Two Foraging Algorithms for Robot Swarms Using Only Local Communication

Citation

Hoff, Nicholas R. III., Amelia Sagoff, Robert J. Wood, and Radhika Nagpal. 2010. Two foraging algorithms for robot swarms using only local communication. In IEEE-ROBIO 2010 : 2010 IEEE International Conference on Robotics and Biomimetics, 123-130. Piscataway, N.J. : Institute of Electrical and Electronics Engineers. December 14-18, 2010, Tianjin, China.

Published Version

doi:10.1109/ROBIO.2010.5723314

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:5345875>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Open Access Policy Articles, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#OAP>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Two Foraging Algorithms for Robot Swarms Using Only Local Communication

Nicholas R. Hoff III, Amelia Sagoff, Robert J. Wood, Radhika Nagpal

Abstract—Large collections of robots have the potential to perform tasks collectively using distributed control algorithms. These algorithms require communication between robots to allow the robots to coordinate their behavior and act as a collective. In this paper we describe two algorithms which allow coordination between robots, but do not require physical environment marks such as pheromones. Instead, these algorithms rely on simple, local, low-bandwidth, direct communication between robots. We describe the algorithms and measure their performance in worlds with and without obstacles.

I. INTRODUCTION

As the cost of robotic hardware has come down and availability has gone up, there has been growing interest in robotic systems which are composed of multiple simple robots rather than one highly-capable robot. This tradeoff reduces the design and hardware complexity of the robots and removes single point failures, but adds complexity in algorithm design. The challenge is to program a swarm of simple robots, with minimal communication and individual capability, to perform a useful task as a collective.

Nature has solved this problem many times over. Schools of fish swim in unison, and are able to execute large scale collective maneuvers to avoid a predator. Termite colonies build large and very complex nests (complete with thermal regulation). Ants collectively search a very large area and are capable of returning food to the nest. In each of these examples, there was no central leader with all the information making decisions for each individual. Leaderlessness is a central aspect of distributed swarm algorithms.

In this work, we take inspiration from natural swarm behaviors to build algorithms for robot swarms. Ant colony foraging is our primary example. Ants use pheromones to mark trails in the environment, which allows them to efficiently communicate the location of food sources and collectively return food to the nest [1]. This work focuses on a way for simple robots with simple sensing capabilities to achieve the same task, without having to implement physical pheromones or physical environment marking.

In our application, we assume there is a swarm of robots that have simple sensing capability: each robot can communicate directionally with other robots within a short range. The task is for the group of robots to search the environment for an object of interest (“food”) and then return the food to the base (“nest”). The robots do not know the location of the food a priori, nor do they have GPS/odometry capabilities,

but the efficiency of the group can be dramatically improved through coordination. We present two algorithms, inspired by ant colony behavior, that use a simple concept: each robot can dynamically take on one of two roles, that of a stationary environment beacon (like a pheromone) or a wandering robot. The two algorithms (termed “virtual pheromone” and “cardinality”) differ in when the beacon role is chosen and what information the beacon emits.

We show that both algorithms are able to effectively construct paths between food and nest, and to adapt to environments with static obstacles. We compare the performance of the algorithms to the performance of robots without coordination and robots with explicit global position and knowledge of the food; these comparisons represent two extremes in hardware complexity of the robots. We show that the virtual pheromone and cardinality algorithms dramatically increase efficiency over no coordination, and performance scales directly as the number of robots increases until congestion impacts all algorithms. Our work shows that for the task of foraging, we can take advantage of the collective efficiency that ant colonies achieve, while relying on hardware implementations that are feasible in simple and small robots.

In the future, when these algorithms can be implemented on large collections of small robots, they will enable applications such as distributed cleanup of environmental toxins and automated repair. For example, a swarm of small robots could search a building for defects, then recruit other robots to bring tools and materials to the site of the damage.

A. Related Work

Multiple researchers have studied foraging and other pheromone-mediated coordination, and have reduced this behavior to algorithms [2], [3], [4], [5]. Pheromone-trail based algorithms sometimes have the ability to dynamically improve their path [6] and can adapt to changing terrain [7].

