Multi-Robot Target Acquisition using Multiple Objective Behavior Coordination

Paolo Pirjanian and Maja Matarić
USC Robotics Research Laboratory
University of Southern California
Los Angeles, CA 90089-0781
phone: +1 (213) 740-4523, fax: +1 (213) 740-7512
{paolo, mataric}@usc.edu

This work is funded by the Office of Naval Research Grant N0014-99-1-0162.

Abstract

In this paper we propose an approach to multi-robot coordination in the context of cooperative target acquisition. The approach is based on multiple objective behavior coordination extended to multiple robots. It provides mechanisms for distributed command fusion across a group of robots to pursue multiple goals of multiple robots in parallel. The mechanisms enable each robot to select actions that not only benefit itself but also benefit the group as a whole. Experimental results with two mobile robots validate that, by using this method, a group of robots can successfully track and acquire a moving target.

1 Introduction

Cooperation of a team of robots in unknown settings poses complex control problems which require solutions that guarantee a suitable trade-off between a multitude of (potentially) conflicting task objectives within and among the robots. For instance, an action that is optimal with respect to a particular robot might be unacceptable with respect to the others. Thus a major issue in the design of systems for control and coordination of multiple robots is the problem of deciding what action to select next. Basically, each robot should make a decision that not only benefits itself but also the group as a whole. Centralized approaches to this problem are not attractive due to their inherent scaling limitations, required communication and communication band-width, global knowledge of the environment etc. Distributed approaches on the other hand are more appealing due to better properties of scaling and reliability. For an overview of approaches and issues in cooperative mobile robotics see [1][3][6]. Related work on multi-robot target observation has been done by [7] using the ALLIANCE architecture, where action selection consists of inhibition (through motivational behaviors). Other approaches to similar tasks include [2] which employs a formation control using a referenced-based (e.g., leader, center, or neighbor) type of cooperation with no explicit task division.

In this paper we propose an approach to multi robot coordination that is based on Multiple Objective Behavior Coordination (MOBC) [8]. As opposed to ALLIANCE, which can be viewed as an extension of subsumptive-style arbitration to multi-robot coordination, the approach proposed here enables *command fusion* across a group of robots. This thus allows for pursuit and satisfaction of multiple goals of multiple robots in parallel. This approach proposes command fusion approaches that find suitable trade-offs between multiple conflicting robot and behavioral objectives. In the past, we have successfully applied this approach to coordination of behaviors on a single robot. In this paper, we extend the approach to multi-robot tasks. In particular, we apply the approach for the task of cooperative target acquisition, where a group of robots collectively find, surround, and capture a moving target.

2 Behavior coordination

Multiple objective behavior coordination is based on multiple objective decision theory [4], which provides approaches for making decisions in complex situations where more than one decision objective should be considered. Objectives define the standards against which the quality of a given alternative may be evaluated. An objective function, o, is thus a mapping that assigns to each alternative, $x \in X$, a value reflecting that alternative's desirability.

In multiple objective behavior coordination, a behavior, b, is formalized as an objective function that defines a mapping from an action space, X, to the interval $[0, 1]$: $b : X \rightarrow [0, 1]$. The action space, $X = \{\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_n\}$, is defined to be a finite set of alternative actions[1]

2.1 Multi-objective decisions

Once the mapping between behaviors and objective functions is established, we can formulate the coordination of multiple behaviors as a multiple objective decision problem that seeks to *optimize multiple objective functions simultaneously*, where the objective functions are produced by multivalued behaviors. From the set of alternative actions X, an action $\hat{\mathbf{x}} \in X$ should be chosen that best satisfies all objective functions. This is formulated as a multiple objective decision problem in the following way:

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{argmax} \qquad [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})] \qquad \text{subject to} \qquad \mathbf{x} \in X, \tag{1}$$

where $o_i(\mathbf{x})$ for $i = 1, ..., n$ are the objective functions estimated by corresponding behaviors b_i and $\mathbf{x}$ is an N-dimensional control vector $(x_1, x_2, ..., x_N)$. Further $X \subseteq R^N$ is the constrained action space[2]. The solution $\hat{\mathbf{x}}$ to equation (1) corresponds to the selected action that best satisfies (the objectives of) the corresponding behaviors.

