Robust Behavior-Based Control for Distributed Multi-Robot Collection Tasks

Dani Goldberg and Maja J Matarić

Abstract—
We demonstrate the effectiveness of behavior-based control in facilitating the development and evaluation of multi-robot controllers that are: (1) robust to robot failures, and (2) easily modified to facilitate development of the controller variation that sufficiently satisfies the design requirements for the task. Our experimental focus here is *distributed multi-robot collection*, a class of tasks that includes de-mining and toxic waste clean-up. We demonstrate a basic multi-robot controller for the collection task, then show how to easily derive two spatio-temporal variations with markedly different performance properties. We evaluate the desirability of these controllers with respect to design requirements involving inter-robot interference, time-to-completion, and energy expenditure. The data for evaluation come from experiments using four physical mobile robots performing the three variations of the collection task.

Keywords— **multi-robot control, behavior-based control, distributed control, distributed object collection**

I. Introduction

Designing and implementing robust controllers for multiple interacting mobile robots is considered something of a black art, often involving a great deal of reprogramming and parameter adjustment. It is difficult enough to develop a multi-robot controller that functions only under the ideal conditions of little noise and no robot failures. The fact that such ideal conditions do not often exist, even in a laboratory setting, places certain practical requirements on the multi-robot controller. In particular, the controller must exhibit group-level robustness to noise and robot failures. This is especially important when physical human intervention is difficult (e.g., a toxic waste spill) or impossible (e.g., an extraterrestrial mission).

Additional design requirements for the controller arise from the fundamental, constrained resources of the system, including energy, time, and the number of robots. Untethered mobile robots are generally powered by batteries and can only perform a limited amount of work before needing recharging. Minimizing energy utilization is thus often required in domains, such as space exploration, where recharging is expensive, difficult or time consuming. In time-critical domains, such as search and rescue, the requirement is for expedient execution of the task. Additionally, regardless of the domain, the fragility of the robots may require the controller to maintain both robot-object and inter-robot collisions at a minimum.

D. Goldberg is with the Computer Science Department, University of Southern California, Los Angeles, CA 90089-0781 USA (e-mail: dani@usc.edu)

M. J Matarić is with the Computer Science Department, University of Southern California, Los Angeles, CA 90089-0781 USA (e-mail: mataric@usc.edu)

For a given task environment and set of robots, the requirements for the controller may not be independent but instead arise as tradeoffs. For example, minimizing both time and inter-robot collisions may not be possible since faster moving robots are less likely to properly sense each other and thus more likely to collide. Different controller variations may have to be tested and compared in order to select one that sufficiently satisfies the requirements given the tradeoffs among them. This places an additional requirement on the controller, namely that it be easily modifiable. The testing and comparison of the variations could potentially be accomplished analytically if an adequate model of the system were developed (a significant challenge in itself), or in simulation (potentially less difficult). In either case, the desire to be able to easily modify the controller remains. Our assumption in this work is that neither an adequate (i.e., very high fidelity) model nor simulation of the physical multi-robot collection task need exist, and thus we performed all tests directly on physical robots.

The controllers we present in this paper are designed to address the requirments above. Specifically, they exhibit group-level robustness to robot failures and noise, and are easily modified. Our focus is on the domain of distributed multi-robot collection (foraging) tasks, including toxic waste clean-up and de-mining. We present a basic *homogeneous* controller for the collection task in which all of the robots have identical behavioral repetoirs and work concurrently. We then derive two other variations, *pack* and *caste*, which respectively modify the robots' temporal and spatial interactions. Finally, we evaluate and compare the performance of the controllers using three spatio-temporal criteria: inter-robot collisions, distance traveled by each robot, and time-to-completion for the task. The latter two criteria also provide an indication of the energy expenditure of the robots. The data for evaluation come from experiments we conducted using four physical mobile robots performing the three variations of the collection task.

After a review of related work in Section II, Section III describes the structure of the collection task as well as the group of physical mobile robots that performed it. Section IV then presents the details of the homogeneous controller including the behaviors it contains and how it achieves robustness. Section V considers spatio-temporal interactions between robots, especially physical interference, and motivates the two interference-modifying controller versions, pack and caste, presented in Sections VI and VII. Section VIII presents an analysis of the controllers using data from physical experiments, and provides a comparative evaluation. Finally, conclusions are presented in Section IX.

II. Related Work

This section gives a partial review of some of the most related robotics work. The reader is encouraged to see [1] and [2] and for a more complete set of references from Artificial Intelligence, Robotics, Distributed AI, and Artificial Life.

Much research has been conducted on the performance and properties of robot collection and foraging tasks. [3] demonstrates simulation work studying the issues of density and critical mass in a collection task using fully homogeneous robots. Density and critical mass of robots has significant effect on the manifested physical interference between robots. [4] describes the general schema-based control architecture, which bears some fundamental similarities to behavior-based control, which we used, and give the critical mass experiments. [5] presents a series of simulation results on related spatial tasks such as foraging, grazing, and herding. This work is similar to the homogeneous controller implementation presented in Section IV.

[6] describes a similar behavior-based approach for minimizing complexity in controlling a collection of robots performing various behaviors including following, aggregation, dispersion, homing, flocking, and foraging (similar to our collection task). The work also includes a simulated dominance hierarchy based on IDs and used to evaluate performance of homogeneous versus ordered aggregation and dispersion behaviors. Our pack controller (Section VII) also utilizes a dominance hierarchy based on robot ID. In other work, we have demonstrated that allowing dynamic reorganization of such a dominance hierarchy can improve group performance [7]. [8] also presents work on multi-robot collection, but focused on issues of critical mass in territorial task division, corresponding to an extreme case of the caste controller we present in Section VI. [9] also presents a caste-like strategy for multi-robot search, but with each robot able to dynamically switch its team and function. [10] and [11] describe multi-robot experiments on foraging R2e robots with *a priori* hard-wired heterogeneous capabilities using the Alliance architecture. [11] describes a temporal division that sends one robot to survey and measure the environment for toxic spills, then has the rest of the group use its information to clean up the spill. This two group division is similar to our caste controller, though ours does not use a temporal division.

[12] describes an approach to maintaining a geometric configuration of a robot group using virtual structures, tested on a group ISR R3 mobile platforms, a later generation of our R2e robots. Similar to our homogeneous implementation, this work also exhibits spatial and temporal homogeneity, though the coupling here is tighter. [13] describes a group of five robots without external sensing or communication effectively clustering pucks through a careful combination of the mechanical design of the robots' puck scoops and the simple controller that moves them forward and in reverse. This work demonstrates a homogeneous controller performing a task similar to our collection task, but where the goal location is not pre-specified, instead emerging during execution. [14] presents more recent results from an expanded study with essentially the same experimental scenario.

Other work on multi-robot collection is inspired by trail formation in ants [15]. [16] describes a foraging robot chain that is constructed and modified using only contact sensing for communication. [17] presents multi-robot ant-like foraging in a simulated environment where efficient foraging trails are dynamically constructed using a mechanism analogous to ant pheromones.

Apart from the collection tasks, behavior-based control has been used in many other applications ranging from multi-robot soccer [18] and service robotics [19], to control of underwater robots [20] and ape-like robots [21]. In all of these behavior-based systems, there is some action selection mechanism that produces a coherent, global behavior. The work described in this paper uses behavior arbitration in which some (possibly small) subset of the behaviors control the motors at any time. [22] describes a number of action selection mechanisms. [23] presents a voting-based action selection mechanism which is extended to multi-robot coordination in [24].

Unlike the work in this paper, our other work involving robot collection focuses not on controller design, but rather on endowing the robots with the ability to adapt to their changing environment. The robots employ augmented Markov models [25] to detect significant shifts in object densities [26], or estimate the state of their environment in order to maximize performance of the collection task [27]. The following section describes the collection task in detail.

III. The Collection Task

The controllers we present implement versions of a multi-robot collection (foraging) task, a prototype for various applications including distributed solutions to de-mining, toxic waste clean-up, and terrain mapping. We present the general structure of the collection task, our multi-robot testbed, and then the controllers.

A. Task Structure

We define the collection task as a two-step repetitive process in which:

1. N ($N \geq 1$) robots search designated regions of space for certain objects, and
2. once found, these objects are brought to a goal region using some form of navigation.