Ant-inspired foraging has been implemented in robots by various groups. One of the chief difficulties is in implementing the pheromone itself, or some way for the robots to interact. There have been many approaches to this problem:

- Physical marks. Robots can physically mark their trails in a variety of ways, such as leaving alcohol [3], heat [8], odor [9], visual marks [10], or RFID tags [11].
- Use existing communication channels. In the work of Vaughan et al., robots maintain an internal pheromone model with trails of waypoints, and share it with other robots over a wireless network [12], [13], [14].

- Virtual pheromones. In [15], authors use direct infrared-based communication between robots to transmit a kind of virtual pheromone. They study the use of these signals to create world-embedded computation and world-embedded displays. It is assumed that the robots that receive the pheromone can measure the intensity of the IR reception to estimate their distance from the transmitter.
- Pre-deployed sensor network. The GNATS project [16] uses robots which operate in an environment which is pre-populated with a regular grid of beacons on which information can be stored. They have had success in finding close-to-optimal paths from one place in the environment to another, even in the presence of obstacles.
- Deployable beacons. Some groups have explored the idea of having each robot be able to deploy beacons as it moves through the environment. The beacons can be movable or non-movable [17], and contain pheromone-like information.
- Robot chains. In [18], [19], the robots form chains and attempt to remain in close proximity with each other, through communication or even by physically gripping the next robot in the chain.

There are several shortcomings of these approaches which we aim to address. Making permanent physical marks in the environment is generally not acceptable, and temporary or decayable marks are difficult to physically implement. Relying on predeployed sensor networks is highly restrictive and prevents operation in a new or unexplored environment. Deploying beacons is a good method but requires building robots which can carry many beacons and can also intelligently recover previously laid beacons. Instead, we use the robots themselves as beacons. (see also [20])

Of the work cited above, our work is most similar to the virtual pheromone approach, except that our robots do not have different behaviors depending on a direct measurement of their distance from each other. This simplifies the communication hardware because distance measurement is not required. Secondly, we show that multiple algorithms can be used with this communication model to achieve coordination and approximate the collective benefit of pheromones.

B. Robot Model and Simulator

Robot Model: We assume a simple non-holonomic robot that moves and turns in continuous space. Each robot has sensors for nest, food, and obstacle detection in direct proximity to the robot. Each robot can also communicate with nearby robots using a simple IR ring architecture. The robots have omnidirectional transmission, and directional reception. This means that when a robot receives a transmission, it knows roughly which direction the transmission came from (see Fig. 1). An example of such communication hardware is described in [21].

Simulator: In order to test the algorithms, we developed a simulator based on Microsoft(R) Robotics Studio (MRDS). The simulator models a continuous world in which the

Fig. 1. This figure shows the communication and sensing structure of the robots. The figure in the upper left shows an example world setup. The blowup details a hypothetical communication situation. When the robot on the left transmits, the robot in the center can receive the transmission, and also knows that the transmission came from the left-facing octant. The light gray range in front of the central robot indicates the approximate area in which the nest, food, and obstacle sensors are sensitive.

robots, food, nest, and obstacles exist, and each object occupies some physical extent. Realistic physics and visualization is provided by MRDS. A snapshot from the simulator is shown in Fig. 2. We chose this simulation environment over a gridded world environment so that the algorithms face the real-world problems such as collisions and congestion. Our simulator closely matches a model of the Kheprea II robots which we plan to use in the future for hardware verification of the algorithm.

In the remaining sections of this paper, we will describe two ant-inspired algorithms capable of foraging without physical pheromones, then present performance results measured in simulated worlds with and without obstacles, and finally conclude with a discussion of future algorithm work and application to small autonomous robots.

II. ALGORITHM DESCRIPTION

In this section, we will describe two algorithms, called the virtual pheromone (VP) algorithm and the cardinality algorithm.

Ant colony foraging: Both algorithms are based on the foraging behavior of ants. Ants mark trails leading from the nest to food and back by depositing a chemical pheromone on the ground. Each ant can both deposit and detect this chemical, and each ant uses the distribution of pheromone in its immediate vicinity to decide where to move. When an ant is carrying food, it lays pheromone as it searches for the nest, and other ants follow this trail outbound to the food. This distributed leaderless pheromone algorithm is the starting point for the VP algorithm, but several important changes have been made.

