
Communication strategies in Multi-Robot Search and Retrieval: Experiences with MinDART

Paul E. Rybski, Amy Larson, Harini Veeraraghavan, Monica LaPoint, and Maria Gini

Department of Computer Science and Engineering, University of Minnesota, 200 Union St. S.E., Minneapolis, MN 55455-0159
{rybski,larson,harini,mlapoint,gini}@cs.umn.edu

Summary. To explore the effects of different simple communications strategies on performance of robot teams, we have conducted a set of foraging experiments using real robots (the *Minnesota Distributed Autonomous Robotic Team*). Our experimental results show that more complex communication strategies do not necessarily improve task completion times, but tend to reduce variance in performance.

1 Introduction

Designing a distributed robotic system using simple units is an attractive engineering solution for many reasons [2]. Each robot in the swarm uses simple local rules to decide its actions without needing any command from a central controller or from any other robot. Obvious advantages to this approach are robustness to individual failure, ability to scale with minimal tractability issues, low unit complexity, and decreased costs.

In this study, we are interested in determining what level of improvement in task performance we can expect by adding simple communications capabilities to the robots in the swarm. In order to explore this question, we built a group of simple robots, the *Minnesota Distributed Autonomous Robot Team* (MinDART) shown in Figure 1, to perform a foraging task and we enhanced them with communication capabilities. We conducted a series of experiments with these robots and compared their performance when using different simple communication strategies.

Although many tasks can serve as a testbed, we chose foraging, which is well studied, so that solutions and results can be compared more easily. In our version of the task, robots locate a target in an enclosed arena, pick up the target, and drop it off at a designated home base. The arena contains some obstacles, and the distribution of the targets varies. The performance




Fig. 1. The Minnesota Distributed Autonomous Robotics Team (MinDART). The MinDART robots searched for the infrared emitting targets in a search and retrieval task. Landmarks were used for homing and localization.

criterion we selected is the time to complete the task, i.e. the time to collect the entire set of targets and return them to home.

All our experiments use physical robots, as opposed to simulated robots. As eloquently explained in [6], we believe that a rigorous study of swarm intelligence warrants physical robots, as opposed to simulated robots.

The simplest control strategy for foraging is random walk. We use reactive behaviors to avoid obstacles and random direction changes at random intervals to increase the probability of complete coverage. This strategy is an attractive choice for simple robotic hardware because it is easy to program and requires little sensor bandwidth.

We analyzed random walk versus control strategies that use communication. The communication methods we chose are forms of indirect communication based on cues from the environment (this is called *stigmergy* in the biology literature). We studied two types of communication (reflexive communication and deliberate communication) and studied how the duration of deliberate communication (10, 20, and 30 seconds) affected the time to complete the task. Our experimental results show that for simple robots such as the MinDART, deliberative strategies help in decreasing the variance of the team’s performance. However, this decrease in variability does not correspond to a significant decrease in the mean time to solve the task. Instead of spending time wandering randomly, the robots spend time recruiting other robots.

2 MinDART Hardware and Software

Each MinDART robot is constructed out of LEGO Technic blocks. The robot is 29 cm long by 24 cm wide by 37 cm tall and has a dual-treaded skid-steer chassis that allows the robot to turn in place and translate at a speed of

0.1693 m/s. The gripper is an articulated cargo bay that grasps and transports targets. Bumpers are used for obstacle avoidance. The robot’s infrared sensors can detect the target IR beacons at a range of approximately 70 cm. A light-bulb beacon serves as binary form of communication among the robots. To detect this beacon, and to identify landmarks for homing, a CMUCam [13] is mounted on top of a servo-controlled turret. The camera processes images at 2-3 frames/second. The robot is controlled by a Handyboard. Power is provided to the camera and Handyboard by two 9.6V NiCad battery packs.