A region in the task is any contiguous, bounded space (in the case of mobile robots, a planar surface) which the robots are capable of moving across. There are three mutually-exclusive, non-overlapping types of regions:

- search regions, S, containing P objects, a fraction of which must be delivered to a goal region;
- goal regions, G, where objects are delivered;
- and, optionally, empty regions, E, that contain no objects and are not goal regions.

The only restrictions placed on the configuration of the regions for the collection task are that there be at least one search and one goal region, and the union of all the regions be contiguous. Figure 1 gives two examples of possible region configurations for the collection task.

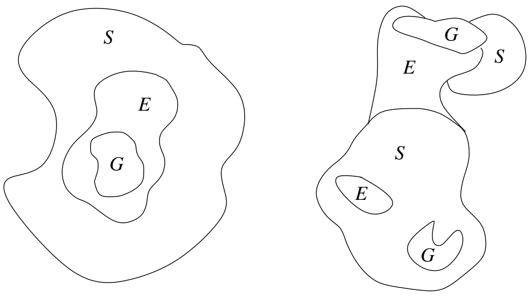

Fig. 1. Two example region configurations for the collection task.

The specific configuration we used is shown in Figure 2. The experiments were performed in an 11 × 14 foot rectangular enclosure (the Corrall). The search region, S, is approximately 126 square feet and has $P = 27$ small metal cylinders (pucks) evenly distributed throughout. The goal region G, also called *Home*, is a ninety degree sector of a circle with a radius of 2 feet, located in one corner of the Corrall. Finally, there is a 25 square foot empty region, E, separating the search and goal regions. E is composed of the Boundary and Buffer zones, whose functions will be described in the next section. $N = 4$ robots are used in the experiments.

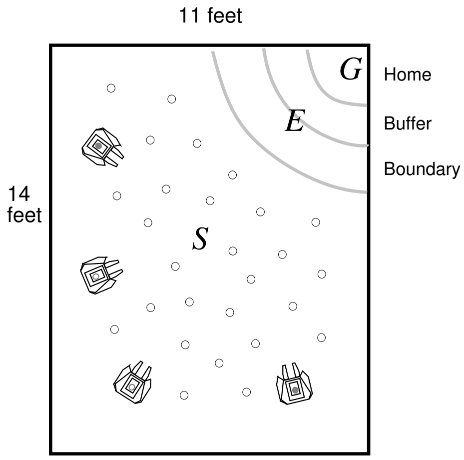

Fig. 2. Actual configuration used in the collection task.

B. The Robots

Four IS Robotics R2e robots were used (Figure 3). Each is a differentially-steered base equipped with two drive motors and a two-fingered gripper. The sensing capabilities of each robot include piezo-electric contact (bump) sensors around the base and in the gripper, five infrared (IR) sensors around the chasis and one on each finger for proximity detection, a color sensor in the gripper, a radio transmitter/receiver for communication and data gathering, and an ultrasound/radio triangulation system for positioning (Figure 4). The robots are programmed in the Behavior Language [28], a parallel, asynchronous, behavior-based programming language inspired by the Subsumption Architecture [29]. The main computational power on each robot is a single Motorola 68332 16-bit microcontroller running at 16 MHz. Even though computationally impoverished by today's standards, the processing capabilities have proven to be adequate for most tasks we have envisioned, helping to show that robust, effective control need not be computationally expensive. Perhaps the greatest drawback of the 68332 is its lack of floating point computation, which, for example, influences our calculation of heading, described in the following section.

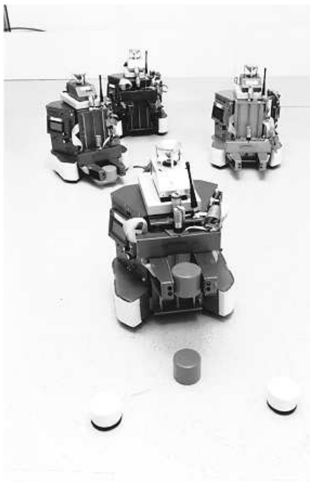

Fig. 3. The four R2e robots used in the experiments.

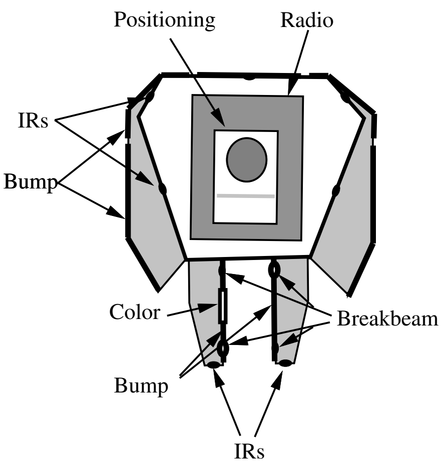

Fig. 4. The sensor configuration of an R2e robot.

C. Behavior-Based Control

The work presented in this paper is couched in the framework of distributed behavior-based control [30], [31]. Behavior-based control has proven to be an effective paradigm for developing single-robot and multi-robot controllers [32]. In behavior-based control, the robot controller is organized as a collection of modules, called behaviors, that receive inputs from sensors and/or other behaviors, process the input, and send outputs to actuators and/or other behaviors. Each behavior generally serves some independent function, such as *avoiding* obstacles or *homing* to a goal location. All behaviors in a controller are executed in parallel, simultaneously receiving inputs and producing outputs. An action selection mechanism prevents conflicts when multiple outputs are sent to actuators or other behaviors [22]. The controllers presented in this paper demonstrate the suitability of the behavior-based paradigm for designing robust and modifiable multi-robot controllers.

In the next section, we present our initial, *homogeneous* controller for the collection task, followed later by two variations, *pack* and *caste*.

IV. The Homogeneous Controller

In this section, we present the first behavior-based controller which implements a homogeneous version of the collection task where the robots' behavioral repetoirs are identical, and the robots act concurrently and independently.

The overall structure of the controller is presented in Figure 5. In the figure, the rounded rectangles represent the robot's sensors, with sensor values being transmitted to behaviors along the dotted lines. The behaviors themselves are drawn as ellipses with text in one of three font styles: italics for behaviors that only receive sensor inputs; bold for behaviors that send actuator outputs; and bold-italics for behaviors that do both. The dashed lines represent commands sent by behaviors to the actuators (rectangles), and the solid lines represent control signals sent between behaviors. These control signals include: inhibition signals that temporarily disable behaviors, or do so permanently until the inhibition is lifted; information about the state of the behaviors; and signals indicating that a behavior should perform a certain action. These control signals establish the hierarchy of actuator commands shown at the right of the diagram. The $\otimes$ represents behavior selection and indicates that only one of relevant actuator command pathways is active at any time. The $\odot$ represents a Subsumption-style priority scheme with the actuator command coming from above taking precedence [29]. The hierarchy of command pathways in the diagram helps illustrate that behavior arbitration is the action selection mechanism for the controller. The next section presents in detail the function of the each behavior in the controller, and the structure of the inter-behavior command pathways. The subsequent section discusses the group-level robustness achieved by this controller.

A. Behaviors

In order to provide a clear picture of the interaction between behaviors, we describe the individual behaviors of the controller in an order that mirrors the progression of the task as the robot performs it. The following twelve behaviors constitute the collection task:

1) *avoiding*: This behavior avoids any object (including other robots) detected by the IR sensors and deemed to be in the path of, or about to collide with, the robot. If the robot has already collided with an object, as detected by the contact sensors, it steers away from it. This behavior is critical to the safety of the robot and therefore takes precedence over most of the behaviors that control the drive motors (*puck detecting*, *wandering*, *homing*, *reverse homing*).

2) *wandering*: The robot moves forward and, at random intervals, turns left or right through some random arc. Using this behavior, the robot searches a region for pucks.

3) *puck detecting*: If an object is detected by the front IR sensors while *wandering*, this behavior, by lifting the gripper, determines whether the object is short enough to be a puck, or whether it is an obstacle that must be avoided. If it is a puck, the robot carefully approaches the object and attempts to place it between its fingers. Otherwise, the robot performs *avoiding*.

4) *puck grabber*: When a puck enters the fingers and is detected by the breakbeam IR sensors, this behavior grasps it and raises the fingers. Raising the fingers above puck height prevents the robot from unnecessarily avoiding pucks while *homing*, and allows the robot to collect up to about four additional pucks with its base.