Virtual Pheromone Algorithm: The first change has to do with the way the robots return to the nest. The pheromone trail is laid by individuals once they leave the food, but the individual still needs a way to return to the nest. Ants have various methods to do this, including visual landmarks and odometry (remembering how far and in what direction

Fig. 2. This shows a screenshot of the simulator. Window a shows a dot for each robot, with beacon robots shown as larger dots and beacon values drawn to the right of beacons. The square near the top is the nest and the square near the bottom is the food. Window b shows the actual MRDS physics simulation and visualization.

the nest is). Because of our focus on miniaturizability and hardware simplicity, neither of these approaches is attractive – vision requires significant computation and memory to interpret the image, and odometry requires high accuracy encoders. As an alternative method to navigate back to the nest, two distinguishable pheromones are used instead of one. One pheromone leads to the food and a second leads to the nest.

A second change from the biological algorithm is that *virtual pheromones* are substituted for real ones (see Fig. 3a). Instead of implementing chemical pheromones, simple local direct communication between robots is used to transmit a virtual pheromone value.

During the execution of the algorithm (see pseudocode in Fig. 4), some robots will decide to stop their normal search behavior and become ‘pheromone robots’, also known as beacons, which means they stop moving and act as locations on which virtual pheromone can be stored. The virtual pheromone is simply a floating point number. Other wandering robots (‘walkers’) can read the pheromone level by receiving a transmission from the pheromone robot, and they can lay virtual pheromone by transmitting to the pheromone robot. So, if there were a network of pheromone robots, the walker robots could use the distribution of virtual pheromone they can sense to decide how to move. This mimics the use of pheromones in real ants. Virtual pheromones decay at a specific rate and robots follow the pheromone distribution by moving toward the strongest value. Several of the parameters of the algorithm, such as the walker-beacon transition probabilities and exploration timer, are shown in Fig. 4. A diagram of an example situation of the VP algorithm is shown in Fig. 3a.

Cardinality Algorithm: The *cardinality* algorithm is similar to the *VP* algorithm in that robots can decide to act as either beacons or walkers — beacons transmit values, and walkers use those values to decide where to move. The main difference is that instead of the values that the beacons store

Fig. 3. This diagram shows example situations in both the *VP* and *cardinality* algorithms. Each black dot represents a beacon (walkers are not shown). The numbers to the right of each beacon indicate the virtual pheromone levels (fig a) or cardinalities (fig b) – food pheromone/cardinality on the top and nest pheromone/cardinality on the bottom. In fig a, the dotted lines indicate approximately level curves in both pheromone values. As can be seen, the area between the nest and food has high concentration of pheromone. In fig b, the solid line indicates the shortest path between the nest and the food. As an example, the dotted lines indicate paths that a robot near the $\frac{5}{4}$ beacon on the left could take to get either to the nest or to the food.

and transmit being real-valued floating-point numbers, they are integers (called “cardinalities”). The first beacon standing next to the nest transmits a 1, then the next beacon a little further out would be able to hear the 1 and so would transmit a 2. In general, each beacon transmits the minimum of all the other beacons it can hear, plus one. In this way, the cardinality of each beacon can be interpreted as the number of beacons between that beacon and the nest. Furthermore, a walker robot can use these cardinalities to find a path to the nest by always moving to the lowest cardinality it can hear. In a similar manner to the *VP* algorithm, each beacon actually transmits two cardinalities — one indicating how many beacons away from the nest it is, and the other indicating the number of beacons away from the food. Pseudocode for

first time step: become walker	FOOD SEARCH: try to pick up food <u>am I doing AVOID?</u> <u>sense obstacle?</u> turn τ degrees END else move forward lay <i>inPheromone</i> (VP only) END <u>am I doing EXPLORE?</u> <u>stepCounter > n?</u> stepCounter = 0 end behavior END else try to pick up food <u>sense obstacle?</u> turn τ degrees END else move forward lay <i>inPheromone</i> (VP only) stepCounter ++ END
begin each timestep: if beacon, goto beacon if walker, goto walker	
walker (either alg): <u>hear 2 or more beacons?</u> <u>carrying food?</u> goto NEST SEARCH else goto FOOD SEARCH else become beacon END	<u>neither AVOID nor EXPLORE:</u> <u>should I explore?</u> set stepCounter = 0 execute EXPLORE else turn to strongest <i>outPheromone</i> or lowest <i>foodCardinality</i> <u>sense obstacle?</u> execute AVOID else move forward lay <i>inPheromone</i> (VP only) END
beacon (VP alg): <u>hear 3 or more beacons?</u> with prob p , become walker END else <u>both pheromones < threshold?</u> become walker END else decay pheromones (rate d) END	
beacon (cardinality alg): <u>hear 3 or more beacons?</u> with prob p , become walker END else set my cardinality to min of all recieved cardinalities + 1	