The MinDART’s control software consists of a finite state machine. Each state in the controller solves one subtask in the robot’s overall task. There are three subtasks comprising the search and retrieval task, which are find a target, grab a target, and return a target to the home base. In the initial state, `FINDTARGET`, a robot searches for targets, or heads toward an activated light-bulb beacon. Once a target is detected with the robot’s infrared sensors, the control system switches to the `GRABTARGET` state which is responsible for maneuvering the robot such that the target fits into the gripper. If the robot successfully grabs the target, the control system switches to `RETURNTARGET`, which returns the robot to the drop-off location.

3 Communication Experiments

Figure 2 shows a view of the experimental setup. The area was 7 m x 8 m and contained uniformly distributed obstacles. The targets were distributed in a single non-uniform distribution in the corner of the environment furthest from the drop-off location. All experiments were run with four robots. Robots communicated by turning on their light-bulb beacons. Beacons could be seen at a maximum range of 2.9 m.

The goal of communication is to reduce the target search time by attracting robots to the area of a sensed target. The experiments were designed to test the robot’s abilities to lead each other to a single clump of targets. The communication method chosen was an attracting light-bulb beacon, which would direct the other robots toward the targets. Communication varied by intent and duration as follows:

- **No Communication.** This was used as a baseline experiment.
- **Reflexive Communication.** A robot turned on its light-bulb beacon while trying to pick up a target (i.e. while in the `GRABTARGET` state). Once the robot grabbed the target, the beacon was deactivated. We consider this a statement of action, not a request for help. This strategy would be of no use in an environment with a uniform distribution of targets, but we hypothesized it would help when targets are clumped and harder to find by random walk.
- **Deliberative Communication.** A robot turned on its light-bulb beacon when a target was sensed, but the robot was unable to pick it up. This form




Fig. 2. Diagrams of the 7 m x 8 m experimental environment showing obstacles, initial placement of targets, and initial starting point for the robots. All experiments contained nine targets and eight obstacles. The obstacles were relatively low and did not block a robot’s view of the landmarks or of each other. However, they did block a robot’s view of the targets. The drop-off area is the same as the robot starting point.

of communication was used if a robot encountered a target while on its way to the home base to drop off one that it had in its gripper. The robot would stay motionless for a fixed amount of time as a deliberate request for assistance. We tested three fixed durations: 10 s, 20 s, and 30 s.

We hypothesized that any form of communication would provide a performance enhancement, due to the reduced time spent in random search, and that some form of communication would provide better results than no communication. Similar findings are reported in the simulation work of [1].

We also predicted that deliberative communication would provide the most benefit and that there would be a peak or plateau in the duration, as seen in the simulation work of Sugawara [19]. In other words, we predicted that there would be an ideal communication duration that would maximize performance, and any duration longer than that would not enhance performance any further. This is because the longer the beacon is left on, the better chance the other robots would see it. However, deliberative communication requires a robot to stay stationary while recruiting others. There is a tradeoff between this delay and the time spent doing random search.

Since one of the goals of the experiments was to measure the effect of the amount of time the light-bulb beacons were on, and since in the reflexive communication experiments the beacons were on for different amounts of time, we recorded the light-on time for each communication occurrence. The average light-on time was approximately 16 s with a standard deviation of 11.6 s, but the distribution is not Gaussian, as seen in Figure 3. Instead, it clusters around

5 and 10 s (the mode of the distribution is 5 s). We will see later how this is correlated with the variance in the experimental results.




Fig. 3. Histogram of the times the light-bulb beacon was on, observed in the reflexive communication experiments. The majority was in the 5 to 10s range. This can be compared to the deliberative communication experiments, where beacon-on times were fixed.

For each of the experiments, we recorded the time a robot returned a target to the drop-off zone. The results were averaged over five runs. Each experiment was run until all nine targets were retrieved. We compared the times between the dropping off of the first and eighth target, to discount the times in the experiment when communication had little effect. Figure 4 shows the means and standard deviations of these times.