2.2 Solving multi-objective problems

There are a number of methods described in the literature for solving multi-objective decision problems, all with certain advantages and disadvantages [8, 4]. Rather than merely giving a description of these methods, we highlight the main characteristics of multiple objective decision making (MODM) that make it suitable for behavior coordination. We start by describing the types of solutions to a MODM problem.

A solution is *optimal* if it results in the maximum value of each of the objective functions simultaneously. That is, $\hat{\mathbf{x}}$ is an optimal solution to equation (1) iff $\mathbf{o}(\hat{\mathbf{x}}) \geq \mathbf{o}(\mathbf{x})$ for all $\mathbf{x} \in X$, where $\mathbf{o}(\mathbf{x}) = [o_1(\mathbf{x}), o_2(\mathbf{x}), ..., o_n(\mathbf{x})]$.

An important characteristic of MODM problems is that the multiple objectives, $o_1, ..., o_n$, can be competitive and conflicting, i.e., the improvement in one can be associated with a deterioration in another. Hence optimal solutions might not exist and thus the concept of optimality should be modified to include similar and useful concepts in MODM problems. In practice we can only hope for solutions that *satisfy* each of the objectives to a *sufficient* degree. This idea is encompassed in the following definition.

Definition 2.1 (A Satisficing Solution) *A satisficing solution of Simon[11] extended to a multiple objective decision problem is a reduced subset of the feasible set, which exceeds all of the aspiration levels (goals) of each objective.*

Simon[11] points out that it might be impractical to search for an optimal solution and argues that in practice it is more efficient to look for solutions that are 'good enough'. A satisficing solution corresponds to the behavior of a decision maker who looks for 'good enough' solutions. Another useful notion of optimality in the context of multiple objective decision making is *Pareto-optimality.*

Definition 2.2 (A Pareto-Optimal Solution) $\mathbf{x}^*$ *is a Pareto-optimal solution to a multiple objective decision problem iff there does not exist any* $\mathbf{x} \in X$ *such that* $o_i(\mathbf{x}^*) \leq o_i(\mathbf{x})$ *for all* i *and* $o_j(\mathbf{x}^*) < o_j(\mathbf{x})$ *for at least one* j.

Intuitively, a Pareto-optimal solution is one in which no other solution can improve one objective function without a simultaneous deterioration of at least one of the other objectives. A Pareto-optimal solution is also known as a nondominated solution, a non-inferior solution, or an efficient solution.

2.3 Action selection

It is not so much the methods for generating solutions but the notions of Pareto-optimality and satisficing that play a significant role in formulating behavior coordination mechanisms with useful properties. See [8]

[1] Behaviors in this framework are equivalent to fuzzy behaviors because objective functions encode the semantics of fuzzy membership functions. See [10] for an excellent overview of application of fuzzy logic to behavior-based robotics.

[2] The action space is constrained due to various physical and geometrical properties, such as a set of non-holonomic constraints imposed by the mechanical construction of the robot, into a subset of *feasible* (e.g., physically possible) actions.

for techniques for generating PO and satisficing solutions for multi-objective decision problems. Based on the definitions of Pareto-optimality and satisficing, we propose a behavior coordination scheme that proceeds in the following step-wise fashion:

1. **Feasible actions**. Based on physical and other hard constraints, find the set of feasible and permissable alternatives

2. **Pareto-optimal actions**. Based on objective criteria find the set of Pareto-optimal actions and reject others to reduce significantly the set of alternatives

3. **Satisficing actions**. Incorporate subjective knowledge (weights, priorities, goals, aspirations levels) to find the set of satisficing alternatives

4. **Most preferred action**. Use other criteria to finally select the most preferred action.