Fig. 4. This figure shows combined pseudocode for both the VP and cardinality algorithms. Robots begin execution at “start” at the top and then proceed according to which algorithm is being run and the role (walker or beacon) of the robot. The normal food search part of the algorithm is shown; the nest search part is nearly identical, the biggest difference being that robots lay *outPheromone* instead of *inPheromone*. Some of the parameters of the algorithm are shown in this pseudocode. In the experiments reported, probability p is set to 30%, angle τ is 45° , and step counter n is 4. The pheromone decay rate d is 1%.

the cardinality algorithm is shown in Fig. 4.

A snapshot of the cardinality algorithm is shown in Fig. 3b. The beacon standing directly next to the nest is transmitting a 1 for its nest cardinality, and the beacon standing next to the food is transmitting a 1 for its food cardinality. All other beacons are updating their cardinalities accordingly. Now, by listening to the cardinalities and always moving to the smallest one, the walker robots can walk to the nest or to the food from wherever they are.

One significant difference to notice between the VP and cardinality algorithms is that VP requires the walkers to transmit to the beacons every time they want to lay a virtual pheromone. The cardinality algorithm, however, only requires beacons to transmit – the walkers only need to receive. Additionally, cardinality has fewer parameters and is simpler.

III. TEST PARAMETERS AND METRICS

In this section, we will discuss the performance of each algorithm in an obstacle-free world.

A. World and Test Parameters

The world setup, including positions of nest and food, is shown in Fig. 8a. (See section V for results in worlds B, C, and D.) There are several test-related parameters which must be chosen, such as world size, food placement, run time, and communication radius.

- **world size** — World size can be non-dimensionalized by dividing by robot body length. The robots we will ultimately use to test this algorithm have a body length of about 8 cm. A square world in which each side is 50 times the body length would be $4m \times 4m$. Real ants have a body length around 8 mm, but operate in roughly the same size areas, so they would have a world size to body length ratio of around 500. We study world with and without obstacles. In this study, we use square worlds for their symmetry and simplicity. In the future, we may experiment with non-square worlds, or worlds with no boundaries at all. (The presence of boundaries makes the problem simpler by focusing the robots on the area of interest; in a world with no boundaries, the robots could get permanently lost.)
- **robot density** — The number of robots in the swarm divided by the total world size gives the average robot density. In this paper, we will primarily explore the effect of this parameter, presenting results as a function of both absolute number of robots and robot density.
- **nest-food separation** — This can be non-dimensionalized by dividing by world size. In this paper, we separate the nest and the food by a distance equal to 60% of the world size. This gives a large enough separation that the problem is hard, but not so large that both the food and nest are next to the world boundaries. In these experiments, there is only one food pile and it never runs out of food.
- **run time** — To non-dimensionalize the run time, we will use the number of direct nest-food traversals that are theoretically possible during the run (assuming the robots could travel straight to the food and back). We will run each simulation long enough for a robot to traverse the nest-food distance 25 times. Assuming the robots can move approximately one body length per time step, that would yield a run time of 750 time steps.
- **communication radius** — The non-dimensional parameter here is communication radius divided by body length. With the hardware we intend to use to implement the local communication, a ratio of 10 seems conservative. So in this work, we assume a communication radius equal to 10 times the body length, or 80 cm.

In this paper, we will focus on the effect of world size and robot density. A study of the sensitivity of the results to variations in the other parameters is beyond the scope of this paper and will be studied in future work.

B. Metrics

To measure the performance of the VP and cardinality algorithms, two metrics will be used:

- **total food returned** — The total amount of food that has been returned to the nest by the entire swarm after 750 time steps. Note that in all experiments, food piles have an infinite amount of food, so this metric is very similar to the rate of food return.
- **cost** — The total energy consumed by the swarm. For the purposes of calculating this cost metric, every communication action will incur a cost of 1 and every movement action will incur a cost of 100. In this way, cost is roughly analogous to battery usage. The total cost is the sum of all the costs of all the robots for the duration of the test.