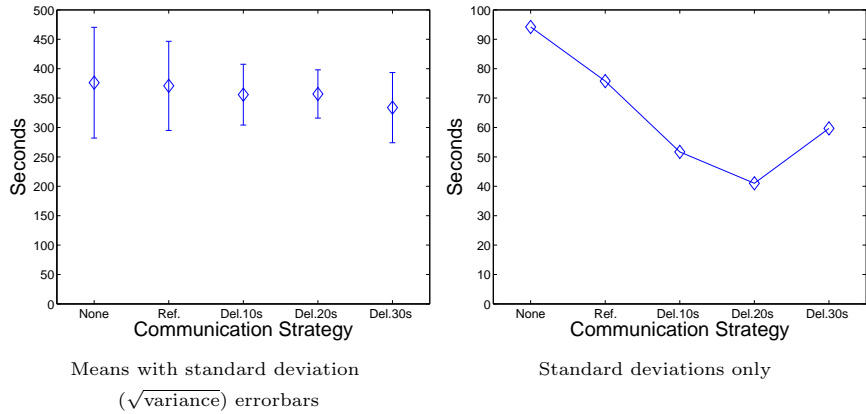


Fig. 4. Means and standard deviations of the times to complete the task for each of the communication strategies. The labels on the x axis stand for the different communication experiments. None=none, Ref.=reflexive, Del.10=10 s deliberative, Del.20=20 s deliberative, Del.30=30 s deliberative.




Fig. 5. Means and standard deviations of the times the robots took to retrieve a new target after dropping one off (i.e. target search time) for each of the communication strategies. The labels on the x axis stand for the different communication experiments. None=none, Ref.=reflexive, Del.10=10 s deliberative, Del.20=20 s deliberative, Del.30=30 s deliberative.

The left graphs of Figure 4 and of Figure 5 plot the means and standard deviations of task completion and of search times, respectively, for the communication experiments. The graphs reflect a slight performance benefit from the use of all forms of communication, but, surprisingly, nothing statistically significant. However, the variance of both show an obvious trend, that can be seen more clearly from the right-hand graphs of Figures 4 and 5. Although f tests show no statistical significance of the difference between the variances at the 95% confidence interval, the variance of the 20 second communication trials were very close to being significant (one-tailed, two-sampled f test with $p=0.0682$ and $p=0.0511$, for time to completion and target search times, respectively).

The beacon-on times recorded in the reflexive communication experiments (shown earlier in Figure 3) suggest that the correlation between the beacon-on times and the variance in completing the task cannot be explained simply by the duration of the beacon-on times.

Robot-to-robot interference and the specifics of how the robots operate are other important factors. The CMUCam turrets can rotate 360° in 5 seconds, but it may take several rotations to detect a beacon. The probability of detection decreases with distance and becomes zero at 2.9 m. A robot can rotate 180° in 5 seconds and can translate at a maximum of 0.17 m/s.

Using these ranges and approximating the probabilities for the time to find a beacon, to rotate, and to home in, we calculated the mean interference time and the mean travel time (i.e. the average time a homing robot traveled toward a communicating robot once it was oriented) for the deliberative experiments:

	10 seconds	20 seconds	30 seconds
Mean Interference Time	0.9971	5.8964	13.8634
Mean Travel Time	2.8500	11.7006	21.6406

Short communication durations (times less than 5 s) do not give enough time to the robots to see the beacon and change their heading to travel towards it. Longer communication times increase the probability of robot interference (mean interference time doubles from 20 s to 30 s). Travel time also increases with long communication.

4 Discussion

To explain the results, we have analyzed some of the failure points during communication. The probability of successful communication (i.e. the ability to see the beacon) is inversely proportional to distance. Even when two robots are in close proximity, successful communication depends upon their relative headings. If the homing robot is facing away from the communicating robot it may not be able to orient itself before the beacon is turned off. If a robot does successfully home in on a communicating robot, the target may be occluded. If the two robots make contact, the homing robot may turn away from the target as it executes obstacle avoidance. Finally, a common source of noise is inter-robot interferences. This becomes particularly troublesome when robots are drawn to the same area by some attractor, such as a beacon.

