GLOBAL TEAM COORDINATION BY LOCAL COMPUTATION

Alessandro Saffiotti*, **Enrique H. Ruspini**†

* Center for Applied Autonomous Sensor Systems
Dept. of Technology, Örebro University, Örebro, Sweden
e-mail: asaffio@aass.oru.se
http://www.aass.oru.se
† AI Center, SRI International
333 Ravenswood Ave, CA-94025 Menlo Park, California
e-mail: ruspini@ai.sri.com
http://www.ai.sri.com/

Keywords: cooperative robotics, desirability functions, team formation, fuzzy control.

Abstract

Desirability functions are an effective way to define group and individual objectives of a team of cooperating mobile robots. By combining desirability functions, we can identify the individual actions that best satisfy both sets of objectives. Combination, however, is global, posing high demands in terms of communication and computation resources. In this paper, we investigate a technique to perform this combination using local computations. Simulated experiments suggest that, under conditions of spatial locality, team control by local computation achieves the same performance than using a global technique.

1 Introduction

An essential problem in the coordinated control of a team of cooperating mobile robots is how to select each robot's actions so that both the team's and the individual objectives are satisfied. For example, consider the situation shown in Fig. 1 involving two mobile robots R1 (above) and R2 (below) that navigate in an obstacle field. Each robot has two *individual objectives*: to avoid the obstacles, and to go "east" (right in the picture). Moreover, the two robots have the *team objective* of keeping a fixed distance between them. 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. From an individual point of view, the best choice for R1 is to turn right, and for R2 to turn left. However, these choices are undesirable for the team objective, since they would bring the robots too close to each other. This problem can only be detected when considering the actions of both robots simultaneously.

The desirability function approach to robot control [5, 7], stressing the principled application of utilitarian notions, provides an attractive avenue for the definition of robot objectives and behavior. In [8], we have shown that desirability functions are also a convenient tool for defining group and individual objectives of a team of robots. Group and individual desirability functions can then be fused to yield the actions that best satisfy both sets of objectives. Intuitively, we can combine both classes of desirability functions on a *global action space*, and then project the resulting global desirability function back on the action space of each robot. Each robot then takes a decision based on this projected desirability function.

The above approach has shown good results in simulated experiments of team coordination for formation control. However, performing the combination on the global action space entails demanding conditions in terms of communication and computation resources. In this paper, we investigate a technique to compute the global desirability function by only performing *local computations*. The starting point is the framework for local computations originally proposed by Shenoy and Shafer for propagating uncertainty over a network [9, 10]. The main advantage of the local solution is that it only involves exchanging and combining information between small subsets of robots.

2 Desirability Functions

In our approach to robot control [7], 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 [5], 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 control u in state s. Equivalently, we can say that D associates each state s with a fuzzy subset D_s of desirable control values.

A desirability function encodes the preferences of a behavior about alternative actions. For instance, let S be the set of per-

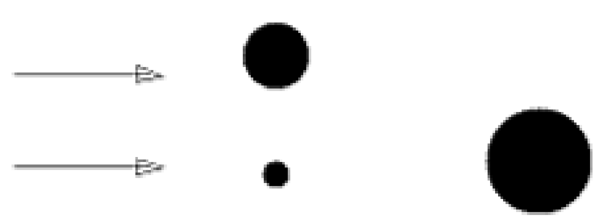

Figure 1: How can the robots negotiate the obstacles and maintain formation?

Proc. of the European Control Conference. Porto, Portugal, September 2001.

ceived positions of a target and let U be the set of steering angles for a robot operating in two-dimensional space. A *go-to-target behavior* could produce, for each situation s, a triangular fuzzy subset D_{goto} of U centered on the direction of the target. (We omit the s subscript for simplicity.)

A desirability function D induces a control law in the obvious way specifying any maximizing value as the most desirable control. Desirability functions corresponding to different goals can be combined using fuzzy logic techniques. In particular, two desirabilities D_1 and D_2 can be combined conjunctively [7] into $D = D_1 \cap D_2$ given by:

$$D(u) = D_1(u) \wedge D_2(u),$$

where $\wedge$ denotes minimum.[1] The combined desirability D expresses the common preferences between D_1 and D_2. 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 [4].