The performance of the algorithms as measured by these metrics, will be compared against two comparison algorithms. The first is a `randomWalk` algorithm, in which each robot decides randomly whether to turn a random amount or to move forward. The second comparison algorithm is called the `foodGPS` algorithm. In this algorithm, each robot has perfect knowledge of the locations of the food and the nest, so it is capable of turning directly to the food source and moving in that direction, then turning directly to the nest and returning. The `foodGPS` algorithm will not always yield perfect or optimal performance because the robots may still be blocked by obstacles or other robots (of which they have no knowledge). The robots do not take optimal paths around obstacles, they simply attempt to go straight to the final goal location, avoiding obstacles as they encounter them.

`randomWalk` was chosen as a lower comparison algorithm because it has the smallest requirements for communication and sensor hardware. It represents a minimum use of coordination among robots — none. `foodGPS`, on the other hand, has very substantial requirements. Every robot must know the position of the food and the nest, and must know its own position at all times. These two algorithms represent extremes of hardware requirements. Our aim is to develop algorithms with significantly increased performance with only moderately increased hardware requirements.

IV. RESULTS IN AN UNCLUTTERED WORLD

Here we show results of the `VP` and `cardinality` algorithms over the course of a single run and compared across multiple runs, operating in obstacle-free worlds. A key aim is to understand how reliable and efficient the two algorithms are.

A. Performance of Algorithms over Time

At each time step during each run, we measured the total food that had been returned to the nest, the fraction of non-beacon robots that were carrying food, and the total number of beacons. Data from example runs of the `VP` and `cardinality` algorithms is shown in Fig. 5. These examples were run in a $4m \times 4m$ world in configuration A (see Fig. 8) with 70 robots. In the first 50-100 time steps, the number of beacons rapidly increases to a steady value as the ad-hoc network of beacons is deployed. Next, the fraction of food-carrying ants becomes non-zero, as the first robots begin finding and picking up the food. Some time after that (about

Fig. 5. This shows an example run of each algorithm in the world configuration of Fig. 8A. Each run was performed with 70 robots.

100-200 time steps in these examples), the first robots are able to return food to the nest. For the rest of the run, the amount of food returned increases and the number of walker robots carrying food approaches roughly 50%. It makes sense that half of the walkers would be carrying food in steady state, because approximately half would be on their way to the food and half would be returning.

The major difference between the two algorithms is that `cardinality` returns food faster than `VP`, and returns more. `VP` requires time for the trail between the food and the nest to be built and reinforced after it is found, whereas no such time is required in `cardinality`. In `cardinality`, once any robot has found the food, all robots have a path to get there, and then once they pick up food, they have a path to get back to the nest. No reinforcement is required.

B. Performance and Comparison of Algorithms

Each of these four algorithms (`randomWalk`, `VP`, `cardinality` and `foodGPS`) were run for 750 simulation time steps, and the total amount of food collected at the end of the run was recorded. The results are shown in Fig. 6a.

`randomWalk` never returns any food. In worlds this big, the probability of a robot randomly finding the food and then randomly finding the nest is so small that it never happened in any of these tests. With this set of test parameters, `foodGPS` returns a large amount of food. `cardinality` and `VP`

A brief 2D random walk analysis, omitted for brevity, shows that the number of steps required for the first robot to find the food is $N \approx \frac{f}{R} \left(\frac{s}{d}\right)^2$, where R is the number of robots, s is the nest-food separation, d is the step size, and f depends on the distance at which robots can pick up food. With appropriate numerical values for these parameters and $R = 120$, N is 2460. Because these tests were only run for 750 time steps, it is not surprising that no food was returned.

are in the middle, and `cardinality` outperforms `VP`. In these tests, the food pile can never run out of food, so there is food available for each algorithm for the whole duration of each run. The algorithms can never be ‘finished’, so `foodGPS` is used as a measure of optimal performance. The effect of congestion can also be seen in the data. As the number of robots increases, the performance also increases because there are more robots available to transport food. However, when there are more than about 100 robots, it is so crowded that they have trouble moving and the performance again decreases. See section IV-C for a more detailed discussion of congestion.