We could claim that these points of failure for communication are implementation details that can be addressed with more sophisticated hardware or better engineering, but discounting implementation details raises an important issue. These implementation details are precisely why we think real robots are necessary for this type of analysis. It is too easy to discount or underestimate the effects of even simple implementations on real hardware. For example, in [1] communication was shown to improve performance, but nearly all of the experiments were done in simulation where the effects of specific actions on the performance of the system (such as cooperative carrying or consuming of a resource) can be abstracted away. The details involved in physically implementing a system which can carry heavy objects or can consume liquid from a spill may affect the performance of the team in ways that those results did not illustrate. Considerable engineering effort may be necessary before the robots would be able to effectively achieve their tasks at the rates reported in this work.

As a point of comparison, consider a MinDART robot that executes a collection of behaviors to align itself to a target when in the GRABTARGET state. The time that it takes a robot to pick up a target is heavily dependent on the interaction between the robot and its environment. To better quantify this, the times the beacons were turned on in the reflexive communication experiments are the same as the times the robots spent in the GRABTARGET

state. These times (see Figure 3) were quite variable. This illustrates the complexity that can arise from a simple operation implemented on real robots.

We believe our findings are validated by work done by others in simulation, particularly by Balch and Arkin [1] mentioned above. They concluded that simple communication often provides the best performing robots, but sometimes no communication performs just as well. We believe that once you carry robotics into the real world, some improvements in performance found in simulation get reduced by the noise and errors of implementation.

5 Related Work

Most research with multiple robots has focused on various forms of collaborative work [3]. While collaboration may be essential for some tasks, we are interested in studying tasks that can be done by a single robot, but where using multiple robots can potentially increase performance by decreasing the time to complete the task and/or by increasing the reliability. Sample tasks include placing a sensor network [16], cleaning up trash [12], pushing boxes [9], or detecting odors [5].

Foraging is a widely used testbed for distributed systems, but there are differences in the way the task is defined. In most studies the goal is to collect a fixed number of objects (roughly half) [4], in other cases objects continue to appear probabilistically and the duration of each experiment is fixed [8]. In our experiments the task is completed when all the objects have been collected, which makes the task more difficult since it is harder for the robots to find targets when they are very sparse. In our previous work we studied the effect of the number of robots [14] and of localization on performance [15]. In addition to these experimental studies, predictive models of foraging behaviors [11] and of robot interference during foraging [10] have been proposed.

There have been a handful of studies to evaluate the efficacy of communication strategies applied to the foraging task. Our work on communication strategies has been inspired mostly by the theoretical model proposed by Sugawara [18, 19] and by the simulation work of Balch and Arkin [1]. Sugawara's model accounts for the effects of indirect communication in foraging tasks. He performed simulation studies and some limited experiments with physical robots to support his model [19]. An interesting aspect of the model is that it predicts that the duration of the communication affects performance, and that there is a critical duration at which the performance is maximized, below and beyond which team performance deteriorates. Our communication experiments were designed to test this specific aspect of the model.

The study by Balch and Arkin [1], which evaluates the effects of various communication strategies on three different tasks, including foraging, was mostly conducted in simulation. The study predicts that communication improves performance by reducing the time spent wandering around. Our communication experiments were designed to verify this improvement

and to quantify it. It is reasonable to assume that communication will assist in foraging, since it is a strategy that has evolved in nature. It is widely known that bees “dance” to communicate the direction of pollen sources [17] and ants communicate the location of prey with pheromone trails [7]. To our knowledge, biologically-inspired communication strategies for foraging on small scale robots have yet to provide performance improvements as predicted by the above mentioned work.

6 Conclusions and Future Work

We studied the effects of communication on a robotic team doing foraging. We compared this against a baseline of a random-walk search strategy.

We hypothesized that communication would decrease the time the robots spent randomly searching their environment and would improve overall performance, but we did not find a statistically significant improvement compared to the baseline. Instead, what we found was a decrease in the variance of the task completion times. We attribute the decrease in variance to the reduction of random search for targets. With communication capabilities, robots have to randomly wander into the communication range of another robot, but are then drawn directly to targets when attracted by a communicating beacon. Analysis of the average homing distances and interference times supports our conclusion that a 20 s communication duration represents a minimal point of variance for our experimental setup. Durations beyond this increase the probability of robot interference which negatively impacts performance.

