Territorial Multi-Robot Task Division

Miguel Schneider-Fontán and Maja J Matarić

Abstract—

This work demonstrates the application of the distributed behavior-based approach [1] to generating a multi-robot controller for a group of mobile robots performing a clean-up and collection task. The paper studies a territorial approach to the task in which the robots are assigned individual territories that can be dynamically resized if one of the robots missfunctions, permitting the completion of the task. The described controller is implemented on a group of four IS Robotics R2e mobile robots. Using a collection of experimental robot data, we empirically derive and demonstrate most effective foraging in our domain, and show the decline of performance of the space division strategy with increased group size.

I. Introduction

Designing controllers for multi-robot systems is a complex problem in robotics and artificial intelligence [2], [3]. Our previous work introduced a methodology for synthesizing a *basis behavior substrate* for generating group behaviors such as *wandering, homing, following, aggregation, dispersion*, and methods for combining such behaviors into higher-level composite behaviors including *flocking* and *foraging*, variants of moving in formation and distributing collection [4], [3]. This paper extends our work on homogeneous agent groups executing identical control strategies over the entire environment to a somewhat more complex case of employing a spatial division of labor based on an ethologically common organizational principle of *territoriality*.

As we have argued in [5], [4], adaptive group behavior is a balance between minimizing interference and maximizing synergy (goal achievement at the level of the group), and interference is a key stumbling block in the way of efficient group interactions. It is likely to be correlated with the spatial density of the agents over

Miguel Schneider-Fontán is with the Instituto de Automática Industrial (IAI) Consejo Superior de Investigaciones Científicas (CSIC), La Poveda, Arganda del Rey, Madrid 28500. E-mail:schneider@iai.csic.es

Maja J Matarić is an Asst. Professor of Computer Science and Associate Director of the Robotics Research Lab at the University of Southern California, 941 West 37th Place, MC 0781, Los Angeles, CA 90089-0781. E-mail: mataric@cs.usc.edu

the lifetime of a task, so various approaches to resource division can be applied to counter its effects. Territoriality is as a stable and recurring behavioral pattern that produces a physical division of space and all associated resources [6], [7].

In this paper, we explore the effects of territoriality on a distributed clean-up and collection task, the prototypical group behavior we have been using for studying issues in multi-robot control and learning [8], [9]. We apply a synthetic approach that consists of implementing ethologically-inspired behaviors on a collection of mobile robots, evaluating their performance on repeated trials, and looking for stable, robust, and generally useful behaviors we can derive that help synthesis for various robotics applications.

We focus on a particular type of territoriality based on equal spatial territories assigned *a priori*. In this paper, we address the question of how does the size of the group affect the effectiveness of the territorial solution.

The rest of the paper is organized as follows. Section II overviews related work. Section III describes our experimental setup, the robots, the spatial assignment of the territories, and gives the key formal definitions on which the controller is based. Section IV presents the distributed controller. Section V discusses the sources of interference, the key metric of system performance. Section VI describes the experimental results, including the experimental layout, the data gathering and analysis methods, the system's performance on the various experiments, and a control experiment with fully homogeneous robots to compare time and interference performance. Section VII concludes the paper.

II. Related Work

The dynamics of group interaction in physical robots have only recently begun to be studied. The review in this section focuses on work on foraging-style collection tasks using either physical robots or realistic simulations implemented with the goal of studying issues of critical mass and control of group behavior.

Matarić [5] describes early work on minimizing complexity in controlling a collection of robots, and outlines the first empirical results with the basis behaviors for control [10]. These served as the foundation for the subsequent demonstrated group behaviors, including following, aggregation, dispersion, homing, flocking, and foraging [4], [3]. The work described here used the same behavior-based methodology, constructing the individual agent controllers out of collections of distributed behavior networks [1]. Unlike our previous work, however, which focused on fully homogeneous systems, this paper addresses issues of task division and thus specialization as an approach to interference minimization.

Arkin [11], and Balch and Arkin[12] describe a simulated foraging task with communication-free robots. In a paper more related to the work described here, Arkin, Balch and Nitz [13] present simulation work with the purpose of studying the issues of critical mass in a multi-agent retrieval task. Unlike in our work, the robots are fully homogeneous, and no strict or dynamic territorial division is implemented. However, complementary results regarding critical numbers of agents for the given task are derived. Both of our pursuits are motivated at least in part by the search for interactive, social strategies that demonstrate linear and superlinear performance improvements through collective strategies. Arkin [14] describes the general schema-based control architecture, which is a form of behavior-based control, and gives the critical mass experiments. Finally, Arkin and Ali [15] present a series of simulation results on related spatial tasks such as foraging, grazing, and herding.

Parker [16], [17] demonstrates work related to ours, on the same set of R2 robots with *a priori* hard-wired heterogeneous capabilities. A toxic waste clean-up task used in the work is equivalent to foraging. The Alliance architecture is proposed as an approach to dynamically assigning tasks to members of a robot groups, and is demonstrated on a group of robots dividing the clean-up task. Unlike the spatial division we employed, Parker [17] describes a temporal division that sends one robot to survey and measure the environment, then return and report to the rest of the group, which uses the information to then clean up the spill. The approach is complementary to ours and demonstrates a clean tradeoff between spatial and temporal task division, both of which are ubiquitous in nature. The same work also describes a simulated office garbage-collection task in which the division of labor is performed based on a first-come first-served basis using a simple and elegant conflict-resolution scheme. This work shares a common philosophy with ours regarding dynamic task assignment; the spatial territories we employed can also be dynamically reassigned according to the adapting needs of the group.

