Anticipation as a Key for Collaboration in a Team of Agents: A Case Study in Robotic Soccer

Manuela Veloso, Peter Stone, and Michael Bowling
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213, U.S.A.

ABSTRACT

We investigate *teams* of complete autonomous agents that can collaborate towards achieving precise objectives in an adversarial dynamic environment. We have pursued this work in the context of robotic soccer both in simulation and with real physical robots. We briefly present these two frameworks emphasizing their different technical challenges. Creating *effective* members of a team is a challenging research problem. We first address this issue by introducing a team architecture organization which allows for a rich task decomposition between team members. The main contribution of this paper is our introduction of an action-selection algorithm that allows for a teammate to *anticipate* the needs of other teammates. Anticipation is critical for maximizing the probability of successful collaboration in teams of agents. We show how our contribution applies to the two concrete robotic soccer frameworks and present controlled empirical results run in simulation. Anticipation was successfully used by both our CMUnited-98 simulator and CMUnited-98 small-robot teams in the RoboCup-98 competition. The two teams are RoboCup-98 world champions each in its own league.

1. INTRODUCTION

We have been pursuing research in the development of *teams* of autonomous agents that need to act in adversarial environments. In these domains, single agents cannot achieve the overall team goals individually. Goal achievement necessarily requires the collaboration between the members of the team. We have used three different testbeds in the robotic soccer domain to pursue this investigation: a rich simulation environment using the RoboCup soccer server, our own-built small wheeled robots, and Sony's fully autonomous legged robots. At the RoboCup-98 competitions we came in first place in each of these three leagues. This paper focuses on the collaborative teamwork algorithms of our CMUnited-98 simulation and small-robot teams.

Several other researchers have developed teamwork theories and opponent modelling in a variety of domains (e.g.[1,2]). One of the main focus of our research is on algorithms for collaboration between agents in a team. We can view our teamwork approach as a method to optimize social utility.[3] We build upon this work in the real-time adversarial robotic soccer domains. In our approach of social utility, an agent, as a member of a team, needs to be capable of individual autonomous decisions while, at the same time, its decisions must contribute towards the team goals.

In many multi-agent systems, one or a few agents are assigned, or assign themselves, the specific task to be solved at a particular moment. We view these agents as the *active* agents. Other team members are *passive* waiting to be needed to achieve some task. Concretely, in the robotic soccer domain, we view the agent that goes to the ball as the "active" agent, while the other teammates are in principle "passive." While the active agent has a clear task assigned and therefore a clear plan to follow (e.g. move towards the ball), it is less clear what is the plan for the passive agents. As the team agents most probably will need to collaborate, it seemed to us that passive agents could not simply be "passive."

Previously, in our CMUnited-97 teams, passive agents flexibly varied their positions within a range, however only as a function of the position of the ball.[4–6] In so doing, their goal was to *anticipate* where they would be most likely to find the ball in the near future. In our CMUnited-97 teams, both simulation and real robots, we effectively used this ball-dependent role-adjustment strategy. Since this strategy did not take into account other agents, we call this first-level of anticipation towards better individual goal achievement, *single-agent anticipation.*

Correspondence: mmv@cs.cmu.edu; http://www.cs.cmu.edu/~mmv.

Recently, we investigated a more elaborate team behavior for the passive agents. For the CMUnited-98 teams, we introduced a team-based notion of *multi-agent anticipation*, which goes beyond single-agent anticipation. The passive team agents position themselves strategically so as to optimize the chances that their teammates can successfully collaborate with them, in particular pass to them. By considering the positions of other agents and the attacking goal, in addition to that of the ball, they are able to position themselves more usefully: they *anticipate* their future contributions to the team, i.e. they anticipate their social utility. This strategic anticipation is the main contribution of this paper.

The paper is organized as follows. Section 2 describes the simulation and the robotic soccer frameworks. Section 4 contributes the anticipation algorithm as a key behavior for the success of team of agents, and reports on the results obtained at the RoboCup-98 robotic soccer competitions as well as controlled empirical tests. Section 5 concludes the paper.

2. SIMULATION AND REAL ROBOTIC SOCCER

Robotic soccer is a challenging domain for studying real-time multi-agent coordination techniques: agents must act quickly and autonomously while contributing to the achievement of the team's overall goal.[7,8]

2.1. Simulator

The RoboCup soccer server[9] has been used as the basis for successful international competitions and research challenges.[10] Though not directly based upon any single robotic system, the soccer server captures several real-world complexities: all players are controlled by separate processes; the players' vision is limited; the players can communicate by posting to a blackboard that is visible to all players; actions and sensors are noisy; each player has limited stamina; and play occurs in *real time*.

- the players' vision is limited;
- the players can communicate by posting to a blackboard that is visible to all players;
- all players are controlled by separate processes;
- each team has 11 members;
- each player has limited stamina;
- actions and sensors are noisy; and
- play occurs in *real time*: the agents must react to their sensory inputs at roughly the same speed as human or robotic soccer players.