This approach systematically narrows down the set of possible solutions, by including increasingly more decision criteria. First, the set of hard constraints given by the nature of the task and physical limitations of the robot (e.g., limited acceleration) are used to exclude the set of infeasible actions. Then, based only on the objective criteria, we can exclude the set of dominated solutions simply because they are not efficient when compared to the Pareto-optimal solutions. By incorporating some subjective preference structure in terms of weights, priorities, etc. we can then select, out of the Pareto-optimal solutions, the set of satisficing, i.e., sufficiently satisfying, solutions. Finally, other criteria or further subjective knowledge can be included to make a final selection. In this way alternatives are evaluated against the behavioral objectives and ranked in terms of permissable, Pareto-optimal, satisficing, and preferred actions.

2.4 Relation to fuzzy control

The (input/output) mechanics of multiple objective behavior coordination are not different from fuzzy approaches, so they can be interchanged as described in [9]. Hence we can formulate the inferencing mechanism using a multiple objective decision-theoretic approach, which provides formal tools for behavior coordination and resolving conflicts in a principled manner. Also with respect to behavior blending [10], fuzzy and multiple objective approaches are compatible and thus can be combined. One straightforward way to combine the two approaches is to: 1) Specify each behavior using a fuzzy rule-base, and generate multivalued behavioral outputs using fuzzy inferencing, 2) Use fuzzy context rules to generate weight values, and 3) Coordinate the behaviors by generating Pareto-optimal and satisficing solutions using multiple objective decision making techniques.

In this way we combine the advantages of approximate reasoning and ease of behavior specification of fuzzy logic with the qualitative and formal behavior coordination capabilities of multiple objective decision making.

3 Cooperative target acquisition

We have applied the described behavior coordination approach to the task of cooperative target acquisition. In its current implementation we have a team of two robots pursuing and surrounding a target. The limitation to two robots only is due to practical considerations such as (lab) space and equipment (robot and sensor) limitations. The approach can, however, easily be extended to the general case of N cooperating robots.

3.1 Group task

The task of the robot team, from a group-level perspective, is to search for a target, surround it, pursue it, and finally capture it. This task can be described, at a high level, by the state diagram in Figure 1(a).

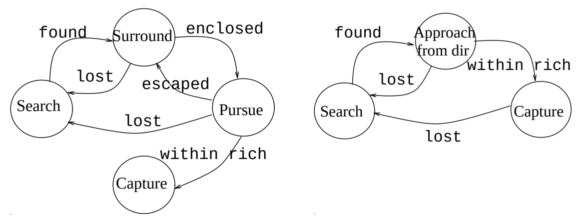

(a) (b)

Figure 1: Finite state description of the target acquisition task. (a) From the group's perspective. (b) From an individual's perspective.

The successful accomplishment of this task is contingent on effective cooperation between the robots. In particular we will focus on the cooperation that is necessary for surrounding the target. Cooperation can also enable an efficient distributed search strategy, however, it is not necessary for target capture and hence will not be the focus of this paper (for an overview of cooperative search strategies see [5]).

3.2 Robot task

From a robot's perspective, the task consists of searching/finding, approaching from a *certain/negotiated direction*, and capturing the target, as depicted in Figure 1(b). In order to decide from which direction it should approach the target, the robot is required to cooperate with its partner. We use techniques from multiple objective behavior coordination to enable such cooperation, as follows.

3.3 Collective vs. individual objectives

In the strategy that we use for surrounding the target, the group's collective objective is to distribute the robots evenly on an arc behind the target. This should results in a formation shown in Figure 2.