The fact that `VP` and `cardinality` perform better than `randomWalk` at all shows that some coordination is being achieved between the robots using the virtual pheromones and cardinalities. These algorithms sacrifice some robots to be immobile beacons, and these robots are no longer picking up and dropping of food. It could be the case that this sacrifice outweighs the benefit derived from the virtual pheromone or cardinality. If that were so, using the `VP` or `cardinality` algorithms at all would be harmful and `randomWalk` would be better. In fact, they outperform `randomWalk`, which shows that the sacrifice of some robots to be beacons confers a net benefit on the algorithm. Furthermore, although each algorithm uses robots as beacons to achieve a coordination benefit, `cardinality` derives more benefit from these beacons. This benefit is reduced and eventually eliminated in both cases at small numbers of robots, however, because in that case the robots are mostly used as beacons so there are few walkers left.

The second metric is cost. Cost is similar to battery usage in the real robots, but with arbitrary units. When using this metric, we divide it by the total food returned by the swarm to arrive at a cost-per-food metric. The results are shown in Fig. 6b. For each number of robots, `cardinality` is capable of returning food at a lower cost than `VP`.

There also appears to be a trend in each algorithm in which smaller swarms are more efficient than larger swarms. This is likely related to congestion. Usually, when a robot encounters another robot in its path, it must do an avoidance maneuver, which requires a sequence of turns and forward movements. If it had not encountered the obstructing robot, it would have been able to simply move forward along its desired path with a single movement action. Because movement is expensive, and avoidance maneuvers require more movement than simple path following, larger swarms are less efficient because of the increased frequency of avoidance.

Both `cardinality` and `VP` exploit the ant-inspired idea of marking the environment, but the success of the `cardinality` algorithm shows that there are more cost effective ways to lay marks than to simply copy the ants’ method. The algorithms are far from the `foodGPS` performance, however, so the ability to have long-range detection of the nest could have a substantial effect on performance.

C. Congestion

Crowds of robots around the nest and the food can impair swarm performance. Increasing the number of robots

Fig. 6. These plots show the performance of `foodGPS`, `cardinality`, and `VP` as measured by both metrics. The horizontal axes are marked both in absolute number of robots and in robot density. `randomWalk` never returns food, so it is not shown. Each point is an average of approximately 15 runs, with standard deviation error bars. Note the log scale on the vertical axes.

improves performance by providing more robots to transport food, but exacerbates the problem of congestion. To measure this effect, we compare the performance of `foodGPS` under normal conditions, and its theoretical maximum performance if robots were able to occupy the same physical space. Since all robots know the exact location of the food and nest, but not other robots, they attempt to move in a straight-line path toward the food or nest, and must avoid collisions with each other as robots appear.

In the absence of physical collisions, the maximum rate at which a swarm of robots could return food to the nest would be $\frac{fns}{2d}$, where f is the amount of food each robot can carry, n is the number of robots in the swarm, s is the distance each robot can travel in one time step, and d is the distance between the food and the nest. (d is doubled because the robots must travel to the food and back.) This rate is measured in food units per time step.

Fig. 7 shows the performance of `foodGPS` vs the the-

Fig. 7. Effect of congestion. When the number of robots is small it is not a problem, when the number of robots is large, it is the dominant effect, and there is an optimum in the middle. Each `foodGPS` point on this plot is an average of approximately 15 trial runs, with standard deviation error bars.

oretical maximum. At low robot densities, congestion has little effect, but even with a density around 4 robots/m², congestion strongly degrades the performance. Congestion is the dominant factor at high robot density.

V. TESTS AND RESULTS IN WORLDS WITH OBSTACLES

Our long term goal is to operate large numbers of autonomous robots in natural obstacle-filled environments. In this study, we used three different obstacle configurations, in addition to an obstacle-free world. All configurations are shown in Fig. 8. In configuration B, a simple obstacle blocks the most direct path, such that the swarm must find a path around the obstacle. Configuration D has two obstacles that require the robots to take a more complex curving path to get to the food and back. Configuration C gives the swarm a choice between two paths: a short path that goes through an obstacle field, or a long but clear path.

VP, cardinality, and `foodGPS` do obstacle avoidance in nearly the same manner. Each walker has a direction it wants to go, determined either by virtual pheromones, cardinalities, or just knowing the correct direction (in the case of `foodGPS`). For each of these three algorithms, when a robot encounters an obstacle, it attempts to avoid it, usually by simply turning left and moving forward.