Becker, Holland and Dennenbourg [18] describe a fully stigmergic approach to the foraging task in which the group of five robots uses no external sensing or communication to collect the pucks. Instead, through a careful combination of the mechanical design of the robots' puck scoops and the simple controller that moves them forward and in reverse, the robots probabilistically move all of the pucks into a single cluster after a certain period of time. The final location of the cluster cannot be determined *a priori* but the collective behavior is highly repeatable and has minimal sensing and computational overhead. The authors also show critical mass effects by comparing task performance on group sizes ranging from one to five robots. These results are complementary to most other work in the field, in that they demonstrate an extreme near-sensorless and totally communication-free approach to foraging. Our approach uses minimal communication in order to ensure the activity of all other robots and thus adapt to any changes in territories. Even this communication could be eliminated and the information could be obtained through direct sensing, but at a relatively high interference cost. Similarly, Parker [16] employs communication for temporal division of labor in order to solve the task and minimize interference.

Agah and Bekey [19] introduce the Tropism-Based Cognitive Architecture for controlling groups of robots in a distributed fashion akin to ours. They demonstrate simulation results of a collection task similar to ours implemented in a homogeneous group and demonstrating effects of stationary vs. mobile obstacles on energy consumption and group performance.

Much work has been done in the field of Adaptive Behavior and Artificial Life on simulating insect colonies and studying their dynamics. This work has served as inspiration for robot experiments, and includes [20], [21], [22], [23], [24], [25], [26], and many others.

III. The Experimental Setup

A. The Robots

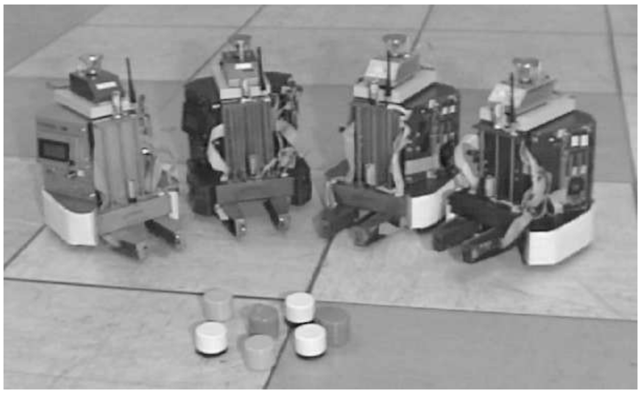

Fig. 1. The R2e robot family: the robots are equipped with a differentially-steerable base, a gripper, piezo-electric bump and gripper sensors used to detect collisions and to grasp pucks, infra-red sensors for collision avoidance, and a radio transmitter for absolute positioning and message passing.

The experiments were conducted with a group of four IS Robotics R2e robots programmed in the Behavior Language, a parallel distributed robot programming language based on the Subsumption Architecture [27], [28]. The robots are fully autonomous and equipped with with on-board power and processing. They consist of a differentially steerable wheeled base and a gripper for grasping and lifting objects. Their sensory capabilities include piezo-electric bump sensors for detecting contact-collisions and monitoring the grasping force on the gripper, and a set of infra-red (IR) sensors for obstacle avoidance (finger sensors) and grasping (break-beam sensors). Figure 1 shows the robot configuration.

The robots are also equipped with radio transceivers, used for determining absolute position and for inter-robot communication. Position information is obtained by triangulating the distance computed from synchronized ultrasound pulses from two fixed beacons, and updated at a rate of 1 Hz. Inter-robot communication consists of broadcasting 6-byte messages at the rate of 0.5 Hz. Once per second, each robot broadcasts a diagnostic "I am alive" message to the whole group, along with its ID number. If a robot fails to broadcast for a fixed period of time (in our experiments, 10 seconds), it is considered "dead" by the rest, who consequently adapt their behavior to the new group size and distribution. Hence, each of the robots is given the largest possible group size and can use it to determine how many others are cooperating on the task at each moment.

B. Definitions

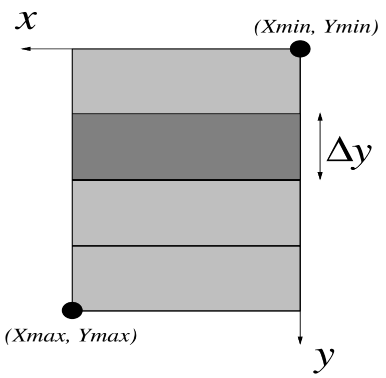

Fig. 2. Division of the robots' workspace.