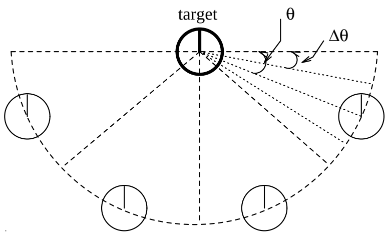

Figure 2: Strategy and group objective for surrounding the target.

For each individual robot this translates to the following objective:

$$\hat{\theta} = \underset{\theta}{argmax} \qquad [o_{self}(\theta), \overline{o_2}(\theta), ..., \overline{o_n}(\theta)] \qquad (2)$$
$$\text{subject to} \qquad \theta \in \cup_{i=1}^{n} [\theta_i - \Delta\theta; \;\; \theta_i + \Delta\theta],$$

where the angle θ describes the direction from the robot to the target and $\Delta\theta$ defines a window of tolerance within which the robot is required to stay (see figure2). In equation (2) it is assumed that $o_1 = o_{self}$ and $\overline{o}(x)$ describes the complement of $o(x)$, meaning NOT $o(x)$ (in this implementation $\overline{o}(x) = 1 - o(x)$).

Equation (2) states that the robot should approach the target from the direction that it desires, and that is not in conflict with the directions of any of its partners. The robot generates its desire for approaching the target from a certain direction based on the target's relative angle to itself. E.g., if the target is on its left then the robot desires to approach from the left. If the target is in front then it desires to approach it either from left or right. The robots do not have a common reference coordinate system, hence we use the target as a common reference point. This is sufficient if the target is moving away from the robots.

To illustrate the dynamics of this strategy consider an example with two cooperating robots. In this case the robots can approach either from left or right. Assume that robot 1 desires to approach from left and robot 2 states that it desires to approach from the left or right. Figure 3 illustrates the objective functions for each robot (left) and their corresponding decision making process (right). Robot 2 will clearly choose to approach from the right. Robot 1 may choose to approach from the left depending on the multiple objective decision making process chosen. E.g., if the robot is required to attain an aspiration level for producing a satisficing solution then it will select one of the two possible solutions (indicated by thick lines) only if the objective function value exceeds the aspiration value and not otherwise. However, in the next (decision making) cycle, robot 2 will state that it desires to approach the target from the right and hence not conflict with robot 1's desire to approach from the left. This resolves the conflict.

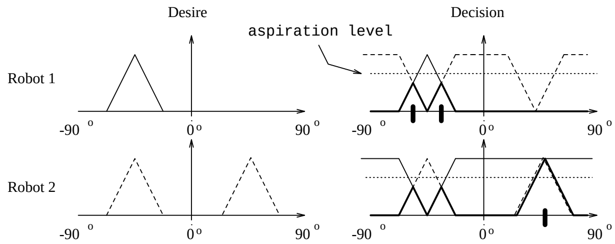

Figure 3: The pursuit direction negotiation technique using multiple objective behavior coordination.

3.4 Context-dependent cooperation

Now consider the case where both robots desire to approach the target from the same direction (say from the left). In this case both will experience a conflict and hence not be able to make a decision. This situation may be resolved over time, e.g., after robots wander around and choose to approach the target from another direction. Another solution would give priorities to the robots, and in cases of conflict the robot with the highest priority is authorized to select a decision.

Alternatively we could assign weight values to the objective functions of the robots, hence the robot with the highest weight would win in a conflicting situation. These solutions, however, impose a hierarchy on the robots, which might be efficient in some cases and not in others. We have chosen an alternative strategy which allows the robots to decide their weight values based on sensory information. Basically, a robot's desire for approaching the target is determined not only by its relative position but is also attenuated as a function of its distance to the target. Thus the objective function of the robot is weighted with a value that is inversely proportional to its distance to the target. This implies that the closer a robot is to the target the higher is its weight value and vice-versa. This approach also makes sense in the overall strategy, because robots that are farther away will have more room to maneuver to get into formation than the robots that are close to the target.

4 System overview

In this section we briefly describe the control architecture of the robots.