The performance of `foodGPS` can be used as a measure of the difficulty of each obstacle field. The field with no obstacles is the easiest, then the single obstacle world, then the obstacle field, and the double obstacle world is the hardest. Taking the `foodGPS` result as a type of performance standard for each obstacle field, we will normalize the performance to the performance of `foodGPS` on that obstacle field. Results of these normalized performances are shown in Fig. 9. These algorithms do quite well compared to `foodGPS` considering their hardware requirements. `foodGPS` requires all robots to have global knowledge, and cardinality can achieve a significant fraction of that performance (about 5% - 20%) with no global knowledge or

Fig. 8. There are four world configurations used in this paper. Results from configuration A are discussed in section IV, and results from configurations B, C, and D are discussed in section V. The obstacle field of configuration D is an area with many small obstacles obstructing simple straight line paths.

Fig. 9. This figure shows the amount of food returned by each algorithm in each obstacle field, normalized to the amount of food returned by `foodGPS` on that obstacle field. For example, in the single obstacle environment, cardinality returns about 12% of the food that `foodGPS` does.

communication at all and only simple local communication. The performance of cardinality follows the expected trend, decreasing with increasing obstacle configuration difficulty. VP, however, performs better in the obstacle field world than in the single obstacle world (although in no case better than cardinality). This is likely because the pheromone distribution will have some randomness so a few robots will penetrate the obstacle field, but in the single obstacle case, all robots must go around. cardinality is better at the single obstacle case because the path can be composed of a small number of straight-line segments which cardinality prefers. This does not mean that VP is better than cardinality at the obstacle field, though – cardinality always outperforms VP.

We would like to measure the ability of the VP and cardinality algorithms to adapt to changing obstacles, and

specifically their ability to find a faster path if an obstacle is removed. To test this, we used the single-obstacle world setup, shown in Fig. 8B and ran the algorithm as usual. Then, 2/3 of the way through the test, we removed the obstacle. This makes a new faster path theoretically available for the robots to use, if they find it and switch to it.

Direct measurements of return rate were taken before and after the obstacle was removed. For `cardinality`, the return rate is in fact faster after the obstacle is removed. Return rates before and after obstacle removal (measured in food units per time step) for each algorithm, along with the percentage increase are shown here:

condition	VP	cardinality
before removal	0.00266	0.0334
after removal	0.00911	0.0536
increase	343%	60%

The speed at which `cardinality` returns food after the obstacle is removed increases by 60%, indicating that it is indeed able to find and switch to a faster path. The specific mechanism of this is simple. When the obstacle is removed, beacons on opposite sides of the obstacle which could not previously communicate are now able to communicate and will therefore update their cardinalities accordingly, connecting the path from the food to the nest. Both before and after obstacle removal, VP performs much worse than `cardinality`. In fact, the pre-removal performance of VP is so poor that any small increase will appear large by comparison, yielding a large increase percentage.

VI. CONCLUSION

In this paper, we presented two ant-inspired algorithms for robot foraging: VP and `cardinality`. Both performed better than `randomWalk`, which indicates that they were able to achieve some coordination benefit from the beacons. They function correctly in worlds with and without obstacles, however congestion has a significant effect.

VP and `cardinality` show that the ant-inspired idea of marking the environment is an effective method of coordination even if one has to sacrifice some robots to act as beacons to hold the values. However, `cardinality` shows that there are better ways of marking the environment than just mimicking the ant-pheromone strategy, as VP does. Lastly, even with the coordination, the performance is much lower than ideal (`foodGPS`) and thus finding ways of increasing the range of detection of food or nest (odor or long-range signal-based navigation to the nest) can have a strong effect on the robot swarm performance and it may be appropriate to pursue a combined strategy.

In the future, we will study ways to reclaim useless beacons for use as walkers, explore more of the parameter space (multiple food piles, varying communication radius), and study the robustness of these algorithms to robot failures.

A parallel line of future development will focus on the hardware required to implement the communication model described in this paper. A preliminary prototype has been designed and built, and testing is under way. The ultimate goal

is to interface the communication structure with the Khepera hardware. This project also fits into a larger microrobotics research effort. These algorithms are designed with a focus on minimal hardware partially because we envision them running on microrobotic insects. Such microrobotic hardware platforms are currently under development at the Harvard Microrobotics Lab [22].

