

Interference as a Tool for Designing and Evaluating Multi-Robot Controllers

Dani Goldberg & Maja J Mataric

Volen Center for Complex Systems, Computer Science Department
Brandeis University
Waltham, MA 02254

dani@cs.brandeis.edu & maja@cs.brandeis.edu

Abstract

Designing and implementing cooperative group behaviors for robots is considered something of a black art involving an extensive amount of reprogramming and parameter adjustment. What seems to be lacking is a pragmatic, practical, general-purpose tool that would both guide the design and structure the evaluation of controllers for distributed real-world multi-robot tasks. In this paper, we propose the use of interference between robots as one such simple tool for designing and evaluating multi-robot controllers. We explore how key issues in multi-robot control can be addressed using interference, a directly measurable property of a multi-robot system. We discuss how behavior arbitration schemes, i.e., the choice of controllers, can be made and adjusted using interference. As an experimental example, we demonstrate three different implementations of a collection clean-up (foraging) task using four physical mobile robots, and present analyses of the experimental data gathered from trials of all three implementations.

Introduction

Designing and implementing controllers for multiple interacting mobile robots is considered something of a black art that involves a great deal of reprogramming and parameter adjustment. In this paper, we propose a pragmatic tool for both guiding the design and structuring the evaluation of such controllers. The idea is based on the use of *interference* between robots as a key diagnostic parameter and a directly measurable property of a multi-robot system. We discuss how behavior arbitration schemes, i.e., the choice of controllers, can be made and adjusted using interference. We start by introducing and defining interference, then discussing its impact on arbitration/controller design. Next, we demonstrate three different implementations of a collection clean-up task with four mobile robots and an-

alyze the experimental data gathered from trials of all three implementations. We conclude the paper with an overview of related work, and a summary.

Interference and Arbitration

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.). This paper focuses on physical interference and demonstrates that, while it is often a source of increased complexity in controller design, it is also an effective tool for system design and evaluation. This work is distinguished from much of the research exploring multi-agent spatial interactions in that it is not computational in nature but rather situated in the dynamics of the physical world. Our work is couched in the framework of distributed behavior-based control (Mataric 1992, Brooks 1991) where controllers consist of a collection of behaviors and an arbitration mechanism. However, the ideas about using interference as a tool are not limited to behavior-based systems.

We consider 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. The result is a surface showing interference peaks that can be used to adjust the controller (or the arbitration scheme) in order to lessen that 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 design

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 maximize overall task efficiency. Multi-robot interactions that we are focusing 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 the four forms of interaction: SPST, SPDT, DPST, and DPDT.

Physical interference fits into the SPST category, covering the cases when two or more robots try to occupy the same location at the same time. However, the other three categories are also useful for deriving and fine-tuning arbitration/controller schemes. For each of the categories, we derived and tested a corresponding arbitration scheme. The SPDT category is associated with **pack arbitration**, and the STDP and DTDP categories are associated with **caste arbitration**, as described next.

The Collection Clean-up Task

In order to explore the issues of interference in multi-robot controller design, we used the categorization described above and implemented three different controllers, i.e., behavior arbitration schemes, on a set of four mobile robots. All of the controllers were versions of a collection clean-up task, a prototype for various applications (distributed solutions to de-mining, toxic waste clean-up, terrain mapping, etc.), but each demonstrated different interference effects, as our analysis demonstrates.

The Robots



Figure 1: The Don Group: the four-robot family used in the experiments.

Four IS Robotics R2e robots were used (Figure 1). Each is a holonomic base with two drive motors equipped with a two-fingered gripper. The sensing capabilities include piezo-electric bump sensors around

the base and in the gripper, five infrared (IR) sensors around the body and one on each finger, a color sensor in the gripper, a radio transmitter/receiver for communication and data gathering, and an ultrasound triangulation system for positioning. The robots are programmed in the Behavior Language (Brooks 1990), based on the Subsumption Architecture (Brooks 1986). The experiments are performed in an 11 by 14 foot rectangular enclosure (the Corral) with 27 small metal cylinders (pucks) randomly but evenly distributed throughout, except in the drop-off area, called *Home*, a ninety degree sector of a circle, with a radius of two feet, located a corner of the Corral. The robots are initially distributed throughout the Corral.

Behaviors

Behavior-based systems are organized as collections of concurrent and sequenced behaviors or processes (typically goal-achieving control laws), each of which takes inputs from sensors or other behaviors, and sends outputs to other behaviors or directly to the robot's actuators. We implemented the following behaviors for the clean-up task:

avoiding: avoid any object (detected by IR and bump sensors) deemed to be in the path of the robot.

wandering: move forward and, at random intervals, turn left or right through some random arc.

puck detecting: if avoiding is not active, and if an object is detected by the front IRs, lift up the fingers to determine whether the object is short enough to be a puck. If it is, approach the object and try to place it between the fingers. If unsuccessful, perform **avoiding**.

picking up: when an object is detected between the fingers, grasp it and raise the fingers.

homing: if carrying a puck, move towards the designated goal location.

dropping off: if in the home region, release the object and lower the fingers.

exiting: if in the home region, exit the region.

boundary: avoid this region if not carrying a puck.

buffer: if carrying a puck, carefully move toward home.