The total workspace ($\Re$) is constrained to a rectangular area defined by the coordinates (X_{min}, Y_{min}) and (X_{max}, Y_{max}) (see Figure 2). The individual robots have different workspaces assigned *a priori*, and $\Re$ is divided into rectangular areas with equal height and width. The width of these regions is fixed to $\mid X_{max} - X_{min} \mid$, and independent of the number of working robots. The height (Δy) is dependent on the number of working robots and is defined as:

$$\Delta y = \frac{\mid Y_{max} - Y_{min} \mid}{Working\ Robots}$$

In order to adapt to the dynamically changing group size, each of the robots continually recomputes its adaptive Logical ID, denoted as i. Initially, each robot is given a pre-assigned unique hard-coded ID from the $[0..n]$ interval. During the lifetime of the task, this ID is used to compute the logical ID dynamically, based on the changing group size. Each robot computes its logical ID by determining how many other robots are not sending diagnostic messages, and how many of those have IDs lower than the robot's own ID ($\sigma Dead\ Robots$). Hence, the logical ID i is defined as:

$$i = ID - \sigma Dead\ Robots \tag{1}$$

The logical ID is used to dynamically compute the concepts of the working area, deposit area and deposit point:

- The Working Area (WA) for a robot (R_i) whose logical ID is i, is defined as the region where it looks for pucks, picks them up, and delivers them to the deposit area:

  $$WA(R_i) = \{(X_{min}, (i)\Delta y),\ (X_{max}, (i+1)\Delta y)\}$$

- The Deposit Area (DA) is defined as the region to which the robot (R_i) delivers the pucks:

  $$DA(R_i) = \{(X_{min}, (i-1)\Delta y), (X_{max}, (i)\Delta y)\};\ i > 0$$

  $$DA(R_i) = \{(X_{min}, Y_{min}), (X_{min}+\delta, Y_{min}+\delta)\}; i = 0$$

  where δ is the size of the "home" region.

- The Deposit Point (DP) is defined as the point towards which the robot will "home" to deliver the puck:

$$DP(R_i) = (\frac{X_{max}}{2}, (i-1)\Delta y + \frac{\Delta y}{2}); \quad \text{i} > 0$$

$$DP(R_i) = (X_{min}, Y_{min}); \quad \text{i} = 0$$

As it can be observed from the above definitions, the behavior of the robot whose logical ID is zero ($i = 0$) will be slightly different than that of the rest, as it will be in charge of placing all the pucks it grabs into the "home" region $DA(R_0)$, located in the corner of its work area (see Figure 3).

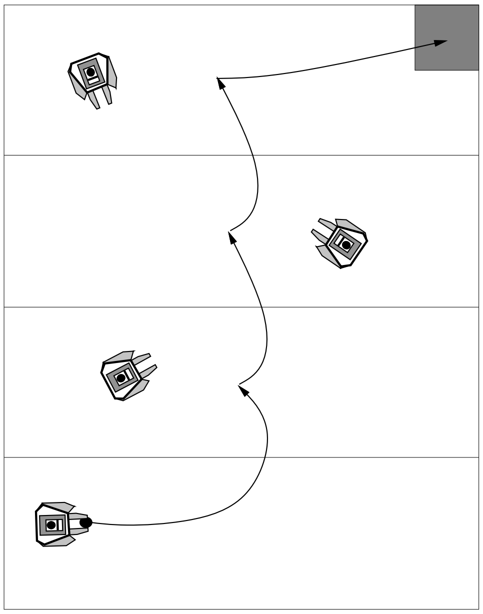

Fig. 3. A puck's path from a distant region to the "home" region (upper right corner).

IV. Behaviors

The main goal of the group is to collect all the pucks that are distributed over the global workspace, and take them to the "home" region. In order to reduce interference between robots, each robot is given a working area and a deposit area. Robots look for pucks while in their working area and deliver them once they reach the deposit area. While in the working area, robots actively search for pucks using the infra-red sensors in the forks. When a puck is detected, it is approached and picked up. Once a robot grabs a puck in its working area, it homes to its deposit point and delivers the puck there. Note that each robot has a different deposit point, determined by its logical ID. When the robot crosses the border between its working area and the deposit area, it leaves the puck, and as the robot is now without a puck and outside its working area, it homes back to its working area. While the robot is outside its working area, it does not try to pick up pucks and considers them obstacles. Therefore, the behaviors that use the sensor inputs from the finger IR sensors do not try to approach the pucks; instead they avoid them and any other obstacles on the way back to the robot's working area.

Figure 4 shows an outline of the robot's control architecture, following the classical Subsumption Architecture organization [27]. The behaviors are classified into three main categories: Survival, Collection, and Navigation. An auxiliary Listener behavior is not a part of the behavior hierarchy; it runs in parallel with the rest, and receives and broadcasts communication messages. Each of the behaviors is described next.

A. Listener

The **Listener** behavior sends and receives "I'm alive" messages, and its core functionality is to update the robot's logical ID using equation 1. The updated logical ID is then used to compute the coordinates of the Working Area (WA), the Deposit Area (DA), and the Deposit Point (DP), according to the expressions given in section III-B.

B. Survival Behaviors

- **Side-IR Avoid**: The Side IRs of the robot can measure objects close to the robot. In the case that an object (obstacle or another robot) is detected, this behavior will turn the robot away from it.
- **Finger-IR Avoid**: This behavior is active either when the robot is outside its WA or when the robot has grabbed a puck and is taking it to the Deposit Point. It uses the input of the finger IRs to avoid

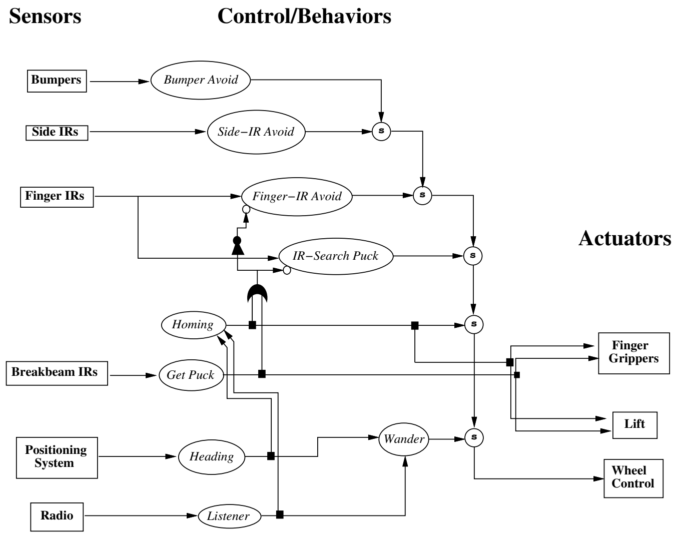

Fig. 4. Control architecture

any object that is in front of the robot.