5) *homing*: If carrying a puck, the robot moves towards the designated goal location, Home. While *homing*, *avoiding* can take precedence in order to avoid obstacles.

6) *boundary*: This behavior monitors how the robot enters the Boundary region. If the robot enters this region without a puck, it returns it to the search region using *reverse homing*. If carrying a puck, the robot is allowed to enter this region and proceed towards Home (see Figure 2). This behavior prevents the robot from collecting pucks that have already been delivered.

7) *buffer*: This behavior monitors entry into the Buffer region. Entering this region triggers the activation of the *creeping* behavior.

8) *creeping*: This behavior is a refined combination of the *homing* and *avoiding* behaviors designed to carefully bring the robot to the very corner of the Corrall where Home is located and where the pucks must be delivered. Under *creeping*, the robot moves more slowly and uses its IR sensors at a closer range appropriate for working within the corner. The standard versions of *homing* and *avoiding* would conflict in a confined corner situation, since *avoiding* would perceive the goal corner as an obstacle and attempt to move the robot away from it. *Creeping* takes precedence over *avoiding* since it already incorporates a version of this behavior.

9) *home detector*: A monitoring behavior for entry into the Home region. Upon entering this region, *home detector*

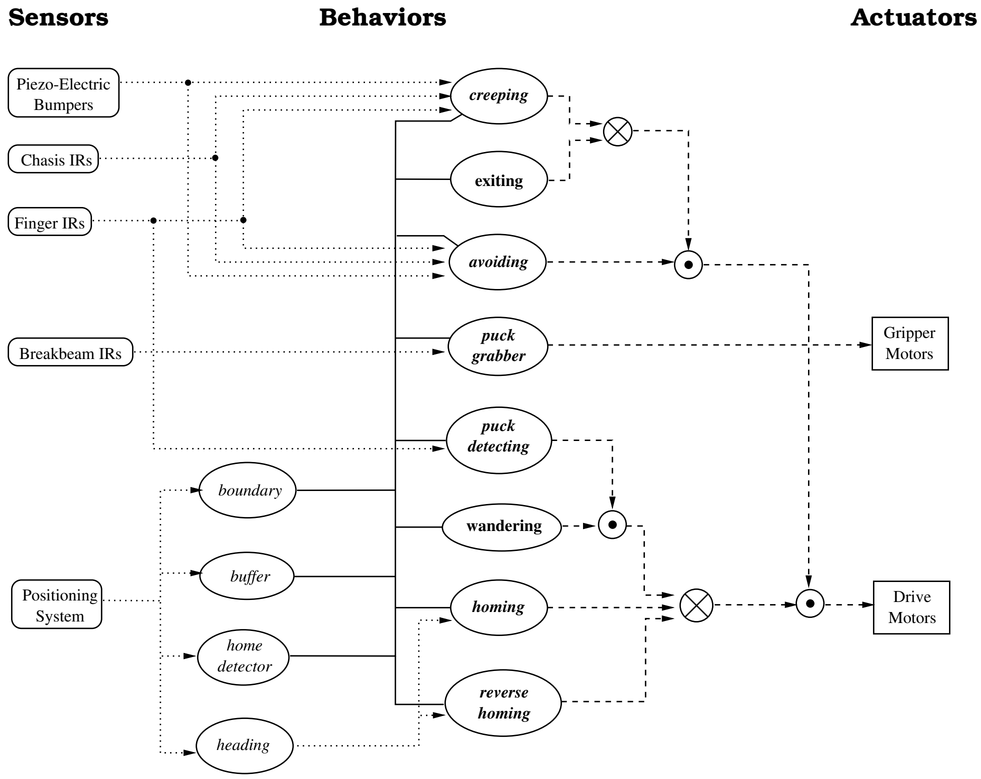

Fig. 5. The homogeneous controller for the collection task. Rounded rectangles represent the robot's sensors, ellipses represent behaviors, and rectangles represent actuators. Sensor values are transmitted along dotted lines, actuator commands along dashed lines, and inter-behavior control signals along solid lines. The $\otimes$ represents behavior selection and the $\odot$ represents Subsumption-style precedence.

sends a signal to *puck grabber* to release the puck.

10) *exiting*: Entering the Home region triggers this behavior which moves the robot several inches backwards, then performs a 180-degree turn in place. This behavior also sends the signal that lowers the gripper. When *exiting* terminates, the robot remains within the Boundary region without a puck. This in turn triggers the *boundary* behavior to begin *reverse homing*.

11) *reverse homing*: Starting from within the Boundary region, this behavior performs the opposite of *homing*; it moves the robot out into the search region. This behavior is essentially identical to *homing* except that the goal location is set to the corner of the Corrall opposite Home. Once the Boundary region has been left, *reverse homing* becomes inactive and the robot once again begins searching for pucks using *wandering*.

12) *heading*: This behavior processes the positioning system data and provides approximate heading values for the *homing* and *reverse homing* behaviors. The positioning system supplies the robot's current (x, y) position at approximately 1–2 Hz. Consecutive position values, (x_0, y_0) and (x_1, y_1), are used in an approximate integer-based calculation of $\arctan(\frac{y_1-y_0}{x_1-x_0})$ adjusted for the quadrant of the angle to provide one of sixteen possible sector headings. The accuracy of this heading calculation is usually within one sector of the true heading, but may be far worse when the robot turns in place. Frequent updates of the heading, with little reliance by the other behaviors on any one heading value, help to compensate for the inaccuracies. (An alternative is to use a physical compass for heading data. In our lab, however, the high variance in magnetic fields makes this an inviable option.)

B. Robustness

In the above described homogeneous controller, group-level robustness is a direct result of the robots behaving identically and independently. No noise-susceptible, or time-critical, radio communication that could be a source of fragility in the system is necessary. Each robot must individually manage the noise and uncertainty associated with its sensors and actuators, and the complexity of a dynamic and basically unknown environment. (Our controller, as is true for most behavior-based controllers, accommodates noise and uncertainty by tightly coupling sensing to action so that no great reliance is placed on any one sensor reading.) The partial, or even complete, failure of any one robot, or a subset of them, does not debilitate the entire group. As long as there is one functioning robot, the task will be accomplished.

As discussed previously, in addition to exhibiting group-level robustness, a multi-robot controller should be easy to modify in order to facilitate the search for a acceptable variation. The desirability of the controller must be measured with respect to any design requirements, such as time-to-completion of the task, energy consumption, or the amount of interference exhibited. Thus, before we present the variations of our homogeneous controller, we discuss the key diagnostic parameters used in evaluation. Our focus here is on inter-robot interference, specifically physical collisions between robots. The goal that motivates the modification of the homogeneous controller is minimization of such interference. The next section provides a discussion of interference and the two spatio-temporal solutions to it which provide the basis for our controller variations.

V. Spatio-Temporal Interactions

In this section, we discuss the nature of physical inter-robot interference (i.e., collisions), and how a multi-robot system may be modified to manipulate this interference. Our discussion here provides the motivation for the two controller variations, *pack* and *caste*, presented later.

Multi-robot systems are by definition physically embodied and embedded in dynamic environments. The types of interference they contain can be distinguished about a physical/non-physical dichotomy. Physical interference manifests itself most overwhelmingly in competition for space. Non-physical interference ranges from the sensory (shared radio bandwidth, crossed infrared or ultrasound sensors) to the algorithmic (the goals of one robot undoing the work of another, competing goals, etc.). Here we focus on physical interference and demonstrate that it is an effective tool for system evaluation and design.

We define the *characteristic interference* of a system at a particular point in space to be the sum, over some finite time period, of all measured interference occurring at that location (see Figure 6). The result is a surface that can be used to adjust the controller in order to reduce interference and thus modify the system's overall performance. Robot density is a critical factor in characteristic interference. Single-robot systems and systems with density so high as to prevent movement produce no characteristic interference. Systems of interest lie in between the two extremes.

A principled multi-step process of controller modification can be implemented by using characteristic interference as a guide indicating where in the robots' physical interaction, and when within the lifetime of the task, behaviors should be switched and the task should be divided to modify overall task interference. Multi-robot interactions we focus on are spatio-temporal in nature and fall into four basic categories. Robots may either be in the *same place* (SP) or in *different places* (DP), both of which can occur at *same time* (ST) or at *different times* (DT), resulting in four forms of interaction: SPST, SPDT, DPST, and DPDT.