REFERENCES

- [1] B. Hölldobler and E.O. Wilson. *The Ants*. Springer-Verlag, 1990.
- [2] F. Adler and D. Gordon. Information collection and spread by networks of patrolling ants. *The American Naturalist*, 140(3):373–400, 1992.
- [3] R. Sharpe and B. Webb. Simulated and situated models of chemical trail following in ants. In *Proceedings of the International Conference on Simulation of Adaptive Behavior*, pages 195–204, 1998.
- [4] D. Lambrinos, R. Möller, R. Labhart, R. Pfeifer, and R. Wehner. A mobile robot employing insect strategies for navigation. *Robotics and Autonomous Systems*, 30:39–64, 2000.
- [5] M. Nakamura and K. Kurumatani. Formation mechanism of pheromone pattern and control of foraging behavior in an ant colony model. In C. G. Langton and K. Shimohara, editors, *Proceedings of the Fifth International Workshop on the Synthesis and Simulation of Living Systems*, pages 67–76, 1997.
- [6] J. Deneubourg, S. Aron, S. Goss, and J. Pasteels. The self-organising exploratory pattern of the argentine ant. *Journal of Insect Behaviour*, 3(2):159–168, 1990.
- [7] M. Resnick. *Turtles, Termites, and Traffic Jams: Explorations in Massively Parallel Microworlds*. MIT Press, 1994.
- [8] R. Russell. Heat trails as short-lived navigational markers for mobile robots. In *Proceedings of the International Conference on Robotics and Automation*, pages 3534–3539, 1997.
- [9] R. Russell. *Odour Sensing for Mobile Robots*. World Scientific, 1999.
- [10] J. Svennebring and S. Koenig. Building Terrain-Covering Ant Robots. *Autonomous Robots*, 16(3):313–332, 2004.
- [11] Mamei et al. Spreading Pheromones in Everyday Environments via RFID Technologies. *2nd IEEE Symposium on Swarm Intelligence*, 2005.
- [12] R. T. Vaughan, K. Stoy, G. S. Sukhatme, and M. J. Mataric. Blazing a trail: insect-inspired resource transportation by a robot team. *Proceedings of the International Symposium on Distributed Autonomous Robot Systems*, 2000.
- [13] R. T. Vaughan, K. Stoy, G. S. Sukhatme, and M. J. Mataric. Whistling in the dark: Cooperative trail following in uncertain localization space. In C. Sierra, M. Gini, and J. S. Rosenschein, editors, *Proceedings of the Fourth International Conference on Autonomous Agents*, pages 187–194, 2000.
- [14] R. Vaughan, K. Stoy, G. Sukhatme, and M. Mataric. LOST: Localization-space trails for robot teams. *IEEE Transactions on Robotics and Automation*, 18(5):796–812, 2002.
- [15] D. Payton, M. Daily, R. Estkowski, M. Howard, and C. Lee. Pheromone Robotics. *Autonomous Robots*, 11(3):319–324, 2001.
- [16] K. OHara, D. Walker, and T. Balch. The GNATs Low-cost Embedded Networks for Supporting Mobile Robots. pages 277–282, 2005.
- [17] E. Barth. A dynamic programming approach to robotic swarm navigation using relay markers. In *Proceedings of the 2003 American Control Conference*, 6:5264–5269, 2003.
- [18] S. Goss and J. L. Deneubourg. Harvesting by a group of robots. In *First European Conference on Artificial Life*, pages 195–204. MIT Press, 1992.
- [19] B. Weger and M. J. Mataric. Robotic food chains: Externalization of state and program for minimal-agent foraging. P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S. W. Wilson, editors, *Proceedings of the International Conference on Simulation of Adaptive Behavior*, pages 625–634, 1996.
- [20] J. Werfel. Building Blocks for Multi-Agent Construction. *Distributed Autonomous Robotic Systems*, 2004.
- [21] A. Gutierrez et. al. Open E-Puck Range & Bearing Miniaturized Board for Local Communication in Swarm Robotics. *International Conference on Robotics and Automation*, 2009.
- [22] A. Baisch, R. Wood. Design and Fabrication of the Harvard Ambulatory Micro-Robot. *International Symposium on Robotics Research*, 2009.