- **Bumper Avoid**: In case the other two Survival behaviors (described above) fail, i.e., if an object is not detected by any of the IR sensors, but one or more bumper sensor are activated, this behavior performs an avoidance manoeuver that moves the robot away from the detected obstacle. For example, if the robot touches an object with the left bumper sensor, it will turn right.

C. Collection Behaviors

- **IR-Search Puck** behavior is only active when the robot is inside the WA and does not have a puck. It receives its inputs from the finger IR sensors, and directs the robot towards the puck.
  Obstacles and pucks are differentiated by this behavior because of the physical properties of the robot grippers. The space between the grippers is bigger than the size of a puck. Therefore, if an object is perceived with the left IR, the robot moves right, and conversely, if an object is perceived with the right IR, the robot moves left. In both cases the robot moves toward the "potential" puck. If both sensor perceive an object, that implies the object is larger than a puck, and is thus avoided, as it could not be picked up. An analogous approach was used by Connell [29], who employed it with a camera and a laser-range finder, in the task of using an arm to pick up soda cans.
- **Get Puck** behavior is based on the philosophy that "the world is the robot's own best representation" [29]. When the break-beam IR sensors detect a puck between the fingers, the behavior gets activated closes the gripper, and lifts up the fingers.

D. Navigation Behaviors

- **Homing** behavior is activated either 1) when the robot has grabbed a puck or 2) when the robot is outside its WA. In the former case, the robot's deposit point is computed as $DP(i)$, whereas in the latter case it is $DP(i+1)$, where i is the robot's logical ID (as defined in section III-B).
- **Wander**: Most of the interference between robots arises when they leave their working area and "invade" other areas, either due to sensor and/or effector noise and errors, or through the process of entering into another's territory when dropping off pucks. In order to minimize the interference while

the robots are looking for pucks, they are equipped with a wandering and searching behavior whose goal is to cover the entire working area without invading the neighbors. Since the search area in our environment is rectangular, the wandering behavior always tries to turn in the direction parallel to the longest side of the working area. Consequently, the robots perform a sweep of the area in the direction parallel with the neighbor workarea boundaries, trying to minimize the possibility of invading the neighbor's working area.

- **Heading**: The R2e robot family is not equipped with odometric sensors that could provide an estimate of the robot's current position and heading. However, the radio system provides the robots with absolute (x, y) data at the rate of 1 Hz. This data can be integrated over time to compute a rough heading estimate. Given two consecutive positions $P_0 = (x_0, y_0)$ and $P_1 = (x_1, y_1)$, the heading of the robot is simply computed as:

$$heading = \arctan\left(\frac{y_1 - y_0}{x_1 - x_0}\right)$$

The robot society utilized in this experiment is a homogeneous group in that all of the members have identical controllers consisting of the behaviors described above, and thus exhibit the same capabilities. The difference between individuals is only displayed in that each confines its behavior to a fixed territory, and can adapt that territory dynamically, depending on the size of the active group.

V. Sources of Interference

As in any multi-agent system, the interaction of a collection of robots produces interference, resulting from a competition for shared resources in the environment, in this case physical space and pucks [4]. In the forthcoming sections we outline those aspects that affect the behavior of the system, and their effect on the overall performance of the foraging task.

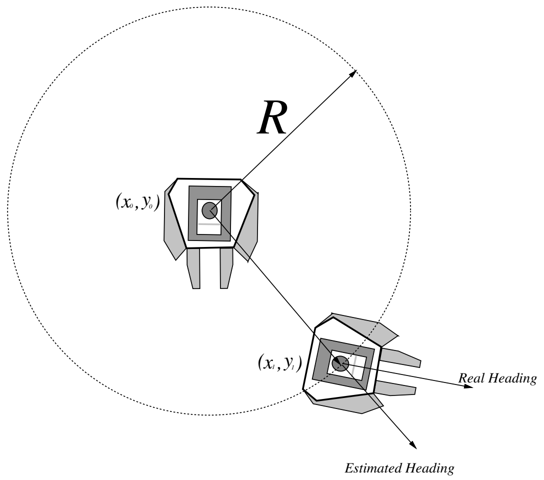

Fig. 5. Heading Estimation.

A. Heading Module

Two sources of errors affect the performance of the heading module:

- The robot must move straight and without turns at least by a fixed distance (R) to compute a new heading. However, during the period of time it takes the robot to move from P_0 to P_1, unexpected events, like obstacle avoidance, may occur, that may alter the heading estimate (see Figure 5).
- Errors in the sonar positioning system also affect the heading computation. Over the lifetime of the task, insignificant errors are filtered away, but in small workspaces even such errors can dramatically affect the overall achievement of the task. For example, false heading estimation might make the robot "invade" other areas more frequently, and thus increase overall interference.