Physical interference fits into the SPST category, covering the case when two or more robots try to occupy the same location at the same time. The other three categories are useful for deriving and fine-tuning controllers that modify SPST interactions. For two of these categories, we implemented and tested a corresponding controller. The SPDT category is associated with our *pack* controller, a temporal modification to the homogeneous controller, while DPST is associated with our *caste* controller, a spatial modification of the homogeneous controller scheme. The DPDT category represents the case where there is little possibility of physical interaction. For example, the robots may occupy non-contiguous regions of space, or only one robot at a time may be activated. Since our focus is on controllers for multiple robots interacting to accomplish a task, the DPDT category does not provide an acceptable solution for interference management.

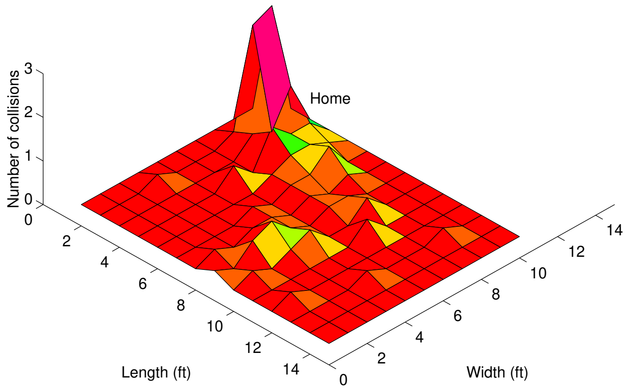

Fig. 6. This plot shows the characteristic interference pattern for the homogeneous implementation of the collection task on the four physical robots. The shading corresponds to the height of the peaks, best seen in color.

Figure 6 presents the characteristic interference pattern for the homogeneous implementation showing the number of collisions between robots within the Corrall. The data for the plot are an average of the collisions observed over five trials with the completion criterion defined as collecting 14 of the 27 pucks at Home. The figure shows high levels of interference near Home resulting from multiple robots simultaneously attempting to deliver pucks. We thus seek to modify the controller in order to reduce this interference using our two spatio-temporal variations, pack and caste. We also present a more detailed comparative evaluation of interference in the Analysis section.

VI. The Pack Controller

In the pack controller, as in the homogeneous version, all individuals have identical behaviors and activation conditions. Unlike the homogeneous controller, however, the robots do not act concurrently and independently. Instead, a dominance hierarchy is imposed, based on some functional criterion such as the robots' different capabilities, or on an arbitrary assignment scheme such as the robot ID, if the robots are functionally identical (as are ours).

The dominance hierarchy induces a temporal structure on the task by allowing only one of the robots to deliver a puck at any time. All of the robots may search for pucks in parallel, as in the homogeneous implementation, but if two or more robots simultaneously find pucks, the one highest in the hierarchy is allowed to deposit its pucks first. The other robot(s) cannot proceed until the first has finished delivering its pucks and has left the Boundary region (Figure 7). This scheme corresponds to SPDT (or temporal) arbitration of SPST interactions.

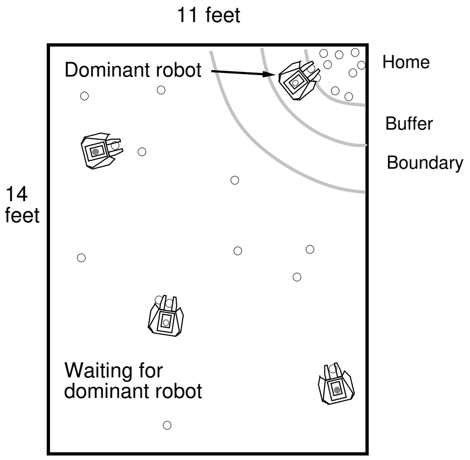

Fig. 7. The pack variation of the collection task.

The pack strategy requires that some form of dominance hierarchy can be assigned and that dominance rank can be recognized between the robots. In our case, rank was communicated over the radios, but in other implementations it could be based on physical characteristics that can be sensed directly.

A. The "message passing" Behavior

Figure 8 presents the controller for the pack implementation. This controller is almost identical to the homogeneous controller (Figure 5), except that it includes a high-precedence *message passing* behavior. The function of *message passing* is to send the robot's status, specifically whether it is delivering a puck, to the other robots, and in turn monitor the status of the other robots. When a robot finds a puck, *message passing* places the robot into a wait state with the motors off and enters the following communications routine:

1. Wait two communication cycles (approximately 6 seconds) to accumulate the most current status information from each robot.
2. If after (1) above, no other robot is currently delivering a puck, transmit the desire to do so. Otherwise return to (1).
3. Wait three communication cycles (approximately 10 seconds) for synchronization with the other robots.
4. If after (3) above, no other robot wishes to deliver a puck, or any that do are less dominant, then proceed to deliver the puck and inform the other robots when finished. Otherwise, return to (1).

B. Robustness

As we have discussed, it is important that multi-robot controllers be robust to noise and robot failures. Similar to the homogeneous controller, the pack controller accommodates robot failures by having each robot able to accomplish the entire task. Unlike the homogeneous controller, the coordinated hierarchy of the pack controller requires special measures by the *message passing* behavior to ensure robustness. If a robot fails while searching for a puck, no special measures are required since no other robot is waiting upon its actions. If, however, the robot fails while delivering a puck, the other robots must be informed so as not to wait indefinitely. The failed robot can send such a message if it is able to detect the failure (a difficult problem in itself). Otherwise, some external agent, such as a human operator, can send the message.

We use a somewhat different approach in our experiments. Whenever a robot fails, it is shut down and restarted by a human operator. (In hazardous conditions, it could be possible to restart the robots remotely.) During this restart period, the other robots receive no communications from the failed robot. The robots consider such periods of protracted radio silence as an indication of the failure of the robot, and once again enter into the communications routine above. Once the failed robot has restarted and begins communicating, it is seamlessly incorporated back into the hierarchy. Since the communications routine only uses relative dominance to decide which robot should deliver a puck, it easily accommodates the attrition or addition of robots.

Another advantage of our communications routine above is that the use of "radio silence" failure detection helps provide group-level robustness to radio noise. As noise levels increase, communication between the robots becomes increasingly difficult. This may lead to protracted periods of radio silence that are incorrectly interpreted as robot failures. In such a situation, two or more robots may deliver pucks at the same time. The degradation of the hierarchy, however, is what prevents the failure of the entire group. Even if the radio system were to fail completely, the task would still be accomplished because every robot would consider every other robot as having failed. Thus, the pack controller would degenerate into the homogeneous controller. We posit that such graceful degradation in group structure, without jeopardizing the entire task, is an important property of controllers for unknown and dynamic environments.

C. Interference

Figure 9 shows the characteristic interference pattern for the pack controller, averaged over 5 trials. The completion criterion was identical to the homogeneous case: delivering 14 of the 27 pucks to Home. As is clear from a comparison to the characteristic interference of the homogeneous controller (Figure 6), the pack controller has reduced interference near Home, as desired.

Not only is the pack controller successful in reducing interference, it is also attractive in its ease of implementation. The pack variation is simply the homogeneous controller with the addition of the dominance hierarchy induced by the *message passing* behavior. Such ease of implementation supports our requirement that controllers be easy to

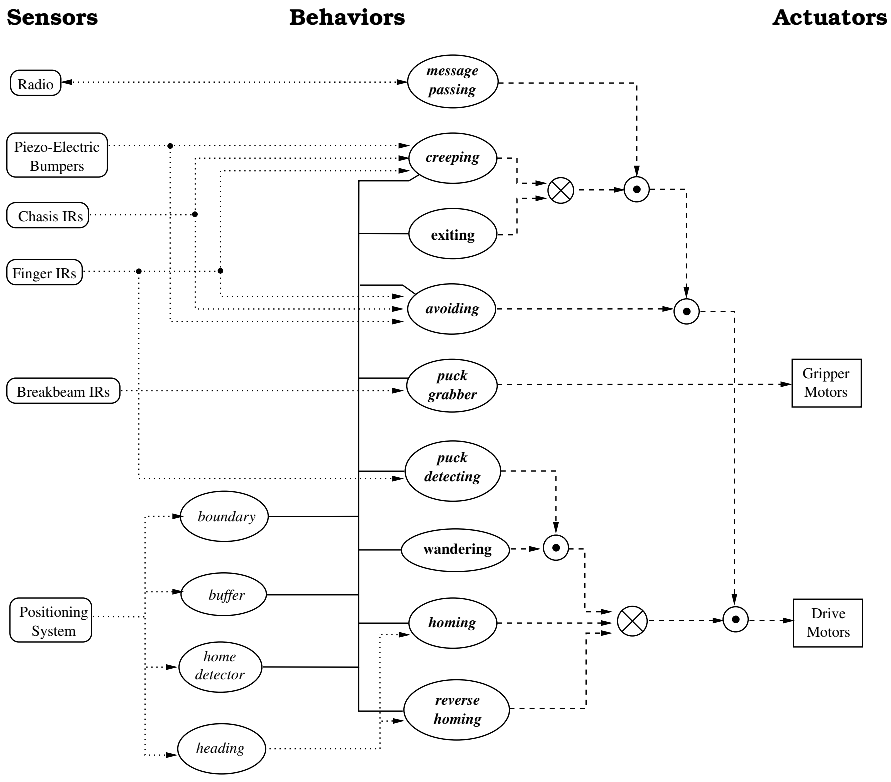

Fig. 8. The pack version of the controller for the collection task.

