous controller (which is $D = \infty$, and as the findings of experiment 1 suggest, an exact correspondence with a matching T value is unnecessary). This is as expected as it was anticipated during the design of the parametrized controller. What is unexpected, however, is that even large D values (40m is large, relative to the size of the arena) can have a marked interference-reducing effect. For large D values, several robots will still attempt to transport a puck toward the home region, but they do not remain in the crowded region as long as homogeneous robots. We believe that this is because, as can be seen in Figure 1, odometric drift forces some robots to yield, having believed to have passed outside the radius D disk. The drift is exacerbated around the crowded home region because much time is spent turning in place. In physical robots, such turns often distort the odometric frame of reference negatively, which is faithfully reproduced in our simulation. This factor appears to be sufficient to break deadlocks within groups of otherwise greedy robots crowding the home region and areas surrounding it.

For each value of D , the width of the performance curve depends on puck density. Zhang and Vaughan [13] provide an excellent discussion of the significance of general performance versus population curves. For a given minimum performance level, they interpret the width (i.e., where the curve intersects line of constant performance) as the range of robot groups sizes that can achieve the task with at least the given performance. Greater width can be useful when wanting to deploy a

system with some redundancy. A trade-off is apparent in Figure 3; in choosing the D parameter value one can either have high performance with few robots, or consistent performance across a range of medium to large numbers of robots. In the latter case, performance depends critically on the puck density. A value like $D = 30\text{m}$ seems like a good compromise: steep initial slope for between 10 and 100 robots, but gradually decreasing performance beyond the optimal number.

The time series plots in Figure 3, particularly for 450 robots, show an initial rate of puck delivery that is unsustainable. The initially steep performance graph flattens out after approximately 200 seconds. This occurs because, initially, the robots are randomly placed within the environment. When early pucks are delivered, there is little interference, because few robots have arrived at the home region. Later deliveries must face exiting robots, and if there are sufficient robots, overcrowding results in a traffic jam scenario. This is the reason for the performance curves flattening with increasing numbers.

V. LIMITATIONS AND FUTURE WORK

There are limits to the applicability of bucket brigading (or more generally DP arbitration). If the cost of picking up or dropping pucks is significant, for example, when precision manipulation is required, then bucket brigading may not be suitable. Interference problems in large-scale multi-robot systems remains significant and further solutions are required.

Our study, like Goldberg [4], considers an open arena in order to study interference caused solely by other robots. Østergaard et al. [8] considered the role of a complex environment which can result in high-density choke points. Additional work is necessary to study non-deterministic and dynamic environments.

Future work will consider significance in the rate of odometric drift in the bucket brigading controller because this affects the “mixing” of sub-regions that occurs.

VI. SUMMARY AND CONCLUSION

This paper is an empirical study of physical interference in large-scale multi-robot systems within the foraging task domain. We defined a parametrized controller that reproduces both homogeneous and bucket brigading strategies. By exploring the parameter space between these two known strategies, we were able to evaluate importance of spatial sub-division in mitigating performance reducing interference. Simulation experiments sample the strategy continuum across varying numbers of robots (up to hundreds of robots) and differing puck densities. The data suggest that treating strategies as different in

degree, rather than of different types, allows for intermediate controllers that possess desirable performance and robustness properties. We highlight the necessary trade-offs.

Acknowledgments

Support for this work was provided by the National Science Foundation Grant No. IIS-0413321.

REFERENCES

- [1] R. C. Arkin, T. Balch, and E. Nitz. Communication of behavioral state in multi-agent retrieval tasks. In *Proceedings IEEE/RSJ International Conference on Robotics and Automation (ICRA'93)*, pages 588–594, Atlanta, GA, 1993.
- [2] T. Balch and R. C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1): 27–52, 1994.
- [3] R. Beckers, O. E. Holland, and J.-L. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In *Artificial Life IV*, pages 181–189, Cambridge, MA., USA, Jul 1994.
- [4] D. Goldberg. *Evaluating the Dynamics of Agent-Environment Interaction*. PhD thesis, Department of Computer Science, University of Southern California, May 2001.
- [5] K. Lerman and A. Galstyan. Mathematical Model of Foraging in a Group of Robots: Effect of Interference. *Autonomous Robots*, 13(2):127–141, 2002.
- [6] A. Martinoli, K. Easton, and W. Agassounon. Modeling of Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. *International Journal of Robotics Research*, 23(4):415–436, 2004.
- [7] M. J. Matarić. *Interaction and intelligent behavior*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, USA, May 1994.
- [8] E. Østergaard, G. S. Sukhatme, and M. J. Matarić. Emergent bucket brigading - a simple mechanism for improving performance in multi-robot constrained-space foraging tasks. In *International Conference on Autonomous Agents*, pages 29–30, Montreal, Canada, May 2001.
- [9] L. E. Parker. ALLIANCE: An architecture for fault-tolerant multi-robot cooperation. *IEEE Transactions on Robotics and Automation*, 14(2):220–240, Apr 1998.
- [10] P. E. Rybski, A. Larson, H. Veeraraghavan, M. LaPoint, and M. Gini. Communication strategies in multi-robot search and retrieval. In *Proceedings of the 7th International Symposium on Distributed Autonomous Robotic Systems (DARS'04)*, pages 301–310, Toulouse, France, June 2004.
- [11] K. Sugawara and M. Sano. Cooperative acceleration of task performance: foraging behavior of interacting multi-robots system. *Physica D*, 100(3-4):343–354, 1997.
- [12] R. T. Vaughan, K. Støy, G. S. Sukhatme, and M. J. Matarić. Go ahead, make my day: Robot conflict resolution by aggressive competition. In *International Conference on Intelligent Autonomous Systems*, pages 491–500, Paris, France, September 2000.
- [13] Y. Zhang and R. T. Vaughan. Ganging up: Team-based aggression expands the population/performance envelope in a multi-robot system. In *Proceedings IEEE/RSJ International Conference on Robotics and Automation (ICRA'06)*, pages 1050–1059, Orlando, Florida, May 2006.