Each robot was programmed with the same behavior set (except in one case, see below). The differences in the controllers resulted from the arbitration mechanism that determined behavior activation and termination conditions, as described next.

Homogeneous Implementation

In the homogeneous implementation, all robots have identical behavior sets (described above) and act in parallel and independently. The goal location for the

homing behavior is the Home region (Figure 2). In this implementation multiple robots may simultaneously attempt to drop pucks off at the same place, resulting in high levels of interference.

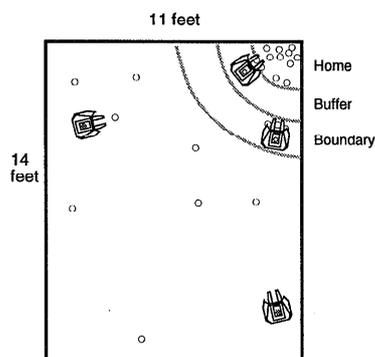


Figure 2: Homogeneous group collection behavior.

Pack Arbitration Implementation

In pack arbitration, as in the homogeneous implementation, all individuals have identical behaviors and activation conditions, but do not act in parallel and independently. Instead, a dominance hierarchy is imposed, based on some functional criterion such as the robots' different capabilities, or on a relatively random scheme, such as robot ID, if the robots are functionally identical as our R2e robots are.

The dominance hierarchy induces a spatial arbitration mechanism by allowing only one of the robots to act at one time. If two or more robots simultaneously pick up pucks, the one highest in the hierarchy is allowed to deposit its pucks first. The other robot(s) cannot proceed until the first has left the Boundary zone (left side of Figure 3). This scheme corresponds to SPDT arbitration.

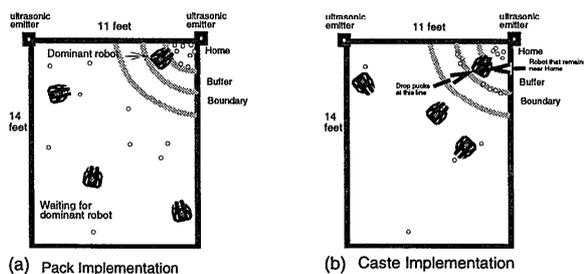


Figure 3: Pack and caste arbitration schemes.

Pack arbitration requires that some form of domi-

nance hierarchy can be assigned and that dominance rank can be recognized between the robots. In our case, rank was communicated by the radios, but in other implementations it could also be based on physical characteristics and sensed directly.

Caste Arbitration Implementation

In caste arbitration, not all of the robots use identical conditions for behavior activation and termination. As a result, the group differentiates into two or more sub-specialties, each of which acts in parallel and independently.

Caste arbitration requires that the global task be divided into several sub-tasks and assigned to different sub-groups (castes). The goal is to minimize interference by an appropriate caste assignment. This is successful if the interaction among individuals of different castes is limited by some form of territorial separation. Thus caste arbitration falls into the DPS'I' and DPDT categories.

In our implementation, three of the four R2e robots have behavior sets identical to the homogeneous implementation. The goal location for **homing** is set to the line separating the Boundary and Buffer regions. The fourth robot collects the pucks within the drop-off area and has the goal location set to Home (right side of Figure 3).

Analysis

We performed five experimental trials with each of the three implementations—homogeneous, caste, and pack. The initial setups for the trials were kept as nearly identical as possible in order to minimize free variables. For each trial, we gathered data on the time of completion of the task, and the location and number of collisions between robots, which shows the characteristic interference. We also calculated the average total number of collisions for each experiment, for relative comparisons of the different arbitration schemes. Finally, we monitored the activity of the internal behaviors of the robots. The *avoiding* behavior was of particular interest since it is the behavior we would expect to be active during highest physical interference between the robots. 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 2).

All of the data presented in this section has 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

Implementation	Time (sec)	Avoid (sec)
Homogeneous	549	143
Caste	1447	442
Pack	1081	229

Table 1: Average time of task completion and average time spent in the avoid behavior for each implementation.

Implementation	Interference	Avoid/Time
Homogeneous	16.4	0.27
Caste	20	0.32
Pack	12.6	0.22

Table 2: Average amount of interference and average amount of avoidance per unit time for each implementation.

ANOVA, in order to 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 .