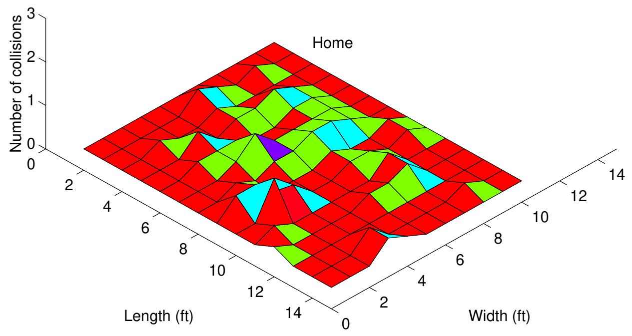

Fig. 9. This plot shows the characteristic interference pattern for the pack implementation of the collection task on the four physical robots. The shading corresponds to the height of the peaks.

modify.

VII. The Caste Controller

In a caste controller, the group of robots differentiates into two or more sub-groups (castes), each of which acts concurrently and independently, but occupies different regions of the task space. The goal is to manipulate interference by an appropriate division of the task space, and assignment of the castes to the sub-regions. This spatial separation of castes limits physical interactions to the territorial boundaries. The caste scheme thus corresponds to DPST arbitration of SPST interference.

Unlike the homogeneous and pack strategies, the subgroups of robots in the caste strategy may have different behavioral repetoirs. Indeed, that is the case with the caste implementation we present in this section. It consists of two sub-groups: the "Search Caste," comprised of three robots which find pucks and bring them near Home, and the "Goal Caste," comprised of one robot which brings the pucks the rest of the way to Home (Figure 10). Each of the two castes has a different controller.

A. The Search Caste

In our implementation, three of the four R2e robots, comprising the "Search Caste," have behavior sets similar to the homogeneous implementation. Each robot searches the region S for pucks, but delivers them to the line separating the Boundary and Buffer zones, rather than all the way to Home. Figure 11 presents the controller for the Search Caste. It is identical to the homogeneous controller (Figure 5), except that it lacks the *creeping* behavior. This more refined combination of *homing* and *avoiding*, designed to bring the robots to the corner of the Corrall, is no longer necessary since pucks are not brought to the corner. The *buffer* behavior is also removed from the controller because it is not needed to activate *creeping*.

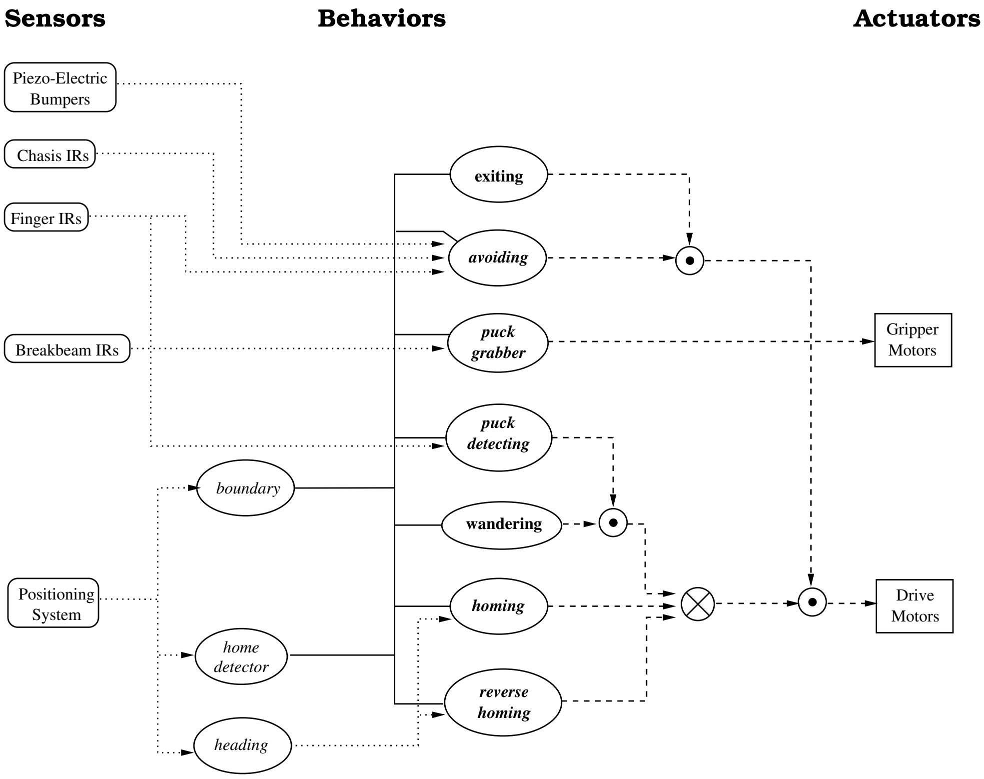

Fig. 11. The controller for the Search Caste, the three robot subgroup that searches for pucks.

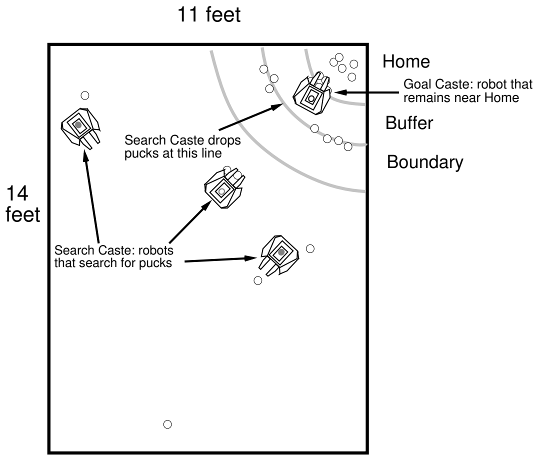

Fig. 10. The caste variation of the collection task.

B. The Goal Caste

The "Goal Caste" consists of one robot that remains in the Home and Buffer regions with the task of transporting to Home the pucks dropped by the Search Caste at the Boundary/Buffer line. The controller for the Goal Caste is presented in Figure 13. The *sweeping* behavior moves the robot away from Home and performs an arc at the Boundary/Buffer line to "scoop up" any pucks left there (Figure 12). The *creeping* behavior then carefully moves the robot to Home, where it performs *exiting* to back up and deliver the pucks. The robot then turns in place 180 degrees to once again begin *sweeping*. During the execution of the controller, the gripper remains lifted allowing the concave front region of the robot's base to scoop up multiple pucks.

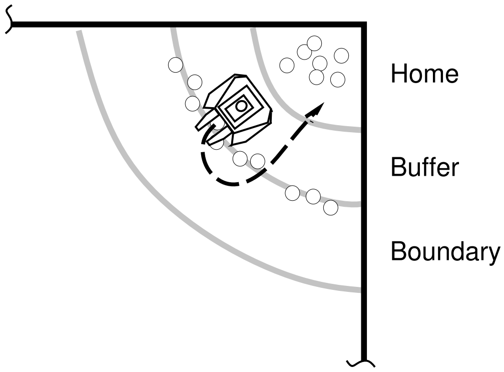

Fig. 12. The *sweeping* behavior of the controller for the Goal Caste.