B. Real World Dynamics

Figure 6 illustrates the situation in which a robot leaves a puck in its deposit area and interferes with the robot that is looking for pucks in its own working area. In this case, robot **a** has to home back to its working area, and may interfere with robot **b**.

The size of the working area effects the probability of interference between robots: the smaller the size of the

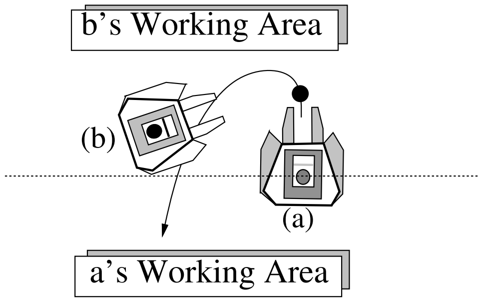

Fig. 6. Invasion, the process in which a robot that is dropping off a puck interferes with another that is searching within its own workspace.

working area the more probable it is that two robots may interfere (see quantitative results in section VI-C). In this case, interference has side effects in the computation of the heading of both robots: as they are avoiding each other they cannot cover the straight distance R.

VI. Experimental Results

A. Experimental Layout

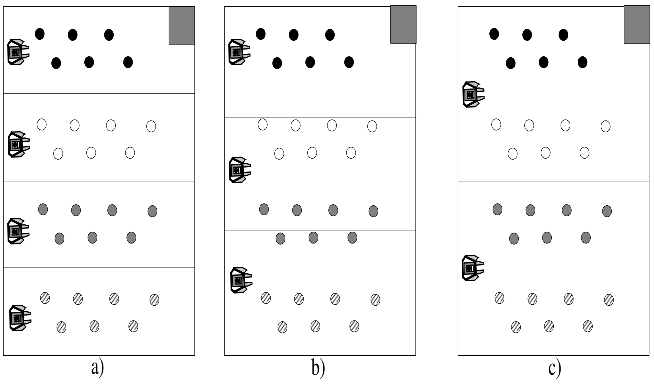

Fig. 7. The experimental layout used for three different sets of experiments.

Three different experiments were performed in order to quantitatively analyze the performance of the robots in the territorial foraging task. Figure 7 gives an overview of the different scenarios. In situation a), four robots were cooperating to perform the task, and the space was appropriately divided into four equally-sized regions. Robot "0", nearest to the "home" region, had six pucks in its working area, while the rest of the robots had seven. Scenarios b) and c) tested three and two robots, respectively. The initial puck distribution is shown, and differently-colored pucks are used in each of the regions in order to track puck movement over the duration of each trial, using video data gathering, in addition to position data. For each of the different layouts, all experiments were repeated five times; all tabulated data show mean values for each experimental setup. The relative error in the worst case was 7%, measured using standard deviation.

B. Data Gathering and Analysis Methods

The following approach was used to compute and analyze the robot data. For each each robot we monitored how long every one of its behaviors was active, and how much time the robot spent in each of the relevant areas around the workspace:

- in the working area $T_{working-area}$,
- outside the working area $T_{outside-working-area}$,
- delivering pucks $T_{delivering}$.

We also gathered two other types of data:

1. the number of pucks that had reached the "home" region 25 minutes after the beginning of the task,
2. the amount of time it took for the robots to place the 80% of the pucks (22 pucks) in the "home" region.

The total time the robot was working T_{total} could be simply computed as the sum of the times spent in each of the areas:

$$T_{total} = T_{working-area} + T_{outside-working-area} + T_{delivering}$$

Hence, the percentage of time spent outside the working area (φ_1) is computed as:

$$\varphi_1 = \frac{T_{outside-working-area} \times 100}{T_{total}}$$

and the percentage of time delivering pucks (φ_2) as:

$$\varphi_2 = \frac{T_{delivering} \times 100}{T_{total}}$$

These percentages were computed for each of the trials and the final results, shown below, are based on rounded mean values.

C. Performance

TABLE I

Percentage of delivered pucks in 25 minutes.

Number of Working Robots	% of Delivered Pucks
2	60 ± 2 %
3	80 ± 2 %
4	60 ± 3 %

Time performance in the three different scenarios is shown in Tables I and II. Table I indicates the percentage of the pucks that were successfully delivered to the "home" region in 25 minutes. Table II shows the mean time (in minutes), the robots needed to deliver the 80% of the pucks to the goal, plus the standard deviation of the mean value (σ).

The errors in the results indicate the deviation with respect to the mean performance, i.e. the repeatability of the experiments.

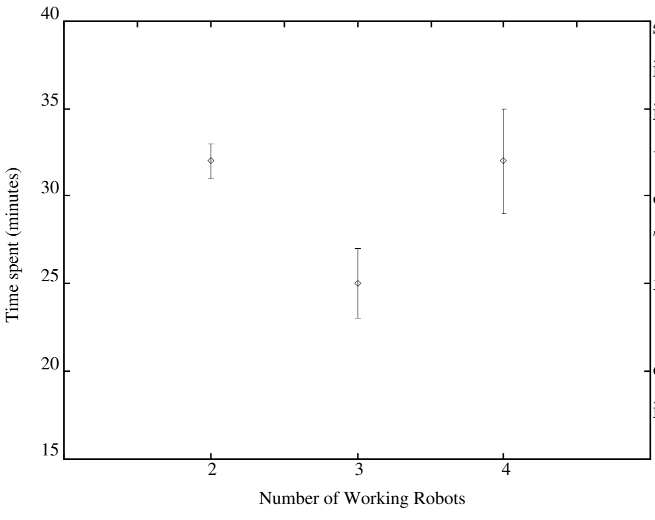

Fig. 8. Time spent in delivering 80% of the pucks.