The simulator, acting as a server, provides a domain and supports users who wish to build their own agents (clients).

By abstracting away the low-level perception and action complexities inherent in robotics, the simulator allows researchers to focus quickly on the multi-agent coordination issues. In this regard, the fact that the simulator enforces a completely distributed approach (each player must be controlled by a separate program) is a crucial feature.

Perception in the simulator is distributed. Each agent sees a portion of the world depending on the direction it is facing. All of the information it receives is in polar coordinates relative to its own position (instead of global Cartesian coordinates). Objects that are farther away are seen with less precision. Actions available to the clients are parameterized movement commands (turn/dash/kick) as well as a communicate "say" command. The effects of all actions are non-deterministic.*

2.2. Real Robots

Our small-size robot team is a fully autonomous system consisting of a global perception system and individual decision-making clients. The team is made up of five robots that we have built.[12] † Figure 1 shows our CMUnited-98 robots.

The decision-making clients select actions based on the perceived world state and send these commands to the physical robots through radio communication. Perception is accomplished through a camera over-looking the playing field. The full view of the world is processed by our vision algorithm to find and track the position and orientation

*For complete details of the simulator, see Stone's PhD thesis.[11]

†We thank Sorin Achim and Kwun Han for designing and building the CMUnited-98 robots.

Figure 1. Our Small Robot Team: RoboCup-98 Champions

of the team's agents, the position of the opponents, and the position and trajectory of the ball.[13] This processing provides a global view of the environment that is shared by all of the agents.

Each physical robot is controlled by a separate client program that makes decisions using the information obtained from the vision system. Each agent is assigned a role and therefore behaves differently. Our robots exhibit three roles: goal-tender, defender, and attacker. Even with this partition of agents, it is still not desirable for all of the agents to be "actively" filling their roles. For example, it is rarely successful for all of the attackers to chase the ball since they will often be hindering each other's progress. Multiple attackers need therefore to coordinate. Each attacking agent uses the perceived world to calculate the value of its own and its teammates' possible actions and acts accordingly.

3. SINGLE-AGENT ANTICIPATION

Our base teamwork structure is situated within a team member architecture suitable for real-time multi-agent domains in which individual agents can cooperate with teammates towards common goals while still acting autonomously.[4] Based on a standard agent paradigm, our team member architecture allows a complete intelligent agent cycle: (i) to sense the environment, (ii) to reason about and select their actions, and (iii) to act in the real world.

Individual agents act autonomously during the games. In the simulator, agents have limited communication capabilities, but in our real robots, agents cannot communicate with each other. In order to achieve the necessary team coordination towards the joint achievement of the team goals, we introduce a "locker-room agreement," as a collection of deliberative plans shared by all the agents. The locker-room agreement specifies team conventions and collaboration protocols (building upon equivalent approaches in multi-agent systems). It alleviates the need for negotiation during time-pressured situations and it allows for team coordination when no communication is available.

In our initial team architecture,[4] we include a behavior-based approach.[14] Agents play different *roles* and are organized in different *formations*. A formation consists primarily of a set of roles.[4] A formation consists of a set of positions and a set of units. The formation and each of the units can also specify inter-position behavior specifications for the member positions. The definition of a role includes *home coordinates*, a *home range*, and a *maximum range*. The home and max ranges of different roles can overlap, even if they are part of the same formations.

When an agent is not in control of the ball, our architecture allows for it to flexibly vary its position within its role. The role's home coordinates are the default location to which the agent should go. Within the flexible role assignment, rather than associating fixed (x, y) coordinates with each role, an agent filling a particular role is given a range of coordinates in which it could position itself. It is this flexibility that is exploited by the anticipation mechanism. In particular, based on the ball's position on the field, the agent positions itself so as to increase its chances of achieving the team goal.

In this initial team architecture approach, agents can vary their position within their home range, only as a function of the position of the ball. The position of the other agents is not used to determine each agent's own position, therefore the agent performs what we call *single-agent anticipation*. When reacting to the ball's position, the agent moves to a location within its range that minimizes its distance to the ball. In so doing, the agent's goal is to *anticipate* where they would be most likely to find the ball. This single-agent anticipation aims towards a better individual goal achievement. It was successfully used in our CMUnited-97 simulation and small-robots teams.[5,6]

4. ANTICIPATION FOR TEAM COLLABORATION

The social utility of agents in a team of agents can be in general very hard to define. We are particularly interested in opportunities that allow the agents to collaborate with each other. The assumption is that collaboration is needed when a task cannot be accomplished by individual agents.

Our anticipation approach for team collaboration presented in this section in the robotic soccer domain, could be in principle generalized.