C. Robustness and Interference

The controller for the Search Caste shares many of the characteristics of the homogeneous controller. It achieves group-level robustness by maintaining a behaviorally iden-

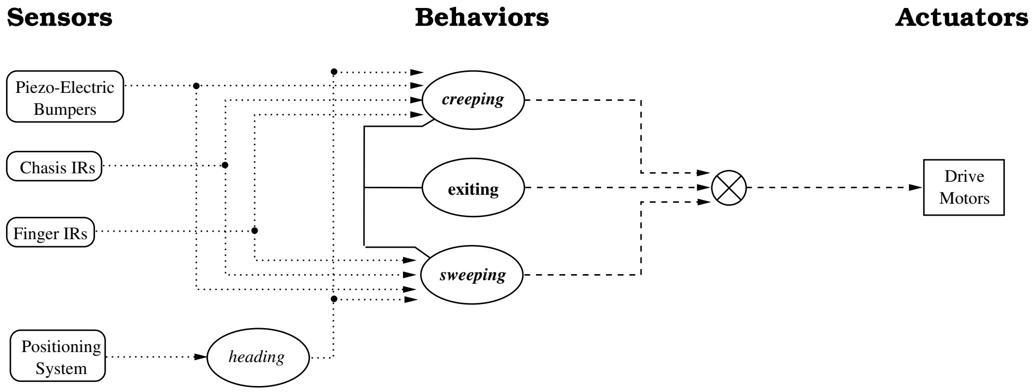

Fig. 13. The controller for the Goal Caste, the one robot subgroup that brings pucks from the Boundary/Buffer line to Home.

tical group with no reliance on explicit communication. Thus, neither high levels of noise nor the failure of a robot debilitates the entire caste. The Search Caste controller also provides a good example of the ease with which the homogeneous controller can be modified.

One of the keys to robustness in the caste controller is the asynchronicity of interaction between the two castes. The Search Caste must deliver pucks to the Boundary/Buffer line, but the Goal Caste is not dependent upon them arriving at a particular time or in a particular order that may be difficult to ensure in such a complex, stochastic system.

Though not implemented in our caste controller, additional robustness could be added by using a variation of the pack communication protocol to transmit the number of active members of each caste. If one caste were to lose too many individuals, members of the other castes could replace them. For example, if the one robot of the Goal Caste were to fail, a member of the Search Caste could substitute. This scheme, while improving robustness, would require each robot to possess all of the individual caste controllers and be able to switch between them as necessary. Such caste switching would be most robust if duplication of the state of the failed robot were not necessary, as would be the case with our controller.

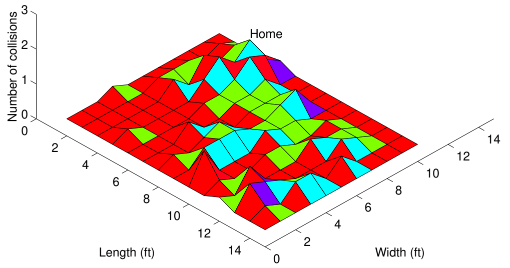

Fig. 14. This plot shows the characteristic interference pattern for the caste implementation of the collection task on the four physical robots. The shading corresponds to the height of the peaks.

Figure 14 shows the average characteristic interference over five trials for the caste implementation. The completion criterion was the same as for the homogeneous and pack controllers: 14 of the 27 pucks collected. It is clear from a comparison to the characteristic interference of the homogeneous controller (Figure 6) that interference near Home is reduced, as was desired. The overall level of physical interference throughout the Corrall, however, is higher with the caste controller.

The following section provides a more detailed quantitative evaluation and comparison of the controllers in terms of interference, as well as time-to-completion and the distance traveled by each robot.

VIII. Analysis

In order to better evaluate the desirability of and tradeoffs between the three controllers — homogeneous, caste, pack — we performed five experimental trials for each of the controllers, gathering both spatial and temporal data. Initial conditions for all trial were as nearly identical as possible in order to minimize free variables, and the completion criterion was 14 out of 27 pucks collected. For each trial, we gathered data on the time-to-completion of the task, and the location and number of collisions between robots, showing the characteristic interference. We calculated the average total number of collisions for each experiment, for relative comparison of the different schemes. Using the positioning system, we also recorded each robot's location at a approximately 0.3 Hz in order to examine the distance traveled and path taken by each. Finally, we monitored the activity of the internal behaviors of the robots. The *avoiding* behavior was of particular interest since it is one directly invoked by physical interference. We hypothesized that time spent avoiding should be correlated with the total amount of interference in each of the implementations, and would thus serve as an alternate measure of interference. As shown below, this hypothesis was validated (see Table II).

All of the data presented in this section have been analyzed with one or more statistical tests. We have performed hypothesis tests using Student's t statistic, 1-factor analysis of variance (ANOVA), and 2-factor ANOVA, in order to

Controller	Time (sec)	Avoiding (sec)
Homogeneous	549	143
Caste	1447	442
Pack	1081	229

TABLE I
AVERAGE TIME OF TASK COMPLETION AND AVERAGE TIME SPENT IN THE *avoiding* BEHAVIOR FOR EACH CONTROLLER.

Controller	Interference (collisions)	Avoiding/Time
Homogeneous	16.4	0.27
Caste	20	0.32
Pack	12.6	0.22

TABLE II
AVERAGE AMOUNT OF INTERFERENCE AND AVERAGE FRACTION OF TIME SPENT IN THE *avoiding* BEHAVIOR.

Controller	Interference/Time (collisions/sec)
Homogeneous	0.030
Caste	0.014
Pack	0.012

TABLE III
AVERAGE AMOUNT OF INTERFERENCE PER UNIT TIME FOR EACH CONTROLLER.

verify that the differences between the results of the implementations were in fact statistically significant. In all cases, these differences were significant with p-values < 0.05.

Our discussion in this section is based on the assumption that the task environment is fixed. Another effective method for altering the spatio-temporal properties considered below is modification of the environment, if such is possible. We could, for example, move Home to the middle of the workspace, thus manipulating the properties like interference and time-to-completion.

A. Interference, Avoiding, and Time

One factor that impacts the total amount of interference observed for each implementation is the time-to-completion of the collection task. One would expect that for any given implementation, the longer the trial continues, the more interference or collisions there would be. One would also expect the total amount of time spent in the *avoiding* behavior to be positively correlated with the time of completion. In Table I we see that this is indeed the case. The homogeneous implementation has the shortest time-to-completion and the least amount of time spent avoiding; the pack implementation has the next larger times; and the caste implementation has the largest times over all.

In their current form, the values for time-to-completion and time-spent-avoiding do not provide much useful information about the amount of interference in each controller. We can observe, however, that the amount of time spent in the *avoiding* behavior is composed of the time spent avoiding other robots (before, during, and after collisions) and the time spent avoiding everything else. Since the environment (discounting the robots) is identical in every trial, we can assume that the amount of avoidance per unit time attributable to non-robot objects is constant between the implementations. This assumption suggests that any differences in the amount of avoidance per unit time between the implementations are primarily due to the avoidance of the other robots, possibly during collisions.

Thus we would expect to see a correlation between the average amount of interference observed in each implementation and the ratio of time spent avoiding to total time. In Table II we observe that such a correlation does exist and it is quite large at $\rho = 0.995$. This indicates an important link between these two values and suggests that the ratio of avoiding and total time is a good estimate of relative average interference levels.

Another potentially useful statistic is the amount of interference per unit time. As shown in Table III, the pack implementation has the most desirable ratio while the homogeneous implementation has the least.

B. Distance Traveled

As mentioned previously, the energy expended by the robots in completing the task may be a concern if recharging is time-consuming or difficult. Time-to-completion provides one approximation of energy expenditure, but it can be inaccurate, especially with a controller such as our pack version where robots can idle for long periods of time. A better approximation is the amount of work accomplished by the robots during the task. Work (W), force (F), and displacement (d) are related through the elementary physics equation

$$W = F \cdot d \cdot \cos\theta$$

or

$$\frac{W}{F \cdot \cos\theta} = d,$$

where θ is the angle between the force and displacement vectors. Since the robots are mechanically identical, we can consider $F \cdot \cos\theta$ to be constant among them. This allows us to compare the work done by the robots solely in terms of d, the distance traveled. Because the robots are identical, d also provides a reasonable, relative indication of the energy expended in performing the work. Finally, it provides a measure of efficiency: the less work required to accomplish the task, the more efficient the controller.

