Multi-Robot Team Coordination using Desirabilities*

Alessandro Saffiotti
Applied Autonomous Sensor Systems
Department of Technology, Örebro University, Örebro, Sweden
asaffio@aass.oru.se

Nina B. Zumel†
Artificial Intelligence Center, SRI International
Menlo Park, California, USA
zumel@resonantcommerce.com

Enrique H. Ruspini
Artificial Intelligence Center, SRI International
Menlo Park, California, USA
ruspini@ai.sri.com

Abstract.
Desirability functions are an effective way to express and implement complex behavior coordination strategies inside a single robot. In this paper, we extend the desirability function approach to deal with behaviors of teams of robots. We show that desirability functions offer a convenient tool to incorporate and blend individual objectives and team objectives. We illustrate our approach on two significant problems of team coordination: reactive formation motion control, and collaborative searching and tracking.

1 Introduction

The coordination of the actions of teams of collaborating autonomous mobile agents presents problems that are considerably more complex than those typically considered when planning and regulating the motion of single robots. The diversity of team interactions and possible formations, inherent communication constraints, and issues related to distribution of the control function complicate the characterization and analysis of the interactions between various global and individual behaviors.

The desirability function approach to robot control [3, 4]—stressing the principled application of utilitarian notions—provides an attractive avenue for the definition of global and individual behaviors and for the study of their interactions. Desirability functions readily permit the description of the consequences of both individual and team actions using the same conceptual structures. These goal-dependent preference measures may then be combined by the same logic-based methods employed to blend the reactive and purposive behaviors of individual mobile agents. Furthermore, reliance on context-dependent activation of behaviors provides explicit mechanisms for

*The work reported in this paper was supported by the United States Office of Naval Research under Contract No. N00014-99-C-0298 and by the Swedish KK Foundation.

†Current affiliation: Resonant Commerce, Redwood City, California

inter-agent regulation of global behavior (e.g., task passing and coordination). Finally, communication requirements may be examined in terms of the ability of the various controllers to compute the underlying desirability structures.

In this paper we introduce basic concepts germane to the definition and blending of reactive and purposive behaviors regulating the individual and global actions of teams of cooperating robots. We illustrate our approach by presentation of results from two significant problems of team coordination: reactive formation motion control and collaborative searching and tracking.

2 Desirability Functions

In our approach to robot control [4], we regard desirable behavioral traits as quantitative *preference* functions defined over the set possible control actions from the perspective of the goal associated with that behavior. Let S be the set of internal states, or knowledge states, of the robot, and U the set of its possible actions. Following [3], we describe each behavior in terms of a desirability function $D(s,u)$—taking values in $[0,1]$—that measures the *goal-specific* desirability of applying the control u in state s. Equivalently, we can say that Des associates each situation s with a fuzzy set $D = \mathit{Des}(s,u)$ of desirable control values.

A desirability function encodes the preferences of a behavior about alternative actions. For instance, let S be the set of perceived positions of a target point and let U be the set of steering angles for a robot operating in two-dimensional euclidean space. A *go-to-target behavior*, for example, could produce, for each situation s, a triangular fuzzy set D_g of steering angles centered on the direction of the target.

When the state s is fixed, a desirability function Des induces in U a distribution D over the set U of possible actions by $D(u) = \mathit{Des}(s,u)$. The desirability function D induces an ordinal meaning among possible actions, that is, performing u_1 is preferable to performing u_2 from the perspective of our goal if and only if $D(u_1) > D(u_2)$. A desirability function D also induces a control law in the obvious way specifying any maximizing value as the most desirable control.

Desirability distributions pertaining to different goals may be interpreted and manipulated employing concepts and techniques derived from fuzzy logic [3]. In particular, two desirabilities D_1 and D_2 may be conjunctively combined into $D = D_1 \cap D_2$ by

$$D(u) = \min(D_1(u), D_2(u)) \ .$$

The combined desirability D expresses the common preferences between D_1 and D_2.[1]

Desirability functions can be used to encode robot behaviors [4]. The conjunctive combination of two behaviors, each aimed at a different purpose, yields a composite behavior that satisfies both purposes at the best possible degree. Maximization of this combined desirability may then be used to generate a tradeoff control value that satisfies both goals as much as possible: a combination strategy that leads to a Pareto optimal behavior [2].

Desirability functions are essentially different from the control functions employed by other robotic approaches to perform local behavior combination since they do not simply determine the most-preferred action but, rather, they produce a full preference valuation over the entire range of possible actions. When behavior combination is solely

[1]More generally, the min operator can be replaced by any triangular norm.

based on consideration of optimal values for each behavior—as is the case with potential-field methods—the resulting strategy is necessarily hampered by the unavailability of information required to make rational tradeoff decisions.