Desirability functions are essentially different from the control functions used in other approaches to perform local behavior combination in autonomous robotics since they do not simply produce the most preferred action but, rather, they give a full preference valuation over the entire range of possible actions. Should this combination have been based on sole consideration of optimal values for each behavior—as is the case with potential-field methods—the resulting strategy would lead to local minima as information necessary to consider tradeoff decisions would be lacking. Desirability functions have been extensively used for solving autonomous robot navigation problems, resulting in robust behaviors and in flexible behavior coordination strategies [7, 6, 2].

3 Team Coordination by Global Computation

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. To say that R performs the joint action $\vec{u} = (u_1, u_2, \ldots, u_n)$ means that each robot R_i performs action u_i.

The desired team 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. Intuitively, $D(\vec{u})$ 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

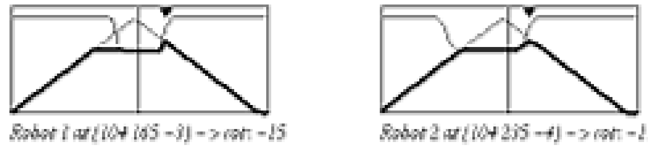

Figure 2: The individual (thin lines) and marginal (thick lines) desirabilities for the steering angle in the situation shown in Figure 1. The small triangles show the preferred control values.

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. For each robot R_i, we extend D_i to the joint space U^n by *vacuous extension*:

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

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

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

Individual decisions to be taken by each robot are obtained by marginalization of this joint desirability to the control space of each individual robot by

$$D^{\downarrow i}(u_i) = \max_{\vec{u} \in U - \{u_i\}} D(\vec{u}, u_i), \tag{3}$$

where the max is taken by varying all the $u_k, k \neq i$, while keeping u_i fixed.

The above procedure is best illustrated by an example. Consider the situation discussed in the Introduction — see Fig. 1. Fig. 2 plots the individual desirability functions of the two robots over the space of possible steering angles. The thin lines show the initial desirabilities for the "avoid obstacles" and the "go east" objectives. From the point of view of these desirabilities, R_2 has a definite preference for a right turn, while R_1 has a preference for a right turn and a slightly stronger one for a left turn. 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 left is undesirable. The thick lines show the desirabilities resulting from the combination of team and individual preferences through equations (1–3): now the most preferred action for both robots is to turn right. Fig. 3 graphically shows the steps involved in this computation. Fig. 4 shows an extended run.