Figure 8 shows the time the robots needed to deliver the 80% of the pucks to the goal. Dots mark mean values and the error bars show the best and the worst results. It is interesting to observe that the uncertainty in the measurements, derived from the height of the bars, and the standard deviation (σ), increase with the number of working robots. More robots working in the same global workspace area tend to interfere more, therefore increasing the uncertainty in the time required for delivering 80% of the pucks.

TABLE II

Time spent in delivering 80% of the pucks.

Number of Working Robots	80 % of the Pucks Delivered in $\pm\sigma$
2	32 ± 1 minutes
3	25 ± 2 minutes
4	32 ± 3 minutes

It may seem surprising that it takes four robots as long as it does two robots to deliver the same percentage of pucks. This effect results from the tradeoffs between interference, search space, and work-load per robot. The above described uncertainty in the heading module leads to an inaccuracy in the *wandering* behavior, which results in the robot leaving its working area and interfering with the others. If a robot leaves its working area, it turns around and tries to return to it. In this case the robot usually does not move the minimum straight distance to compute a new heading, as described earlier. Therefore, the smaller the working area, the more dramatically errors in the heading module affect the *wandering* behavior, and therefore increase the probability of a robot leaving its working area and interfering with its neighbors.

Table III presents the mean and standard deviation (σ) of the percentage of time spent by robots outside their working area. The results indicate that with narrower working areas robots tend to "invade" other areas more frequently.

Figure 9 is complementary to Table III, and shows the mean, best, and worst values of the percentages of time spent by each of the robots (marked with different IDs) outside their working area in the three different scenarios (2, 3, and 4 working robots).

From Figure 9 and Table III we can infer that robots belonging to working areas whose boundaries are connected to other robot's working areas (e.g., in the case of for four working robots, those with IDs 1 and 2), tend to leave their working area more frequently than others. Once the robots start to work, the system's dynamics (e.g., other robot's location) cannot be known in advance, and vary during the time course of one experiment and across different experiments. Hence, the uncertainty in the measurement (given by the error bars and the standard deviation) increases for robots whose working areas are surrounded by those of other robots, being therefore subject to higher interference.

TABLE III

PERCENTAGE OF THE TIME SPENT BY THE ROBOTS OUTSIDE THE WORKING AREA.

4 Robots

ID	Time $\pm\sigma$
0	26 ± 6 %
1	29 ± 7 %
2	29 ± 7 %
3	25 ± 3 %

3 Robots

ID	Time $\pm\sigma$
0	20 ± 3 %
1	24 ± 4 %
2	18 ± 2 %

2 Robots

ID	Time $\pm\sigma$
0	8 ± 3 %
1	8 ± 1 %

The robots are programmed to avoid any kind of object detected with their lateral IR sensors and their finger-IR sensors while they are outside their own working area. However, when obstacles are too close, or dif-

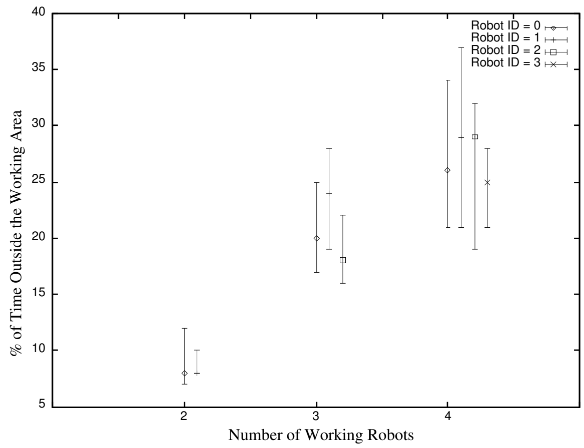

Fig. 9. Percentage of the time spent by the robots outside the working area.

ficult to detect (like other moving robots), the bumper-based avoid behavior takes control. Activation of this behavior is a good measurement of interference, because bumper hits are mostly due to inter-robot interference rather than interaction with other objects and boundaries of the environment, which are more likely to be detected with the IRs.

TABLE IV

BUMPER HITS.

Number of Working Robots	**Average of Bumper Hits** $\pm\sigma$
2	4 ± 1
3	7 ± 1
4	20 ± 5

Results shown in Table IV, combined with those shown in Tables I, II and III, highlight the fact that with increased group sizes the performance of the task division strategy declines due to inter-robot interference, explicitly expressed in terms of bumper hits which grows steeply with increased group size. For example, three robots accrue seven hits while four robots accrue twenty. Three robots perform the task better than two robots (as can be seen in Tables I and II), but interference increases, as is clearly demonstrated in Table IV. Overall, of the two-, tree-, or four-robot sets we tested, three robots are the most efficient choice given the trade-off between interference and work-load in the particular territorial division.