3 Multi-robot Coordination

Desirability-function approaches may be readily extended to the coordination of groups of collaborating autonomous mobile real robots. The major conceptual idea behind this generalization is the computation of joint preferences—considering potential desirable and undesirable inter-agent interactions—defined over the space of potential actions of all the mobile agents. In other words, joint actions of the team as a whole may also be ordered in terms of their relative preference as promoters of overall group goals. In general, however, the related joint desirability functions will defined over complex control/decision spaces of high dimensionality. We present examples of two methods, each based on a different approach to the computation of the joint desirability function, for the treatment of problems related to the complexity of joint decision spaces.

We discuss first, in the context of a problem arising from the control of the formation of a team of multiple robots moving through a field of obstacles, a technique based on the combination, in the decision-space of the actions of all agents, of simple agent-specific and global preference functions along lines suggested by local uncertainty propagation methods [5].

In another example—suggested by a tracking and pursuit problem—the determination of group preferences is made on a robot by robot basis by combination of each agent preferences with preferences derived from other agents. The *wandering* behavior of Agent A may, for example, be modified by the action of Agent B, who wishes A to move towards the position of a target known to B (*target handoff*). In this case, computation of the preferences of a particular agent simply reduces to the combination, for each mobile agent, of several desirability functions defined on the decision space of that particular agent rather than on more complex decision spaces.

It is important to remark, however, that—regardless of the computational approach employed to produce joint preference functions and its related control actions—these approaches are conceptually similar in that they rely on combination of agent-specific and coordinative desirability functions to combine elastic constraints restricting the extent of the overall actions of each agent and of the team as a whole.

3.1 Reactive Formation Motion Control

Desirability functions can be extended to express coordination strategies between multiple robots. Let $R = \{R_1, R_2, \ldots, R_n\}$ be a team of robots. For notational simplicity, we assume that all robots have the same S and U spaces, although extension to heterogeneous robots is straightforward. We can treat R as one single robot with internal state S^n and set of possible actions U^n. For R to perform the joint action $\vec{u} = (u_1, u_2, \ldots, u_n)$ means that each robot R_i performs action u_i.

The desired joint behavior of the R team can be expressed by associating each (joint) state $\vec{s} \in S^n$ a to a *team desirability function*

$$D_{\text{team}} : U^n \to [0, 1]$$

defined over the (joint) set U^n of global robot control actions, such that $D(\vec{u})$ measures the desirability of R's performing the combined action $\vec{u}$ from the point of view of the desired team behavior. For instance, D_{team} can give high desirability to vectors of steering actions that maintain a given relative configuration, or *formation*, between the robots.

The desirability D_{team} can be combined with those of the individual robots to produce a combined desirability function for each robot that takes into account both their own individual goals—such as avoiding nearby obstacles—and the collective goals—such as maintaining a formation. Let $\vec{s} \in S^n$ be the current joint state. We denote by D_i the individual desirability of R_i. D_i measures the desirability of individual actions from the point of view of R_i alone, e.g., avoiding an incoming obstacle. (D_i is identically 1 if R_i does not have preferences.) For each robot R_i, we extend D_i to the joint space U^n by *vacuous extension*, i.e.,

$$D^{\uparrow i}(\vec{u}) = D_i(u_i). \tag{1}$$

All the extended desirabilities $D^{\uparrow i}$'s can then be combined with the team desirability D_{team} to obtain an overall joint desirability D on U^n by

$$D(\vec{u}) \;=\; D_{\text{team}}(\vec{u}) \;\otimes\; \min_{i=1,\ldots,n} D^{\uparrow i}(\vec{u}), \tag{2}$$

where $\otimes$ is a T-norm (min in our experiments). Individual decisions to be taken by each robot R_i are obtained by projection of this joint desirability into the control spaces of the individual robot by

$$D^{\downarrow i}(u_i) = \max_{\substack{u_k \in U \\ k \neq i}} \; D(u_1, \ldots, u_n), \tag{3}$$

where the max is taken by varying all the u_k while keeping u_i fixed.