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

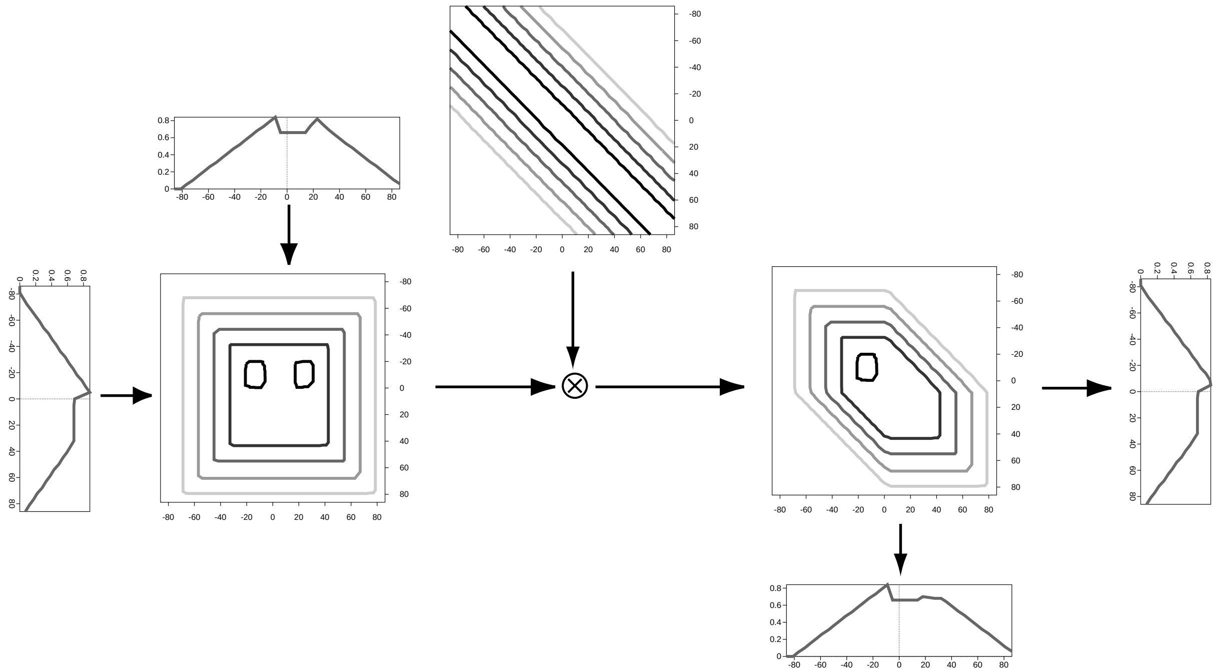

Figure 3: Global computation of joint and marginal desirability functions. Left: the individual desirabilities are expanded on the joint action space and combined. Middle: the team desirability function is also combined. Right: the result is projected on each robot's individual action space.

4 Team Coordination by Local Computation

The above approach to team coordination requires that all (individual and team) desirability functions be collected and combined together on the joint action space in order to decide each robot's actions. When the number of robot increases, this may become extremely demanding in terms of communication and computation resources. In this section, we show how we can compute the global desirability function in an implicit way by only exchanging and combining local information. Intuitively, we only work on sub-spaces of the full joint space that correspond to sub-groups of closely related robots.

4.1 Valuation Algebras

Valuation Algebra (VA's) are an abstract framework for computing marginals using local computation originally defined by Shenoy and Shafer [9, 10, 3]. In VA's the process of propagating values over a graph has been abstracted from the specific nature of the values being propagated. VA's capture many propagation schemes as special cases, including constraint satisfaction systems and Bayesian networks.

A VA is a tuple $\langle \mathcal{X}, \mathcal{V}, \otimes, \downarrow \rangle$ where: $\mathcal{X} = \{x_1, \ldots, x_n\}$ is a set of *variables*, each associated to a set w_i of values (its *frame*); $\mathcal{V}$ is a set of *valuations*, objects that represent information about (sets of) variables; $\otimes : \mathcal{V} \times \mathcal{V} \to \mathcal{V}$ (*combination*) is an operator intended to capture aggregation of information; and $\downarrow : \mathcal{V} \times 2^{\mathcal{X}} \to \mathcal{V}$ (*marginalization*) is an operator that projects a valuation over a smaller frame. For instance, a probabilistic system can be seen as a VA where valuations are probability distributions, $\otimes$ is point-wise product, and $\downarrow$ is the usual marginalization.

Once we have a valuation algebra, we can *evaluate* it. This means to: (i) expand all the valuations in $\mathcal{V}$ to the global frame $W = w_1 \times \ldots \times w_n$; (ii) combine, via $\otimes$, all the expanded valuations into a new valuation V_W; and (iii) compute the marginals of V_W on the (subsets of) variables of interest. Formally, evaluating a VA with respect to $h \subseteq \mathcal{X}$ means to compute

$$\Big(\bigotimes_{v \in \mathcal{V}} v^{\uparrow W}\Big)^{\downarrow h}. \qquad (4)$$

Evaluation thus finds the impact of the available information on the variables of interest. Notice that this process is analogous to the one described by equations (1–3) above.

4.2 Local Computation

The evaluation of a VA involves a $\otimes$-combination over the entire frame W. This may easily become computationally intractable, since the size of W grows exponentially with the number of variables in $\mathcal{X}$. Shenoy and colleagues have defined a local computation scheme to evaluate a VA by only performing local combinations over restricted sub-frames of W, that is, without explicitly computing the overall V_W.

The basic observation is that, in order to compute (4), we can ignore all the variables not in h, provided that every time we delete a variable x we modify the valuations in $\mathcal{V}$ that included

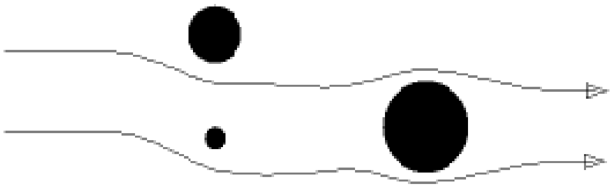

Figure 4: Combined obstacle avoidance and formation control with two robots using global computation

that variable by

$$\mathrm{Del}_x(\mathcal{V}) = \{ \big(\oplus \{ v \mid x \in d(v) \} \big)^{\downarrow W - x} \} \;\cup\; \{ v \mid x \notin d(v) \}$$

where $d(v) \subseteq \mathcal{X}$ denotes the domain of v [3]. Thus, we can stepwise reduce our VA until $\mathcal{X}$ only includes h. What is important here is that each step only involves a $\otimes$-combination on a sub-frame of W. These sub-frames are in general much smaller than W, thus reducing the computational complexity; the amount of this reduction depends on the size of the largest frame of any valuation in $\mathcal{V}$.