4.1 Robot platform

Our system is implemented on two Pioneer-2DX mobile robots with on-board PC-104 stacked Pentium processor running Linux. Each robot is equipped with a Sony Pan-Tilt-Zoom active camera connected to a PXC200 frame grabber, a ring of 16 ultrasonic sensors, 2 position encoders, and a compass. Only the 6 front sonars and the Sony camera are used in our controller. The degrees of freedom that we control are 1) the angular velocity, $\omega \in [-90; 90]$ degs/s; 2) velocity, $v \in [-400; 400]$ mm/s; and 3) absolute camera angle, $\theta \in [-90; 90]$ degrees.

4.2 Software architecture

The control approach we have taken is behavior-based. Multivalued behaviors are generated using fuzzy expressions and action selection is performed using multiple objective behavior coordination mechanisms, in particular the lexicographic approach is used (see [8]). We have previously shown (see [9]) the advantages of this combination of fuzzy and multiple objective behavior coordination.

4.3 Behavior repertoire

In the following we describe each behavior.

Safety This behavior controls the velocity of the robot, based on the front sonar reading, to maintain a safe speed. The rule-base of this behavior is:

```
IF  front_sonar = high_dist    THEN  v = high_speed
IF  front_sonar = low_dist     THEN  v = low_speed
IF  front_sonar = medium_dist  THEN  v = medium_speed
```

Avoid This behavior controls the heading of the robot based on the sonar readings to maintain a safe travel direction.

```
IF  front_sonar = low_dist     THEN  w = sharp_left
                                     OR sharp_right
IF front_sonar  = medium_dist  THEN  w = left OR right
IF left_sonar   = low_dist     THEN  w = sharp_right
IF left_sonar   = medium_dist  THEN  w = right
IF right_sonar  = low_dist     THEN  w = sharp_left
IF right_sonar  = medium_dist  THEN  w = left
```

Face This behavior servos the heading of the robot (w), based on the input from vision system, so as to face the target. The vision system provides the image position of the target (based on chromatic color segmentation). The difference between the target position and the center of the image generate an error vector which this behavior attempts to minimize. As a result the robot will rotate until the target is centered in the image.

```
IF  target_pos = front  THEN  w = forward
IF  target_pos = left   THEN  w = left
IF  target_pos = right  THEN  w = right
```

LookAt This behavior controls the pan angle of the camera (p) to a given angle (ang), (determined by negotiating with other robots) towards the target as it approaches the target. Since the robot's sensors do not provide distance information to the target, we use the appeared size of the target in the image as a value proportional to it. This behavior is triggered only when the target is in the robot's field of view.

```
IF  target_size = large  THEN  p = around(ang)
```

where `around` is a triangular membership function centered at `ang`. The effect of combining this behavior with the Approach behavior is that when the robot is far from the target it will direct the camera to look straight ahead and thus drive the robot closer to the target. As the robot gets closer it will turn the camera head toward the target and hence align it accordingly. Thus by controlling the camera pan angle we can align the robot to the target as desired.

Surround This behavior controls the camera pan angle to express the desire for which direction to approach the target from.

```
IF target_pos = forward THEN p = left or right
IF target_pos = left    THEN p = left
IF target_pos = right   THEN p = right
```

The output from this behavior and the corresponding one from the partner robot is used to select the direction.

Rotate The Rotate behavior controls the angular velocity of the robot to rotate left or right. It merely causes the robot to turn in one or the other direction to run in a circle and help the robot look for the target.

4.4 Action selection

For coordination of the activities of the behaviors we use a combination of arbitration, command fusion, and context-dependent blending mechanisms. Figure 4 depicts the behavior hierarchy and the relationship between the behaviors. First, the phase between Searching and Approaching the target is clearly distinguished by the visibility of the target in the robot's (camera's) field of view. When the target is not visible, the Search behavior is activated and when it is, the Approach behavior is activated. In either case the robot's safety behaviors (Safety and Avoid) are activated, and MOBC command fusion is used for action selection.