The above procedure is best illustrated by an example. Consider the situation shown in Figure 1. Each of the two robots R_1 (below) and R_2 (above) has two individual objectives: avoiding obstacles, and going "east" (right in the picture). Moreover, the two robots have the team objective of keeping a fixed approximate distance between them. Actions are pairs of steering controls. In the situation shown in the figure, the robots are facing two obstacles, and must decide which way to turn in order to avoid them. The thin lines in the graphs on the right plot the individual desirabilities for avoiding obstacles and for going east for each robot. From the point of view of the individual desirabilities, R_2 has a strong preference to turn right, while R_1 can indifferently turn in either direction. However, in order to satisfy the team desirability of maintaining a fixed distance, both robots should turn in the same direction; hence, the option for R_1 to turn right is undesirable. The thick lines show the resulting desirabilities after combination of team and individual preferences. The preferred controls, indicated by the black arrows, steer both robots to the right. Note that most common approaches, like those based on potential-fields, summarize preferences solely in terms of a "most preferred" action, and combine these actions instead of combining the full desirability functions (as done, for example, by Balch and Arkin, [1]). In our example, application of these approaches might result in each robot individually deciding to steer toward the other one when first approaching the obstacle, which would eventually lead them to collide.

Figure 2 shows the steps involved in this combination, where we only consider the steering angle for simplicity. Each desirability function is generated by a specific *behavior* that considers one specific objective [4]. In particular, the individual obstacle avoidance behaviors generate desirability functions that assign lower desirability to controls that steer the robot in the direction of an obstacle; and the team desirability behavior generates a function that assigns higher desirabilities to pairs of steering controls that bring the robots closer to the desired distance.

Fig. 3 shows a full run in that obstacle field. By using elastic contraints and desirability functions, the robots could negotiate the obstacles while maitaining the desired heading and the desired formation as much as possible.

3.2 Collaborative Searching and Tracking

Another important problem involving team and individual behaviors involves multiple trackers cooperating to ensure that a target is continually covered, that is, there is always a tracker who is within a given distance from the target. We assume a three-dimensional operational environment where, for simplicity, the target moves at a different height than that of the trackers. The target is assumed to have speed and maneuverability comparable to that of the trackers, each of which has a primary sector of responsibility (which overlaps slightly those of other trackers). Sensors are modeled as being omnidirectional with finite range: the agent can see an area of the ground equivalent to about 20% of the sector it is responsible for. Communication from one agent to another is limited to conveying the target's location and estimated velocity.

Figure 4 shows an example configuration. Each behavior active at any time generates an individual desirability function over the agent's possible direction and speed. Each such preference function is weighted by a context-dependent *activation*, indicating how important a particular goal or constraint is in a given situation, then combined to produce an action that considers individual and global goals. In Figure 4, the activations for each agent at the moment shown are displayed on the left-hand bar chart.

Each agent A_i has the same strategy, i.e.,

- *Follow* the target while in its sector
- If capable, when alerted by another tracking agent that the target is in its sector, *take over* tracking duties
- When target leaves own sector, notify agent in new sector and continue to track target until the latter takes over
- When relieved of tracking duties outside own sector, *go home*
- If target is not seen or reported, just *wander*
- Avoid colliding with other agents (*repel*)

This agent strategy (Case 2 in Figure 5), was compared to two idealized cases. In Case 1, all the agents know the location of the target all of the time (Infinite Sensing). In Case 3, the agents have the same sensor range as in Case 2 but no there was no communication between the agents. Figure 5 shows a comparison of the target coverage for the same initial conditions and target motion in the three cases. For this test, the Infinite Sensing case had the best performance in terms of total time of target coverage. Surprisingly, though, Case 2 outperformed Case 1 in terms of the target coverage being handled by the correct agent(s)—that is, for Case 2, the target was more often covered by an agent responsible for the current sector. The third bar graph in Figure 5 shows the median time that any agent spent outside of its sector for the three cases.

Figure 6 illustrates the target coverage by agent, showing the ratio of target coverage by agent over the time spent by the target spent in that agent's sector. Ideally, all ratios should be near unity. For this experiment, Case 2 was much better than Case 3 (no communication), and even somewhat more evenly distributed than in the Infinite Sensing case. A possible reason for this is that in the Infinite Sensing case, the agents could see “too much,” and would all rush to cover the target when it was unattended, regardless of which sector it was in. As a result, the agents would all interfere with each other at the target, and often would be badly positioned when the target did in fact end up in their area of responsibility. Explicit negotiation behavior would alleviate this problem but it is interesting to see that with the finite visual range (and some information exchange), this problem is less severe. In the No Communication case, it was often true that the agent responsible for a sector would be badly positioned by the time the target arrived in its sector.