Evaluation of a VA by stepwise reduction can be effectively implemented on graph-based representations [10, 3]. The interest of VA's comes from the fact that no strong assumptions must be made on the nature of elements of the $\langle \mathcal{X}, \mathcal{V}, \otimes, \downarrow \rangle$ algebra in order to apply this local computation scheme. All what is required is that the following axioms be satisfied:

A0. $G^{\downarrow g} = G$
A1. $G^{\downarrow (h \cap k)} = (G^{\downarrow h})^{\downarrow k}$
A2. $G \otimes H = H \otimes G$
A3. $G \otimes (H \otimes K) = (G \otimes H) \otimes K$
A4. $(G \otimes H)^{\downarrow g} = G \otimes (H^{\downarrow g \cap h})$

where G, H, K are valuations on $g, h, k \subseteq \mathcal{X}$. For example, probability distributions, possibilities, and belief functions, all satisfy these axioms [10].

4.3 Valuation Algebra of Desirabilities

We can use the elements defined in Section 3 to build the valuation algebra $\langle U^n, \mathcal{D}, \wedge, \downarrow \rangle$, where $\mathcal{D} = \{D_1, \ldots, D_n, D_{\text{team}}\}$. The desirability function $D^{\downarrow i}(u_i)$ computed by (3) corresponds to the result of evaluating this VA with respect to the control variables u_i of robot R_i, i.e., to equation (3) where $h = u_i$.

It is easy to see that this VA satisfies axioms A0–A4. This means that $D^{\downarrow i}(u_i)$ can be computed by the above local computation scheme, thus eliminating the need to compute the joint desirability function D explicitly by (2). Unfortunately, this does not reduces the complexity of our problem, since one of the valuations (D_{team}) is defined on the full frame U^n.

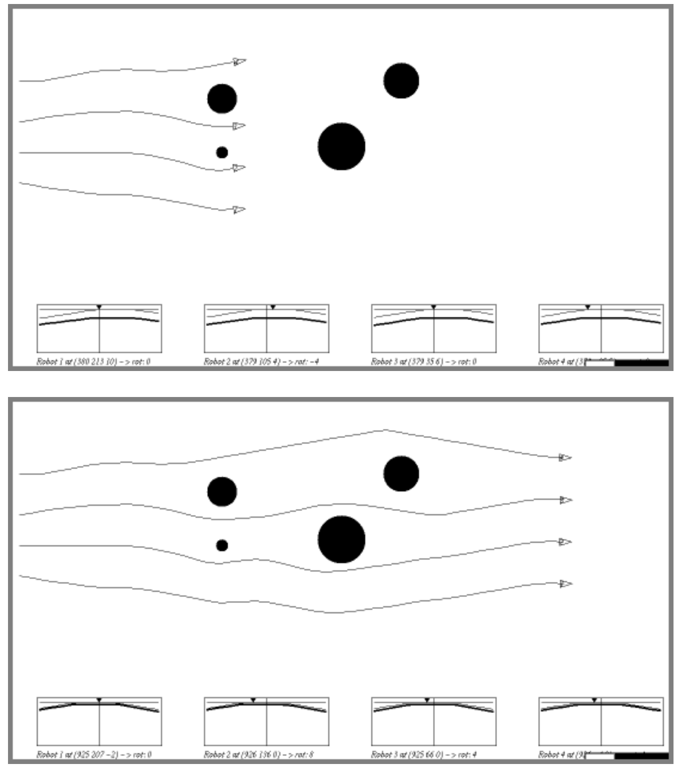

Figure 5: Combined obstacle avoidance and formation control with four robots using local computation

In order to profit from the local computation scheme, we need to decompose D_{team} into a set of smaller team desirability functions over subsets of robots. One way to do so is to represent the team objective by a set of local objectives for clusters of neighboring robots. In a four-robot formation control example, we can express the team objective by three pairwise desirability functions D^{12}_{team}, D^{23}_{team}, and D^{34}_{team}, each having the same shape as the D_{team} used in the two robot example above. Local computation of the VA propagates the impact of each local team desirability function to the entire team: this ensures, for instance, that the interactions between the individual objectives of $R1$ and $R4$ are correctly taken into account.