For future work, we will explore how robots might dynamically adapt to their environment and tune their communication durations to optimize the team’s overall performance. This learning capability would require upgrades in the processing and communication systems of the robots. Such upgrades would facilitate a robot’s ability to share more information such as intentions, therefore teams could collaborate at a higher level.

7 Acknowledgments

Work supported in part by the Doctoral Dissertation, the Louise T. Doshall Fellowships of the University of Minnesota, and by NSF under grants EIA-0224363 and EIA-0324864. We would like to thank Heather Metcalf and Devon Skyllingstad, who were supported by the CRA Distributed Mentor Program, for their help in constructing the electronics for the robot’s beacons.

References

1. T. Balch and R. C. Arkin. Communication in reactive multiagent robotic systems. *Autonomous Robots*, 1(1):27–52, 1994.

2. E. Bonabeau, M. Dorigo, and G. Theraulaz. *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, Oxford, England, 1999.
3. G. Dudek, M. Jenkin, and E. Milios. A taxonomy for multi-agent robotics. In T. Balch and L. E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*, chapter 1. A K Peters Ltd, Natick, MA, 2002.
4. D. Goldberg and M. J. Matarić. Design and evaluation of robust behavior-based controllers. In T. Balch and L. E. Parker, editors, *Robot Teams: From Diversity to Polymorphism*. A K Peters Ltd, Natick, MA, Apr. 2002.
5. A. T. Hayes, A. Martinoli, and R. M. Goodman. Distributed odor source localization. *IEEE Sensors*, 2(3):260–271, 2002.
6. O. Holland and C. Melhuish. Stigmergy, self-organisation, and sorting in collective robotics. *Artificial Life*, 5:173–202, 2000.
7. B. Hölldobler and E. Wilson. The multiple recruitment systems of the african weaver ant *ocophylla longinoda* (latreille). *Behavioral Ecology and Sociobiology*, 3:19–60, 1978.
8. C. Jones and M. J. Matarić. Adaptive division of labor in large-scale minimalist multi-robot systems. In *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, pages 1969–1974, Las Vegas, Nevada, Oct. 2003.
9. R. C. Kube and E. Bonabeau. Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 30(1/2):85–101, 2000.
10. K. Lerman and A. Galstyan. Mathematical model of foraging in a group of robots: effect of interference. *Autonomous Robots*, 13:127–141, 2002.
11. A. Martinoli, A. Ijspeert, and F. Mondada. Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems*, 29:51–63, 1999.
12. L. E. Parker. On the design of behavior-based multi-robot teams. *Journal of Advanced Robotics*, 10(6):547–578, 1996.
13. A. Rowe, C. Rosenberg, and I. Nourbakhsh. A low cost embedded color vision system. In *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, Lausanne, Switzerland, 2002.
14. P. E. Rybski, A. Larson, M. Lindahl, and M. Gini. Performance evaluation of multiple robots in a search and retrieval task. In *Workshop on Artificial Intelligence and Manufacturing*, pages 153–160. AAAI Press, Aug. 1998.
15. P. E. Rybski, A. Larson, A. Schoolcraft, S. Osentoski, and M. Gini. Evaluation of control strategies for multi-robot search and retrieval. In *Proc. Int'l Conf. on Intelligent Autonomous Systems*, pages 281–288, Marina Del Rey, CA, March 2002.
16. P. E. Rybski, S. A. Stoeter, M. Gini, D. F. Hougen, and N. Papanikolopoulos. Performance of a distributed robotic system using shared communications channels. *IEEE Trans. on Robotics and Automation*, 22(5):713–727, Oct. 2002.
17. T. Seeley. The honeybee colony as a superorganism. *American Scientist*, 77:546–553, 1989.
18. K. Sugawara and M. Sano. Cooperative acceleration of task performance: Foraging behavior of interacting multi-robots system. *Physica*, D100:343–354, 1997.
19. K. Sugawara and T. Watanabe. Swarming robots – foraging behavior of simple multi-robot systems. In *Proc. IEEE/RSJ Int'l Conf. on Intelligent Robots and Systems*, Lausanne, Switzerland, 2002.