In this example, global behavior blending is simplified by the nature of the policy being followed for coordination as this integration is equivalent to the combination, for each robot and on its decision space, of the desirabilities induced by its active behaviors (regardless of the interactions leading to the activation of those behaviors). If, on the other hand, more complex policies are considered (e.g., continuous task reallocation as a function of the ongoing situation and available resources), the dependencies that are created would, generally, necessitate the consideration of more complex combination schemes (i.e., involving, as is the case with formation control, the consideration of nontrivial constraints on the joint actions of several agents).

4 Conclusion

The desirability-function approach provides an effective framework for the integration of goals and constraints restricting the joint actions of a team of collaborating robots with individual objectives and restraints placed on individual members. The resulting methodology allows the computational of rational tradeoffs between competing objectives while being less liable to problems, such as local minima, that are common in approaches that summarize preferences solely in terms of a “most preferred” action.

We are currently extending our approach in various directions, including the consideration of constraints and goals regulating wider classes of actions and its generalization to the reformulation and integration of global strategies (e.g., path replanning). Finally, we are developing analytical tools that permit determination of properties of integrated behaviors, such as performance degradation due to behavior combination.

References

[1] T. Balch and R.C. Arkin. Behavior-based formation control for multi-robot teams. *IEEE T. on Robotics and Automation*, 14(6):926–939, 1998.

[2] P. Pirjanian and M. Mataric. Optimality issues in fuzzy behavior coordination. In D. Driankov and A. Saffiotti, editors, *Fuzzy logic in autonomous robot navigation*. Springer-Verlag, 2000.

[3] E. H. Ruspini. Truth as utility: A conceptual synthesis. In *Proc. of the Conf. on Uncertainty in Artif. Intell.*, pages 316–322, Los Angeles, CA, 1991.

[4] A. Saffiotti, K. Konolige, and E. H. Ruspini. A multivalued-logic approach to integrating planning and control. *Artificial Intelligence*, 76(1-2):481–526, 1995.

[5] Prakash P. Shenoy and Glen Shafer. Propagating belief functions with local computations. *IEEE Expert*, 1(3):43–52, Fall 1986.

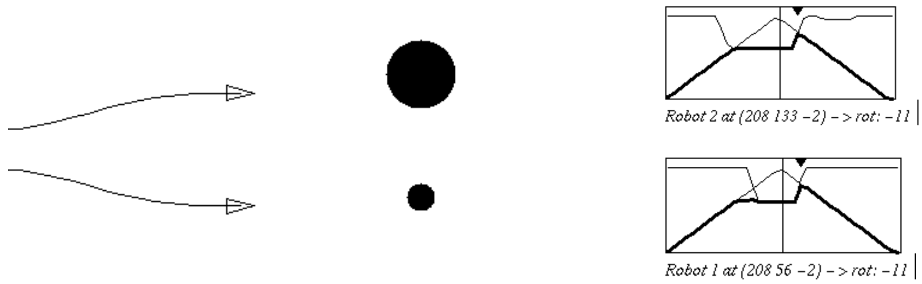

Figure 1: Negotiating two obstacles while maintaining desired inter-robot distance