D. Scalability

TABLE V

PERCENTAGE OF THE TIME SPENT BY THE ROBOTS DELIVERING PUCKS

4 Robots

ID	Time $\pm\sigma$
0	28 ± 2 %
1	11 ± 2 %
2	8 ± 3 %
3	3 ± 1 %

3 Robots

ID	Time $\pm\sigma$
0	27 ± 6 %
1	12 ± 3 %
2	6 ± 1 %

2 Robots

ID	Time $\pm\sigma$
0	27 ± 3 %
1	10 ± 1 %

Table V shows the percentage of time spent by the robots delivering pucks. Note the fact that this amount of time remains almost constant for the robot labelled with $ID = 0$, which is in charge of placing the pucks in the definitive "home" region, and thus must handle *all* of the pucks, regardless of the size of the group.

E. Homogeneous case: Comparison

In order to compare time and interference performance of the space division strategy, we performed control experiments using homogeneous groups of two, three and four robots with identical behavioral abilities and no space division. Each of the robots placed the found pucks in the same home region.

Experiments with only one robot are not shown in Table VI, as it was unable to deliver 80% of the pucks in less than 30 minutes, due to the limited battery lifetime relative to the size of the working area. A single robot was able to collect a 40% of the pucks in 25 minutes.

TABLE VI

TIME SPENT IN DELIVERING 80% OF THE PUCKS. (HOMOGENEOUS CASE)

Number of Working Robots	80 % of the Pucks Delivered in $\pm\sigma$
2	30 ± 2 minutes
3	24 ± 2 minutes
4	20 ± 3 minutes

Table VI shows the amount of time needed by the robots to deliver the 80% of the pucks. Time performance is almost the same with 2 and 3 robots in the homogeneous case as in the space-division strategy. Four robots perform much better in the homogeneous case than when using the space-division strategy. However, 50% of the experiments with 4 homogeneous robots had to be repeated because the robots collided and became trapped (and unable to move) in the home region. The results in Table VI only show the data gathered during the experiments that finished successfully. Data in Table VII, showing the average of bumper hits in different trials, demonstrate a much higher interference rate among the robots.

TABLE VII

BUMPER HITS (HOMOGENEOUS CASE)

Number of Working Robots	Average of Bumper Hits $\pm\sigma$
2	14 ± 3
3	26 ± 7
4	40 ± 9

VII. DISCUSSION AND CONCLUSIONS

Our experimental setup was designed to minimize inter-robot interference, and thus to achieve high task efficiency. This efficiency is directly dependent on the spatial density of the agents, since the most fundamental type of resource competition the robots encounter is that for space. Based on the examination of the quantitative data recorded during repeated trials, we found the sources of uncertainty, inherent to embodied agents, that demonstrate the size of the working areas and robot density for most effective foraging in this particular experiment. Furthermore, we postulate that robot data was critical in locating these sources, since simulation results alone would have most likely failed to accurately reproduce the relevant effects, as detailed in Section VI.

We have empirically demonstrated that an increased group size can, in physical robots, negatively impact the effectiveness of the territorial solution. Intuitively, in a constrained environment, small groups can perform their task better than larger groups simply as a result of the direct relationship between robot density and interference. In future experiments we plan to explore the effects of much lower density on the effectiveness of the group, by significantly increasing the workspace and holding the group size constant.

Matarić [4] describes a methodology for estimating interference through computing the spatial density of the particular environment-task combination using the robot's physical size, kinematics of movement, and sensory range. This computation must take into account realistic sensory models with appropriate error and uncertainty. The computed value provides a fixed approximation of an efficient group size for the given workspace, assuming no drastic changes in either. However, the computation must be performed every time the size of the environment changes or one or more robots fail. Thus, the density value is best used as a rough approximation.

We are considering an extension that allows superfluous robots to be removed from the workspace. As shown, given the specific parameters of our physical robots and the size of their workspace, the ideal number of robots for the foraging task is smaller than the complete set of available robots. If this fact cannot be computed before the robots are assigned the tasks, then it would be useful to be able to establish it dynamically, as we have demonstrated.

This work focuses on one approach to minimizing interference: the use of spatial division of the task. Interagent interference is minimized as a consequence of spatial isolation of the agents' territories. Note that the agents themselves are not further specialized, but are homogeneous across the group. This approach is neatly complementary to the alternatives of using heterogeneous agents (i.e., employing different types of agents that may or may not be spatially isolated) and temporal isolation (i.e., employing division of labor in time rather than space). Other work by Goldberg and Matarić [30] demonstrates the heterogeneous agent alternative using an implementation of the foraging task on the same robot family.

We plan to compare and tie our territoriality and robot density results to task division behavior in animal societies in order to find common principles of group organization. We hope that our results can be informative and useful toward the challenging goal of synthesizing group behavior in mobile robots.

Acknowledgements

Miguel Schneider Fontán's work is partially supported by the Spanish Ministry of Science under project I.Ares: CICYT, TAP 816193-C03-2. Maja Matarić's work is supported by the Office of Naval Research Grant N00014-95-1-0759 and the National Science Foundation Infrastructure Grant CDA-9512448.

The authors wish to thank Dani Goldberg for his tireless help with the misbehaving robots, Lourdes Agapito for her useful comments on earlier versions of the paper, and Barry Werger for editing the video. They are also grateful to the anonymous reviewers for their comments.

References

[1] Maja J. Matarić, "Behavior-based systems: Key properties and implications", in *IEEE International Conference on Robotics and Automation, Workshop on Architectures for Intelligent Control Systems*, Nice, France, May 1992, pp. 46–54.

[2] Rodney A. Brooks, "Challenges for complete creature architectures", in *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, Jean A. Meyer and Stewart Wilson, Eds. 1990, MIT Press.