Table IV presents the average distance traveled by each robot, and the total over all robots, for each of the three controllers. The values were calculated from the robot position data gathered during the experiments. The results indicate that the homogeneous controller performs the least work in completing the task, and thus is the most efficient,

Controller	Robot0	Robot1	Robot2	Robot3	Total (ft)
Homogeneous	123	120	113	119	475
Caste	353	370	385	119	1227
Pack	112	145	188	178	623

TABLE IV
Average distance (in feet) traveled by the robots for each controller.

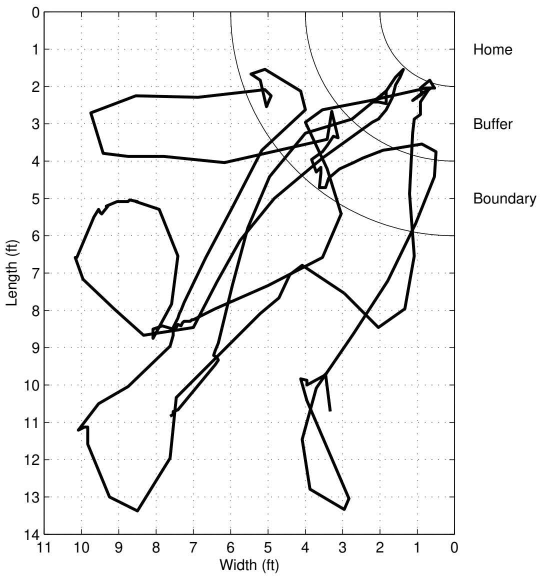

Fig. 15. A typical path taken by one physical robot in the homogeneous controller.

whereas the caste controller performs the most work and is least efficient.

Although the total distances traveled for the three controllers are statistically different, this is not necessarily true of the distances traveled by the individual robots within a controller. This follows intuitively from the structure of the controllers. In the homogeneous controller where all four robots are behaviorally identical, there is no statistical difference in the distances traveled. In the caste controller, Robot0, Robot1, and Robot2, which comprise the Search Caste, travel similar distances, whereas Robot3 of Goal Caste moves significantly less, as might be expected. In the pack controller, one would expect the less dominant robots to travel less since they spend more time waiting for the dominant robots to deliver pucks. Table IV, with Robot0 as the least dominant and Robot4 as the most dominant, shows that this is the general trend. Although a one-way analysis of variance indicates that there is significant difference among these values, there are too few trials to provide further discrimination using a t-test. (The exception is that Robot0 is shown to travel significantly less than Robot2 and Robot3.)

A more qualitative, visual examination of the execution of the controllers is also possible. Figure 15 shows a typical path of one robot in the homogeneous controller. It is clear that the robot searches for pucks, delivers several to Home, and sometimes enters the Boundary without pucks and promptly leaves. Figure 16 (Left) shows a similar path taken by a robot in the Search Caste of the caste controller. The path is much longer than that of the homogeneous controller due to the protracted time of the trial. We also note that the Search Caste very clearly delivers pucks to the Boundary/Buffer line. Figure 16 (Right) shows the complementary path of the Goal Caste collecting pucks from the Boundary/Buffer line and taking them to Home. Figure 17 provides a juxtaposition of typical paths taken by the least dominant and most dominant robot of the pack controller. As expected, the most dominant robot has a path (Right) very similar to that of the homogeneous controller. The least dominant robot, however, has a severely stunted path demonstrating the significance of the time it waits for the more dominant robots to deliver their pucks.

C. Robustness

During the experimental trials for each controller, we had the opportunity to evaluate group-level robustness. The R2e robots used in the experiments are quite fragile and prone to failure from something as simple as a buildup of static electricity corrupting memory or causing the robot's computer to crash. There was seldom a trial without multiple failures requiring the failed robots to be restarted. With the homogeneous controller, we noted very clearly that the failure of one robot did not effect the activity of the others. In the pack controller, the less dominant robots of the hierarchy were able to compensate for the failure a dominant robot by using the message passing protocol. If a dominant robot failed while delivering a puck (which occurred at least once per trial), the less dominant robots would stop waiting and begin delivering their pucks. In the caste controller, the Search Caste exhibited group-level robustness similar to the homogeneous controller: the failure of one robot did not affect the other members of the caste. In addition, due to the asynchronicity of interaction between the two castes, the failure of the robot in the Goal Caste did not debilitate the activity of the Search Caste.

D. Evaluation

Using the analyses presented above we can now discuss the relative desirability of the three controllers. All three are desirable in that they exhibit good group-level robustness. The tradeoff between time and interference captures the relative performance. The homogeneous implementation requires the least time but does not result in the least

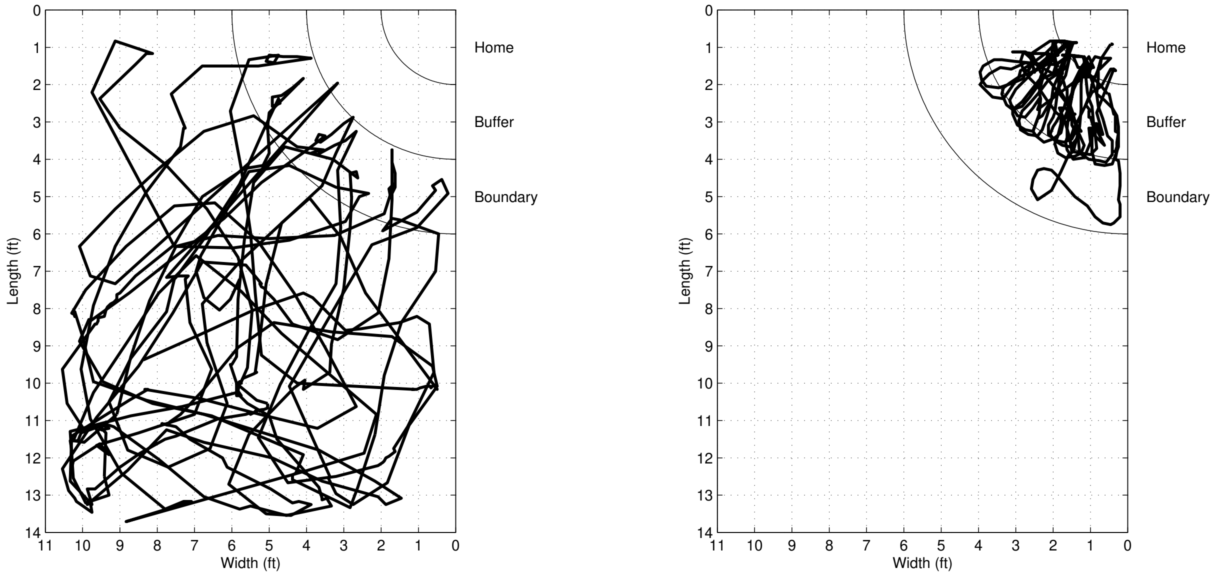

Fig. 16. (Left) A typical path of a physical robot in the Search Caste of the caste controller; (Right) a typical path of the robot in the Goal Caste.