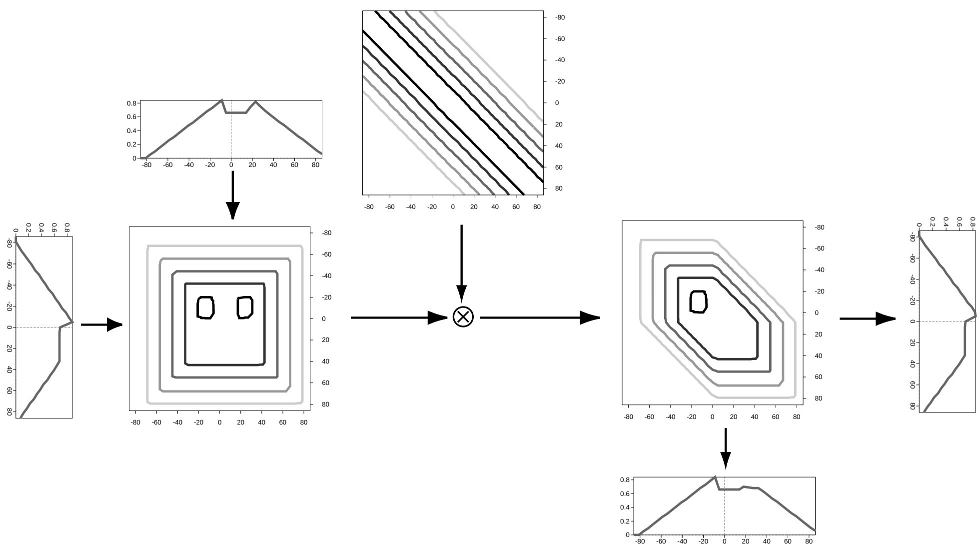

Figure 2: Computation of Joint and Marginal Desirability Functions

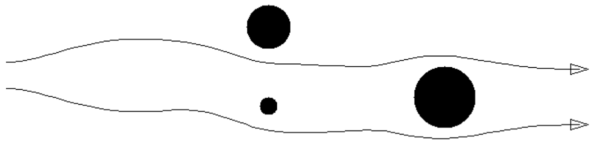

Figure 3: Full Run in an Obstacle Field

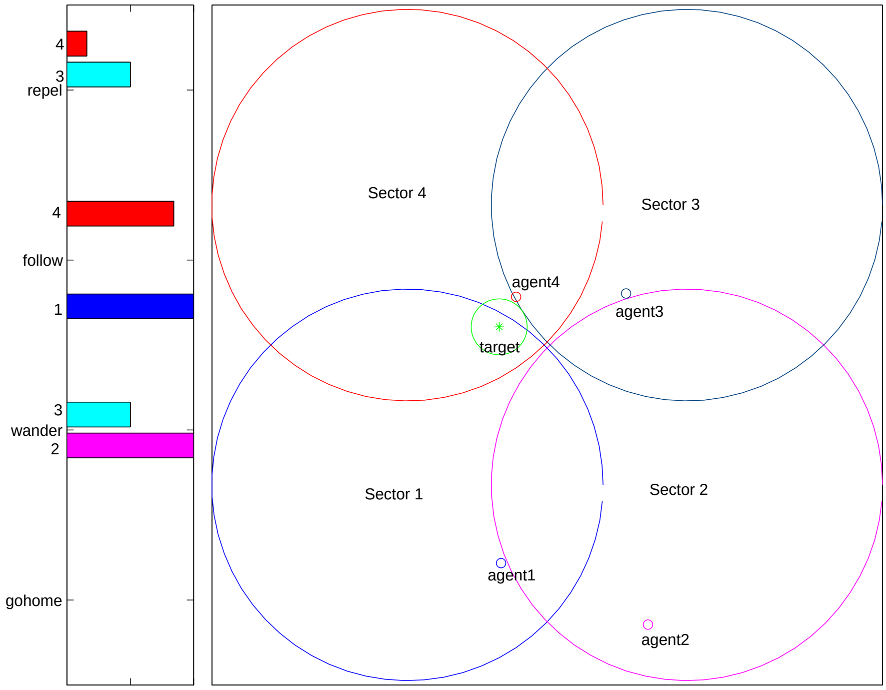

Figure 4: Multiple agents covering a target

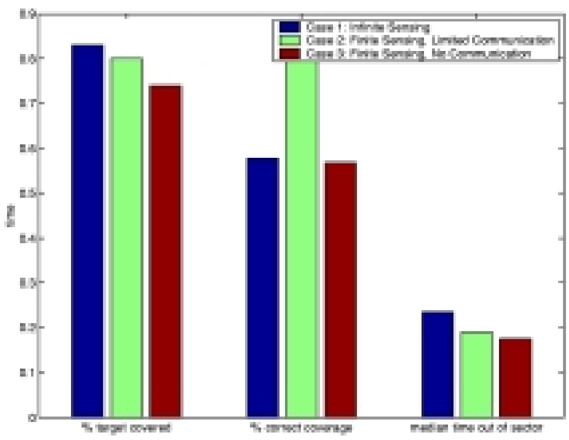

Figure 5: Overall Target Coverage

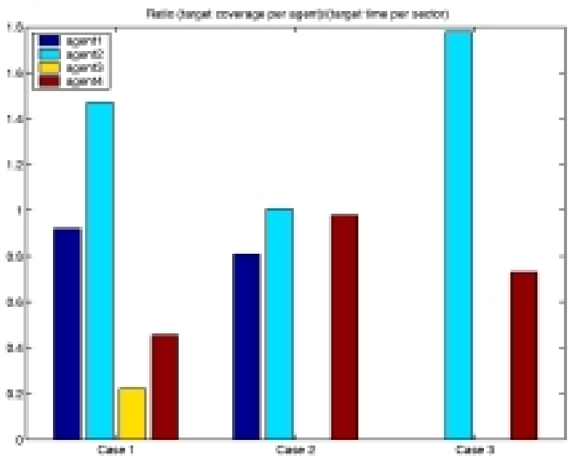

Figure 6: Coverage by Agent