Fig. 5 shows a simulated run with four mobile robots in an obstacle field. Each mobile robot has the individual objectives to go east and to avoid the obstacles. In additions, the robots have the team objective of maintaining a line formation. In this example, we use local computation on the valuation algebra $\langle U^n, \mathcal{D}, \wedge, \downarrow \rangle$, where $\mathcal{D} = \{D_1, \ldots, D_n, D^{12}_{\text{team}}, D^{23}_{\text{team}}, D^{34}_{\text{team}}\}$. The four plots at the bottom of each image show the individual desirability functions of each robot before and after evaluation of the algebra by local computation. The small black triangles show the resulting steering control for each robot.

As it appears, the robots can maintain a global formation although the team objective is defined locally, as long as the spatial order of the robots is preserved. If we assume that each robot computes the preferred controls locally, then at each cycle the robots need to transmit a total of six desirability functions (two between each pair of neighboring robots) and to compute three combinations on the frame U^2 (one for each pair of neighboring robots). By contrast, global computation using the full D_{team} would require the overall transmission of 12 desirability functions and the computation of one combination on the full frame U^4.

5 Conclusions

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 computation of rational tradeoffs between competing objectives.

The desirability-function approach is essentially different from most common approaches to formation control, like those based on potential-fields, which summarize preferences solely in terms of a "most preferred" action, and combine these actions instead of combining the full desirability function (e.g., [1]). In the example shown in the Introduction, combining, in each robot, the most preferred action for avoiding obstacles with the most preferred action for keeping formation (here, zero steering) would lead each robot to steer toward the other one: obviously an undesirable choice in our case.

A potential drawback of the desirability-function approach is its computational complexity. In this paper, we have shown that desirability functions can be formalized in the framework of valuation algebras, thus allowing us to use local computations to dramatically reduce the amount of communication and computation resources needed for team coordination.

In our current work, we are studying the problem of how to find suitable decompositions of the team desirability function into local (sub-team) desirability functions for a given task.

Acknowledgements

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.

References

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

[2] Driankov, D. and Saffiotti, A., editors. Fuzzy Logic Techniques for Autonomous Vehicle Navigation. Phisica-Springer-Verlag, Heidelberg, DE, 2001. ISBN 3-7908-1341-9.

[3] Kohlas, J. and Shenoy, P.P. Computation in Valuation Algebras. In J. Kohlas and S. Moral, editors, *Handbook of Defeasible Reasoning and Uncertainty Management Systems*, Vol. 5. Kluwer Academic Press, Amsterdam (in press).

[4] Pirjanian, P. and Mataric, M. Optimality issues in fuzzy behavior coordination. In D. Driankov and A. Saffiotti, editors, *Fuzzy logic techniques for autonomous vehicle navigation* pp. 235–253. Phisica/Springer-Verlag, 2001.

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

[6] Saffiotti, A. Fuzzy Logic in Autonomous Robotics: behavior coordination. *Proc of the IEEE Intl Conf on Fuzzy Systems*, pp. 573–578. Barcelona, Spain, 1997.

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

[8] Saffiotti, A., Zumel, N.B., and Ruspini, E.H. Multi-robot team coordination using desirabilities. *Proc of the 6th Intl Conf on Inteligent Autonomous Systems* (IAS-6), pp. 107–114. Venice, Italy, 2000.

[9] Shenoy, P.P., and Shafer, G. Propagating Belief Functions with Local Computations. IEEE Expert, 1(3):43-52, 1986.

[10] Shenoy, P.P., and Shafer, G. Axioms for Probability and Belief-Function Propagation. In R. D. Shachter, T. Levitt, J. F. Lemmer and L. N. Kanal (eds.) *Uncertainty in Artificial Intelligence*, Vol. 4, pp. 169–198. North-Holland, Amsterdam, 1990.