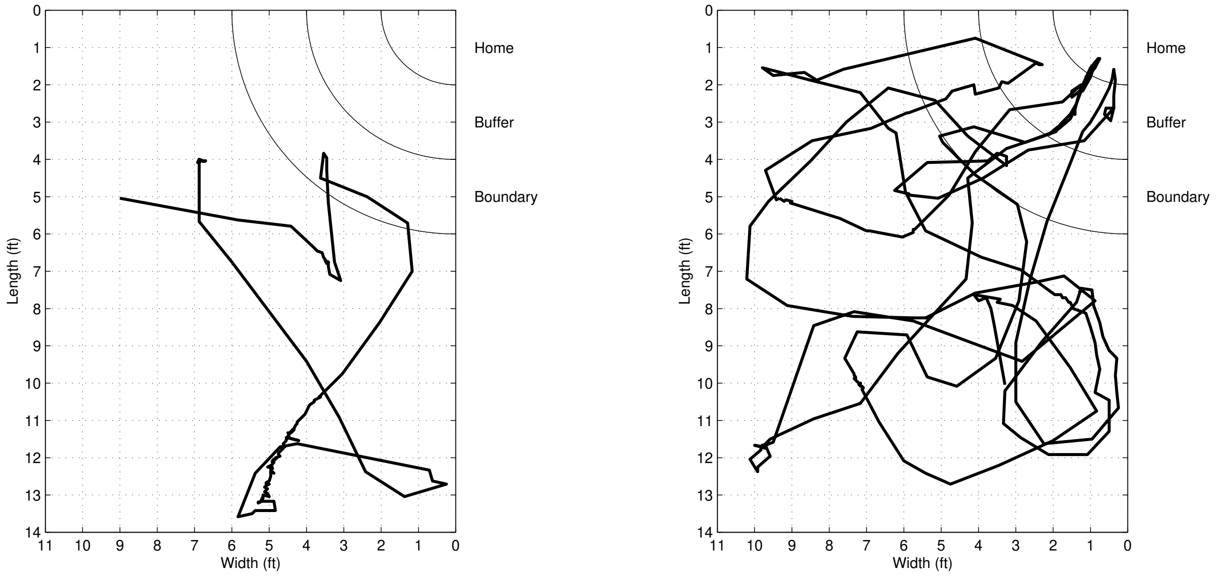

Fig. 17. (Left) A typical path of the least dominant robot of the pack controller; (Right) a typical path of the most dominant robot.

interference, whereas the pack implementation exhibits the least total interference and least interference per unit time, but takes longer overall. Thus, we must decide which criterion is more important or what kind of compromise we wish to make in the final controller choice. If we can sacrifice some performance time for decreased robot interference, then the pack implementation appears to be the best choice. This solution applies to conservative systems where collisions and the possibility of equipment damage outweighs the required time. In contrast, if total time or energy expenditure is the critical factor, such as in domains where the items to be collected are toxic or dangerous, or robot power is limited, then the homogeneous implementation is the better choice. From this analysis we also observe that the caste implementation does not appear to be a satisfactory solution under any criterion, and may be discarded from consideration.

Although our analysis does not identify one controller that is clearly superior in all respects, it does provide information to make an intelligent decision regarding the tradeoffs between the homogeneous and pack controllers. The designer may decide that one of the controllers sufficiently satisfies the requirements for the task, or might wish to investigate other variations for a more suitable controller. The latter decision is facilitated by the ability to build behavior-based controllers that are easy to modify and evaluate in an expeditious manner, as we have demonstrated here.

IX. Conclusions

We have demonstrated the successful application of behavior-based control to the task of distributed multi-robot collection. Our focus has been on developing controllers that are robust to noise and robot failures, and easily modified to facilitate development of the variation that sufficiently satisfies the requirements for the task. Three versions of the collection task were presented: an initial homogeneous controller, and two variations (pack and caste) derived from the spatio-temporal manipulation of physical interference. All three versions were evaluated in a spatio-temporal context using interference, time-to-completion, and distance traveled as the main diagnostic parameters. This work demonstrates that given a good substrate for development (e.g., a useful set of behaviors), it can be relatively easy to implement, evaluate, and compare multi-robot controllers.

Acknowledgments

The research reported here was conducted at the Interaction Laboratory at the Brandeis University Volen Center for Complex Systems and the Computer Science Department, and at the University of Southern California Robotics Research Labs and Computer Science Department. The work was supported in part by the Office of Naval Research Grant N00014-95-1-0759, and by the National Science Foundation Infrastructure Grant CDA-9512448, and by the Office of Naval Research Grant N00014-00-1-0140, and by the Sandia National Labs Grant 3065.

References

[1] Y. Uny Cao, Alex S. Fukunaga, and Andrew B. Kahng, "Cooperative mobile robotics: Antecedents and directions," *Autonomous Robots*, vol. 4, pp. 1–23, 1997.

[2] 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, 1995.

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

[4] 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*. 1993, pp. 486–493, MIT Press.

[5] 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*. 1994, pp. 473–478, MIT Press.

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

[7] Dani Goldberg and Maja J Matarić, "Coordinating mobile robot group behavior using a model of interaction dynamics," in *Proceedings, The Third International Conference on Autonomous Agents (Agents '99)*, Seattle, Washington, May 1–5 1999, ACM Press.

[8] Miguel Schneider Fontán and Maja J Matarić, "Territorial multi-robot task division," *IEEE Transactions on Robotics and Automation*, vol. 14, no. 5, October 1998.

[9] Steven Y. Goldsmith, John T. Feddema, and Rush D. Robinett, "Analysis of decentralized variable structure control for collective search by mobile robots," in *Sensor Fusion and Decentralized Control in Robotic Systems*, Boston, Massachusetts, 1998, vol. 3523 of *SPIE Proceedings*, pp. 40–47, SPIE.

[10] 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, MIT Press.

[11] Lynne E. Parker, *Heterogeneous Multi-Robot Cooperation*, Ph.D. thesis, MIT, 1994.

[12] K.-H. Tan and M. Anthony Lewis, "Virtual structures for high-precision cooperative robotic control," in *IEEE/TSJ International Conference on Intelligent Robots and Systems*. 1996, IEEE Computer Society Press.

[13] R. Beckers, O.E. Holland, and J.L. 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*. 1994, pp. 181–189, MIT Press.

[14] Owen Holland and Chris Melhuish, "Stigmergy, self-organization, and sorting in collective robotics," *Artificial Life*, vol. 5, no. 2, 2000.

[15] Alexis Drogoul and Jacques Ferber, "From tom thumb to the dockers: Some experiments with foraging robots," in *From Animals to Animats II*. 1992, pp. 451–459, The MIT Press: Cambridge, Massachusetts.

[16] Barry Brian Werger and Maja J Matarić, "Robotic "food" chains: Externalization of state and program for minimal-agent foraging," in *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*. 1996, pp. 625–634, The MIT Press: Cambridge, Massachusetts.

[17] Richard T. Vaughan, Kasper Støy, Gaurav S. Sukhatme, and Maja J Matarić, "Whistling in the dark: Cooperative trail following in uncertain localization space," in *Proceedings of the Fourth International Conference on Autonomous Agents*. 2000, pp. 187–194, ACM Press: New York.

[18] Henrik Hautop Lund and Luigi Pagliarini, "Robocup jr. with lego mindstorms," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000, IEEE.

[19] Mattias Lindström, Anders Orebäck, and Henrik I. Christensen, "Berra: A research architecture for service robots," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000, IEEE.

[20] Julio Rosenblatt, Stefan Willams, and Hugh Durrant-Whyte, "Behavior-based control for autonomous underwater exploration," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000, IEEE.

[21] Yasuhisa Hasegawa, Yoshikuni Ito, and Toshio Fukuda, "Behavior coordination and its modification on brachiation-type mobile robot," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000, IEEE.

[22] Paolo Pirjanian, *Multiple Objective Action Selection & Behavior Fusion using Voting*, Ph.D. thesis, Institute of Electronic Systems, Alborg University, Denmark, 1998.

[23] Paolo Pirjanian, Henrik I. Christensen, and J. A. Fayman, "Application of voting to fusion of purposive modules: An experimental investigation," *Robotics and Automation*, vol. 23, pp. 253–266, 1998.

[24] Paolo Pirjanian and Maja J Matarić, "Multi-robot target acquisition using multiple objective behavior coordination," in *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, San Francisco, CA, April 2000, IEEE.

[25] Dani Goldberg and Maja J Matarić, "Augmented markov models," Tech. Rep. IRIS-99-367, University of Southern California, 1999.

[26] Dani Goldberg and Maja J Matarić, "Detecting regime changes with a mobile robot using multiple models," Tech. Rep. IRIS-00-382, USC Institute for Robotics and Intelligent Systems, 2000.

[27] Dani Goldberg and Maja J Matarić, "Reward maximization in a non-stationary mobile robot environment," in *Proceedings, Agents 2000*, 2000.

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

[29] Rodney A. Brooks, "A robust layered control system for a mobile robot," *IEEE Journal of Robotics and Automation*, vol. RA-2, no. 1, pp. 14–23, April 1986.

[30] Rodney A. Brooks, "Intelligence without reason," in *Proceedings, IJCAI-91*, 1991.

[31] 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, 1992, pp. 46–54.

[32] Ronald C. Arkin, *Behavior-Based Robotics*, The MIT Press: Cambridge, Massachusetts, 1998.