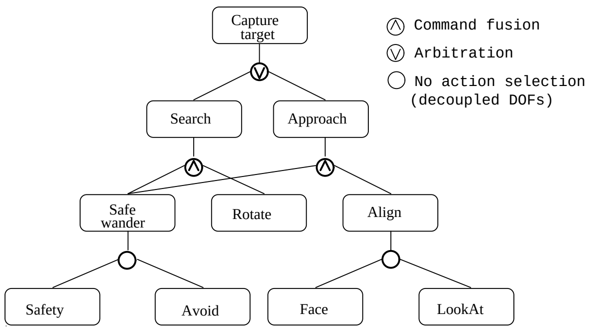

Figure 4: Behavior hierarchy and action selection in the target acquisition task.

5 Experimental results

In this section we describe experimental validation results from trials in a confined lab environment. In the experimental setup we used two Pioneer robots (see section 4.1 for detailed description) for target acquisition and a manually guided dummy as the target. (In the future we plan to have an autonomous robot replacing the dummy). In order to log the position and orientation of the robots and the target, the experiments where conducted in a corral ($6 \times 4m^2$) with an overhead camera tracking system capable of simultaneously tracking up to eight markers' positions and orientations. Our lab experiments, described next, show very reliable results and successful cooperative target acquisition in most of the trials. A more complete battery of experiments is in progress.

5.1 Validation experiments

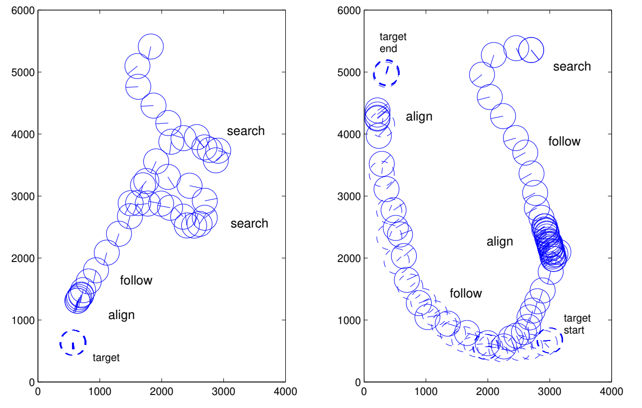

(a) (b)

Figure 5: Aligning behavior. (a) Static target. (b) Moving target.

First we conducted some simple experiments using only one robot to follow the target. The results show that the robot searches for the target, finds it and then approaches it from some (predetermined) direction (see figure 5(a)). Since the field of view of the robot camera is narrow (about 30^o), the target is easy to lose when the robot is close to it. However, the robot immediately searches in the direction where it last detected the target and usually finds it within a few seconds (see figure 5(b)). These results validate that a single robot can follow and pursue a target.

5.2 Group experiments

Allowing two robots to perform their task should result in the robots surrounding the target, which includes cooperative decision making to negotiate from which direction each robot should approach the target.

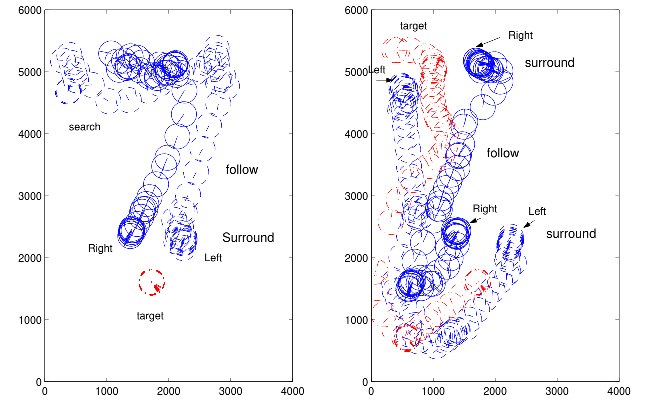

(a) (b)

Figure 6: Surrounding behavior. (a) Static target. (b) Moving target.

