# Towards Single-Operator Control of Tightly-Coordinated Robot Teams

Gal A. Kaminka and Yehuda Elmaliach *

The MAVERICK Group
Computer Science Department
Bar Ilan University, Israel
{galk, elmaley}@cs.biu.ac.il

## 1. Introduction

There is growing recognition that many applications of robots will require a human operator to direct multiple robots that collaborate to achieve the operator's goals. However, the bulk of existing work in this area assume that robots are independent of each other and thus key challenges and opportunities are ignored in monitoring and operating tightly-coordinating teams.

This paper takes steps to address these open issues. First, we introduce a graphical display that explicitly shows the coordination in the team, in terms of the robots' state with respect to each other (*socially-attentive monitoring*). As a result, the operator can easily detect coordination failures, even before these cause overall failure in the task. We also take advantage of the robots' teamwork, to allow the robots to actively assist the operator in maintaining her control of the team (*Distributed Call-Request Resolution*). We evaluate *Distributed Call-Request Resolution* techniques in several multi-robot task. Preliminary results show that they can lead to significant improvements in task completion time, and number of coordination failures during execution.

## 2. Socially-Attentive Display

We focus on coordinated tasks–robotic team tasks that require tight coordination between the robots. The hypothesis underlying our investigation is that such tasks require monitoring of the coordination in the team, i.e., explicit monitoring of the team-members' state with respect to each other. Such monitoring is called *socially-attentive monitoring* because it focuses on inter-agent relations, rather than their goals. For example, in formation maintenance tasks, socially-attentive monitoring may include information about the relative positions of the robots within the formation (without reference to where the formation is heading).

A corollary of our hypothesis is that when an operator controls robots in a coordinated task, she will need to infer socially-attentive information if it is not directly available. Unfortunately, existing displays are not necessarily well suited to display socially-attentive information. Instead, they provide information about the current state of each robot, individually. Thus the operator is burdened with inferring the socially-attentive information that is required.

The key to the monitoring approach we advocate is to provide the operator with a socially-attentive display that integrates the raw information coming from each individual robot into a coherent visual display of the social relationship within the monitoring system. Using this display, the operator can easily identify coordination faults (if there are any) within the monitored team, with little or no need for inferring this information from the other displays. This eases the cognitive load on the operator in coordinated tasks. The socially-attentive display must complement the monitoring display associated with the task.

Towards this end, we developed a display called *the relation tool*. The relation tool is a graphical display of multiple robot states on a two-dimensional (2D) plane. Colored dots denote different robots. The positions of the dots denote their states, and thus the shape they make up—their relative positioning—denotes their relative states. In principle, a target relative positioning of the dots must be defined for each application, which signifies correct coordination. Every application requires its own method of projecting robot state onto a 2D plane, and a target shape that defines normative coordination. The key is that the operator should be able to see, at a glance, whether the shape being maintained corresponds to correct coordinated execution.

The relation tool is simple and effective. It has three main advantages. First, it significantly reduces the amount of inference needed by the operator to infer the relative state of

* This research was supported in part by BSF grant #2002401, and by MOD grant 2935762101


AAMAS'04, July 19-23, 2004, New York, New York, USA.

robots—and thus the state of coordination between them. Second, it is not limited to reproducing a global view of the robots, but instead its dimensions can be used to directly provide the operator with information about what is going wrong (in case of failure). Third, it can easily complement other types of displays useful for the task, such as any that show the heading or distance left to the destination.

## 3. Distributed Call-Request Resolution

Previous approaches to human control of multiple robots treat the operator's attention as a centralized resource, which is time-shared between the robots[2, 3, 1, 4]. Robots that require operator's assistance initiate or are issued *call-requests*, which are queued for the operator. The operator switches control between robots, and uses single-robot teleoperation with individual robots to resolve the call requests in some (prioritized) sequence. This method works well in settings where the task of each robot is independent of its peers, and thus the resolution of call requests can be done in sequence, independently of other call-requests.

Centralized resolution of call requests, by the operator, may work well when robots' tasks are independent of each other. However, in coordinated tasks, many robots may have to stop their task execution until a call request is resolved, because their own task execution depends on that of the robot that requires the resolution. In such cases, it is critical to minimize the time it takes to resolve a call request.

We propose a distributed control approach, whereby the robots who depend on the resolution of the call-request take active steps to resolve it, in collaboration with the operator. This approach takes advantage of the robot teamwork, by turning the resolution of the call-request into a distributed collaborative task for all involved. Moreover, the active robots (that do not require assistance) are involved in a coordinated effort with the robot requiring assistance, and thus may be in a better position to assist it.

We investigate distributed resolution in repairing broken formations. Three robots were to be led in a triangular formation. However, we disabled one of the robots (to simulate a catastrophic failure), not letting it move or communicate. In accordance with previous approaches, a call-request was issued to determine the whereabouts of the failing robot. The prototypical method to resolve such a situation, based on previous approaches, would be for the operator to halt the operation of all remaining robots in the formation, and then teleoperate the other robots until the position of the failing robot was determined.

Instead, in distributed resolution, all affected robots actively seek the failing robot. If an operator switches control to one of these active robots, the others coordinate with it (for instance, to cover more ground in the search). However, even if the operator is not involved (e.g., because she may be handling other, unrelated, call-requests), the active robots will still seek out the failing robot. When the operator is ready, they would have hopefully found it.

The robots use their knowledge of the robot's role in the formation to attempt to locate it. They first head out directly towards where the robot would have been if it simply lagged too far behind in the formation. If they fail to find it there, they begin a general search pattern (spiral) that is guaranteed to find the robot, but may take relatively long time.

The key objective in this techniques is too speed up call-request resolution by distributing it, and by using organizational knowledge in the team. Although we tested this technique with triangular formation, the same principle (search at the position where robot should have been, before falling back to a general–but slow–search) can easily transfer to other formations and other spatial tasks.

## 4. Evaluation

To evaluate the techniques we presented, we tested their performance with novice human operators, and contrasted them with other approaches. To evaluate the socially-attentive display we tested it in a task where the operator leads three robots in triangle formation, and in a cooperative box-pushing task. We contrasted the use of socially-attentive display with an approach using the robots' cameras (without a social display). The results support our hypothesis that socially-attentive display can significantly improve monitoring of robots in coordinated tasks.

In a second set of experiments, we contrasted centralized call-request resolution with distributed resolution, in which both robots and operator cooperate. We experimented with several failure scenarios, varying in resolution difficulty, and found that distribution call-request resolution resulted in faster resolution time than both (i) centralized resolution, and (ii) autonomous robotic resolution.

## References


[1] Julie A. Adams. *Human Management of a Hierarchical System for the Control of Multiple Mobile Robots*. PhD thesis, University of Pennsylvania, 1995.

[2] MaryAnne Fields. Modeling the human/robot interaction in onesaf. In *Proceedings of the 23rd Army Science Conference*, 2002. (Poster).

[3] Terrence Fong, Charles Thorpe, and Charles Baur. Multi-robot remote driving with collaborative control. *IEEE Transactions on Industrial Electronics*, 50(4):699–704, August 2003.

[4] Ashlley D. Tews, Maja J. Mataric, and Gaurav S. Sukhatme. A scalable approach to human-robot interaction. In *ICRA-03*, 2003.