One factor that impacts the total amount of interference observed for each implementation is the time of completion for 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 that the total amount of time spent in the avoid behavior to be positively correlated with the time of completion. In Table 1 we see that this is indeed the case. The homogeneous implementation has the shortest time of 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 of completion and time spent avoiding do not provide any useful information about the amount of interference in each implementation. We can observe, however, that the amount of time spent in the avoid 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	Interference/Time
Homogeneous	0.030
Caste	0.014
Pack	0.012

Table 3: Average amount of interference per unit time for each implementation.

implementation and the ratio of time spent avoiding to total time. In Table 2 we observe that such a correlation does exist and it is very large at $\rho = 0.995$. This indicates a significant link between these two values and suggests that the ratio avoiding/total-time is a good estimate of relative average interference levels.

Another useful statistic is the amount of interference per unit time. As shown in Table 3, the pack implementation has the most desirable ratio while the homogeneous implementation has the least.

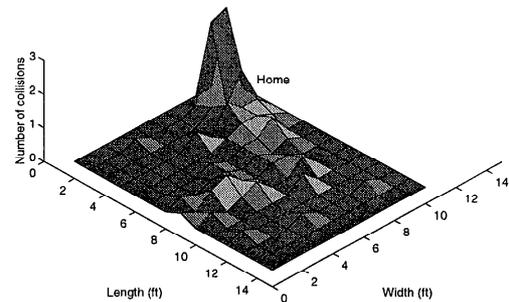


Figure 4: This plot shows the characteristic interference pattern for the homogeneous implementation of the collection task.

A more qualitative look at the interference characteristic of the three implementations is presented in Figures 4, 5, and 6. The homogeneous implementation exhibits a large degree of interference near Home whereas both the caste and pack implementations do not. Though spread out, the interference in the caste implementation is still substantial, while the overall interference in the pack implementation is decreased.

Using the analyses presented above we can now discuss the relative efficiency of the three implementations. The tradeoff between time and interference captures the relative performance. The homogeneous implementation requires the least time but does not result in the least interference, whereas the pack implementation exhibits the least total interference and least interference per unit time but takes longer overall. Thus,

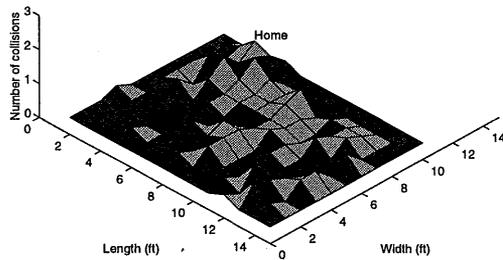


Figure 5: This plot shows the characteristic interference pattern for the caste implementation of the collection task.

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 is the critical factor, such as in domains where robot power is limited or the items to be collected are toxic or dangerous, 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 either criterion.

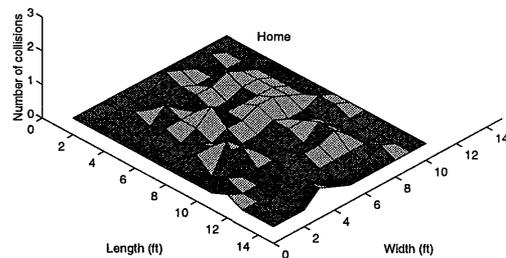


Figure 6: This plot shows the characteristic interference pattern for the pack implementation of the collection task.

The described arbitration mechanisms were based on the assumption that the task environment is fixed. Another effective method for altering interference prop-

erties is to modify the environment, if possible. We could, for example, move Home to the middle of the workspace, thus eliminating locally peaked interference in that region but potentially moving it to puck zones or elsewhere in the area.

Related Work

This section gives only a partial review of some of the related robotics work. Mataric (1994) 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. Arkin, Balch & Nitz (1993) demonstrate simulation work studying the issues of density and critical mass in a hoarding task using fully homogeneous robots. Arkin & Hobbs (1993) describe the general schema-based control architecture, which bears some fundamental similarities to behavior-based control, and give the critical mass experiments. Finally, Arkin & Ali (1994) present a series of simulation results on related spatial tasks such as foraging, grazing, and herding. This work is equivalent to our fully homogeneous initial case. Fontán & Mataric (1996) also worked on multi-robot collection, but focused on issues of critical mass in task division which would correspond to a DPST category. Parker (1992) and Parker (1994) describe multi-robot experiments also on foraging R2e robots with *a priori* hardwired heterogeneous capabilities using the Alliance architecture. Parker (1994) describes a temporal division that sends one robot to survey and measure the environment, and has the rest of the group use its information to clean up the spill in a form of a caste. Cao, Fukunaga, Kahng & Meng (1995) offer a view of the multi-agent and multi-robot research applied to 10 ISR R3 mobile platforms. Tan & Lewis (1996) describes an approach to maintaining a geometric configurations of a robot group using virtual structures, tested on the same homogeneous robot group. Beckers, Holland & Deneubourg (1994) describe 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 is an extreme instance of a homogeneous controller. Donald, Jennings & Rus (1993) and continuing work describes a theoretical framework, based on information invariants, for organizing control of multiple robots. The approach is

successfully demonstrated on physical robots cooperating on tightly-coupled object manipulation tasks. Continuing work by Asada, Uchibe, Noda, Tawaratsumida & Hosoda (1994) deals with multiple special-purpose vision based mobile robots that play a simplified version of soccer using reinforcement learning.