Figure 6 shows two segments of a single trial. In the left the target is initially static 6(a). After a short search both robots detect the target and approach it, and align to it ending in a stable formation surrounding the target. In figure 6(b) the target starts moving and the robots follow it in the same formation. What cannot be seen from these trail plots are the temporal occurrences of events. In the example above when the target starts moving the robot on its right follows immediately while the other holds back until the path is clear and then proceeds. The system does not have any mechanisms for explicit coordination of such situations; they result from the interaction of the robots with each other and their environment, within the constraints imposed by the action selection mechanism.

6 Conclusions and future work

In this paper we have extended the multiple objective behavior coordination approach to cooperative multi robot tasks. The approach can easily be extended to any alternative command fusion mechanism (e.g., voting, fuzzy etc.). We can thus conclude that the work described here extends command fusion action selection mechanisms to multi-robot coordination. In contrast to behavior arbitration mechanisms (e.g., ALLIANCE [7]) that essentially allow for coordination in 'turn-taking' activities, multi-robot command fusion allows for pursuit of multiple goals by multiple robots in parallel.

We have demonstrated how multi-robot command fusion can be used to produce dynamic task division in the context of cooperative target acquisition. Our experimental results validate that the approach produces reliable cooperation. Our current implementation employs two robots, but the approach can easily be scaled to an arbitrary number of robots (disregarding any problems caused by interference).

One problem that we anticipate with large numbers of robots is the amount of information that has to be exchanged between them. In the future we will investigate strategies where each robot will have to communicate with its k nearest neighboring robots. For many tasks we believe this will produce acceptable results. Further, we intend to investigate the applicability of this approach to other tasks, e.g., distributed cooperative searching.

References

[1] Ronald C. Arkin. *Behavior-Based Robotics.* Intelligent Robotics and Autonomous Agents series. MIT Press, May 1998.

[2] Tucker Balch and Ronald C. Arkin. Behavior-based formation control for multi-robot teams. *IEEE Transaction on Robotics and Automation (To appear)*, 1999.

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

[4] Vira Chankong and Yacov Y. Haimes. *Multiobjective Decision Making - Theory and Methodology*, volume 8. North-Holland, 1983.

[5] Miguel Schneider Fontan and Maja J. Mataric. Territorial multi-robot task division. *IEEE Transactions of Robotics and Automation*, 14(5), 1998.

[6] Maja J Matarić. Issues and approaches in the design of collective autonomous agents. *Robotics and Autonomous Systems*, 16(2–4):321–331, December 1995.

[7] Lynne P. Parker. Cooperative robotics for multi-target observation. *Intelligent Automation and Soft Computing, special issue on Robotics Research at Oak Ridge National Laboratory*, 5(1):5–19, 1999.

[8] Paolo Pirjanian. *Multiple Objective Action Selection & Behavior Fusion using Voting.* PhD thesis, Department of Medical Informatics and Image Analysis, Institute of Electronic Systems, Aalborg University, Fredrik Bajers Vej 7, DK-9220 Aalborg, Denmark, August 1998. *Available online: http://www-robotics.usc.edu/˜paolo/publications.*

[9] Paolo Pirjanian and Maja Mataric. A decision-theoretic approach to fuzzy behavior coordination. In *IEEE International Symposium on Computational Intelligence in Robotics and Automation*, November 1999.

[10] Alessandro Saffiotti. The Uses of Fuzzy Logic in Autonomous Robot Navigation: a catalogue raisonné. Technical Report 2.1, IRIDIA, Universite Libré de Bruxelles, 50 av. F. Rossevelt, CP 194/6, B-1050 Brussels, Belgium, November 1997.

[11] Herbert A. Simon. *The New Science of Management Decision.* Harper and Row, New York, 1960.