[3] Maja J Matarić, "Designing and understanding adaptive group behavior", *Adaptive Behavior*, vol. 4, no. 1, pp. 50–81, December 1995.

[4] Maja J Matarić, "Interaction and intelligent behavior", Tech. Rep. AI-TR-1495, MIT Artificial Intelligence Lab, 1994.

[5] Maja J. Matarić, "Minimizing complexity in controlling a collection of mobile robots", in *IEEE International Conference on Robotics and Automation*, Nice, France, May 1992, pp. 830–835.

[6] James L. Gould, *Ethology; The Mechanisms and Evolution of Behavior*, W. W. Norton & Co., NY, 1982.

[7] David McFarland, "The oxford companion to animal behavior", in *Oxford, University Press*, 1987.

[8] Maja J Matarić, "Issues and approaches in the design of collective autonomous agents", *Robotics and Autonomous Systems*, vol. 16, no. 2–4, pp. 321–331, December 1995.

[9] Maja J Matarić, "Reinforcement learning in the multi-robot domain", *to appear in Autonomous Robots*, 1996.

[10] Maja J. Matarić, "Designing emergent behaviors: From local interactions to collective intelligence", in *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, J-A. Meyer, H. Roitblat, and S. Wilson, Eds., 1992.

[11] Ronald C. Arkin, "Cooperation without communication: Multiagent schema based robot navigation", *Journal of Robotic Systems*, vol. 9, no. 3, April 1992.

[12] Tucker Balch and Ronald C. Arkin, "Communication in reactive multiagent robotic systems", *Autonomous Robots*, vol. 1, no. 1, 1995.

[13] Ronald C. Arkin, Tucker Balch, and Elizabeth Nitz, "Communication of behavioral state in multi-agent retrieval tasks", in *IEEE International Conference on Robotics and Automation*, Los Alamitos, CA, 1993, pp. 588–594, IEEE Computer Society Press.

[14] Ronald C. Arkin and J. David Hobbs, "Communication and social organization in multi-agent systems", in *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pp. 486–493. MIT Press, 1993.

[15] Ronald C. Arkin and Khaled S. Ali, "Reactive and telerobotic control in multi-agent systems", in *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pp. 473–478. MIT Press, 1994.

[16] Lynne E. Parker, "Adaptive action selection for cooperative agent teams", in *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, 1992, pp. 442–450.

[17] Lynne E. Parker, *Heterogeneous Multi–Robot Cooperation*, PhD thesis, MIT, February 1994.

[18] Ralph Beckers, Owen E. Holland, and Jean Louis Deneubourg, "From local actions to global tasks: Stigmergy and collective robotics", in *Artificial Life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems*, Rodney Brooks and Pattie Maes, Eds. 1994, MIT Press.

[19] Arvin Agah and George A. Bekey, "Autonomous mobile robot teams", in *AIAA/NASA Conference on Intelligent Robots in Field, Factory, Service, and Space*, Houston, Texas, 1994, vol. 1, pp. 246–251.

[20] Jean L. Deneubourg, Simon Goss, J. M. Pasteels, D. Fresneau, and J. P. Lachaud, "Self-organization mechanisms in ant societies, ii: Learning in foraging and division of labor", *From Individual to Collective Behavior in Social Insects*, vol. 54, pp. 177–196, 1987.

[21] Jean-Louis Deneubourg and Simon Goss, "Collective patterns and decision-making", in *Ethology, Ecology and Evolution 1*, 1989, pp. 295–311.

[22] Jean L. Deneubourg, Simon Goss, Nigel Franks, A. Sendova-Franks, C. Detrain, and L. Chretien, "The dynamics of collective sorting", in *From Animals to Animats: International Conference on Simulation of Adaptive Behavior*, pp. 356–363. MIT Press, 1990.

[23] Jean L. Deneubourg, G. Theraulaz, and R. Beckers, "Swarm-made architectures", in *Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, F. Varela and P. Bourgine, Eds., pp. 123–133. MIT Press, 1992.

[24] B. Corbara, A. Drogoul, D. Fresneau, and S. Lalande, "Simulating the sociogenesis process in ant colonies with manta", in *Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, 1993, pp. 224–235.

[25] Alberto Colorni, Marco Dorigo, and Vittorio Maniezzo, "Distributed optimization by ant colonies", in *Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, F. Varela and P. Bourgine, Eds., pp. 134–142. MIT Press, 1992.

[26] Alexis Drogoul, Jacques Ferber, Bruno Corbara, and Dominique Fresneau, "A behavioral simulation model for the study of emergent social structures", in *Toward A Practice of Autonomous Systems: Proceedings of the First European Conference on Artificial Life*, F. Varela and P. Bourgine, Eds., pp. 161–170. MIT Press, 1992.

[27] Rodney A. Brooks, "A robust layered control system for a mobile robot", *IEEE Journal of Robotics and Automation*, vol. RA-2, pp. 14–23, April 1986.

[28] Rodney A. Brooks, "The behavior language; user's guide", Tech. Rep. AIM-1227, MIT Artificial Intelligence Lab, 1990.

[29] Jonathan H. Connell, *Minimalist Mobile Robotics: A Colony Architecture for an Artificial Creature*, Academic Press, 1990.

[30] Dani Goldberg and Maja J. Matarić, "Interference as a Tool for Designing and Evaluating Multi-Robot Controllers", in *Proceedings, AAAI-97*, Providence, RI, AAAI Press, pp. 637–642, 1997.