Conclusions

We have presented an approach for evaluating multi-robot controllers and selecting alternatives, based on using interference as a tool. Physical interference, a side-effect of embodied robot interaction, is directly measurable, and a useful metric of system performance. We measured interference both externally, through total task time, and internally, through avoidance time. Three implementations of a collection clean-up task were demonstrated, and interference used to analyze and compare the relative effectiveness of the three arbitration schemes in terms of total time required to complete the task, and in terms of characteristic interference. The work reported here is the beginning of an effort toward a framework for principled task classification and controller design in the domain of distributed physically embodied multi-agent systems.

Acknowledgments The research reported here was done at the Interaction Laboratory at the Brandeis University Volen Center for Complex Systems and the Computer Science Department. The work is supported by the Office of Naval Research Grant N00014-95-1-0759 and the National Science Foundation Infrastructure Grant CDA-9512448.

References

Arkin, R. C. & Ali, K. S. (1994), Reactive and Telerobotic Control in Multi-Agent Systems, in 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', MIT Press, pp. 473-478.

Arkin, R. C. & Hobbs, J. D. (1993), Communication and Social Organization in Multi-Agent Systems, in 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', MIT Press, pp. 486-493.

Arkin, R. C., Balch, T. & Nitz, E. (1993), Communication of Behavioral State in Multi-agent Retrieval Tasks, in 'IEEE International Conference on Robotics and Automation', IEEE Computer Society Press, Los Alamitos, CA, pp. 588-594.

Asada, M., Uchibe, E., Noda, S., Tawaratsumida, S. & Hosoda, K. (1994), Coordination of Multiple Behaviors Acquired by A Vision-Based Reinforcement Learning, in 'Proceedings, IEEE/RSJ/GI Inter-

national Conference on Intelligent Robots and Systems', Munich, Germany.

Beckers, R., Holland, O. E. & Deneubourg, J. L. (1994), From Local Actions to Global Tasks: Stigmergy and Collective Robotics, in R. Brooks & P. Maes, eds, 'Artificial Life IV, Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems', MIT Press.

Brooks, R. A. (1986), 'A Robust Layered Control System for a Mobile Robot', *IEEE Journal of Robotics and Automation* **RA-2**, 14-23.

Brooks, R. A. (1990), The Behavior Language; User's Guide, Technical Report AIM-1227, MIT Artificial Intelligence Lab.

Brooks, R. A. (1991), Intelligence Without Reason, in 'Proceedings, IJCAI-91', Sydney, Australia.

Cao, Y., Fukunaga, A. S., Kahng, A. B. & Meng, F. (1995), Cooperative Mobile Robotics: Antecedents and Directions, in 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Yokohama, Japan.

Donald, B. R., Jennings, J. & Rus, D. (1993), Information Invariants for Cooperating Autonomous Mobile Robots, in 'Proc. International Symposium on Robotics Research', Hidden Valley, PA.

Fontán, M. S. & Matarić, M. J. (1996), A Study of Territoriality: The Role of Critical Mass in Adaptive Task Division, in P. Maes, M. Matarić, J.-A. Meyer, J. Pollack & S. Wilson, eds, 'From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior', MIT Press, pp. 553-561.

Matarić, M. J. (1992), Behavior-Based Systems: Key Properties and Implications, in 'IEEE International Conference on Robotics and Automation, Workshop on Architectures for Intelligent Control Systems', Nice, France, pp. 46-54.

Matarić, M. J. (1994), Interaction and Intelligent Behavior, Technical Report AI-TR-1495, MIT Artificial Intelligence Lab.

Parker, L. E. (1992), Adaptive Action Selection for Cooperative Agent Teams, in 'From Animals to Animats: International Conference on Simulation of Adaptive Behavior', pp. 442-450.

Parker, L. E. (1994), Heterogeneous Multi-Robot Cooperation, PhD thesis, MIT.

Tan, K.-H. & Lewis, M. A. (1996), Virtual Structures for High-Precision Cooperative Mobile Robotic Control, in 'IEEE/TSJ International Conference on Intelligent Robots and Systems', Kyoto, Japan.