Consider that for each agent, for each state, and at each time, there is a computable value for the probability that an active agent could successfully collaborate with a passive agent. As the world is constantly changing, the values for the probability of collaboration are computed as a function of the dynamic world.

Assuming that the transitions between states for each agent take time (or some other cost function), then anticipation consists of the selection of a new state that maximizes the probability of future collaboration.

Anticipation therefore allows for a flexible adjustment of a team agent to maximize the probability of its social utility, i.e., to be useful for the team. We now formally present our anticipation algorithm within the robotic soccer domain.

4.1. Anticipation in a Robotic Soccer Team

The CMUnited-98 small robot team significantly extends the initial CMUnited-97 single-agent anticipation in two ways: (i) we maintain the role assignments, but we don't use pre-defined formations, (ii) we use a decision theoretic algorithm to select the active agent; and (ii) we use a technique for passive agents *to anticipate* future *collaboration*. Passive agents are therefore not actually "passive;" instead, they actively *anticipate* opportunities for collaboration.

Collaboration is built on robust individual behaviors. We first developed individual behaviors for passing and shooting. Passing and shooting in CMUnited-98 is handled effectively by the motion controller.[15] The target configuration is specified to be the ball (using its estimated trajectory) and the target direction is either towards the goal or another teammate. This gives us robust and accurate individual behaviors that can handle obstacles as well as intercepting a moving ball.

4.2. Decision Theoretic Action Selection

Given the agents' individual behaviors, we must select an active agent and appropriate behavior. This choice is made using a decision theoretic analysis using a single step look-ahead. With n agents this amounts to n^2 choices of actions involving shooting or a pass to another agent followed by that agent shooting. An estimated probability of success for each pass and shot is computed along with the time estimate to complete the action, which is provided by the motion controller. A value for each action is computed:

$$\text{Value} = \frac{\text{Pr}_{\text{pass}}\text{Pr}_{\text{shoot}}}{\text{time}}.$$

The action with the largest value is selected, which determines both the active agent and its behavior. Table 1 illustrates an example of the values for the selection considering two attackers, 1 and 2.

		Probability of Success			
Attacker	Action	Pass	Shoot	Time(s)	Value
1	Shoot	–	60%	2.0	0.30
1*	Pass to 2	60%	90%	1.0	0.54
2	Shoot	–	80%	1.5	0.53
2	Pass to 1	50%	40%	0.8	0.25

Table 1. Action choices and computed values are based on the probability of success and estimate of time. The largest-valued action (marked with an *) is selected.

This decision theoretic action selection is implemented on our real robots. It is important to note that this action selection is occurring on each iteration of control, i.e., approximately 30 times per second. The probabilities of success, estimates of time, and values of actions, are being continuously recomputed. This allows for quick changes of actions if shooting opportunities become available or collaboration with another agent appears more useful.

4.3. Dynamic Positioning: SPAR

Although the above technique determines what action should be taken by the active agent, it is unclear what the passive agents should be doing. However, in a team multi-agent system such as robotic soccer, success and goal achievement often depend upon collaboration. Therefore, we introduce in CMUnited-98, the concept that team agents should not actually be "passive."

Here we introduce a team-based notion of *anticipation*, which goes beyond individual single-agent anticipation. The passive team members position themselves strategically so as to optimize the chances that their teammates can successfully collaborate with them, in particular pass to them. By considering the positions of other agents and the attacking goal, in addition to that of the ball, they are able to position themselves more usefully: they *anticipate* their future contributions to the team.

This strategic positioning takes into account the position of the other robots (teammates and opponents), the ball, and the opponent's goal. The best position is determined by the solution to a multiple-objective function with repulsion and attraction points. To present our algorithm, we introduce the following variables:

- n - the number of agents on each team;
- O_i - the current position of each opponent, $i = 1, \ldots, n$;
- T_i - the current position of each teammate, $i = 1, \ldots, (n-1)$;
- B - the current position of the active teammate and ball;
- G - the position of the opponent's goal;
- P - the desired position for the passive agent in anticipation of a pass.

Using these variables, we formalize our algorithm for strategic positioning which we call SPAR for *Strategic Positioning with Attraction and Repulsion.* SPAR extends similar approaches using potential fields[16] to our highly-dynamic, multi-agent domain. The probability of collaboration is directly related to how "open" a position is to allow for a successful pass. SPAR maximizes the repulsion from other agents and minimizes attraction to the ball and to the goal. Using d as the *Euclidean distance* function, SPAR considers the following set of objectives:

- *Repulsion* from opponents. Maximize the distance to each opponent: $\forall i, \max d(P, O_i)$.
- *Repulsion* from teammates. Maximize the distance to other passive teammates: $\forall i, \max d(P, T_i)$.
- *Attraction* to the ball: $\min d(P, B)$.
- *Attraction* to the opponent's goal: $\min d(P, G)$.

This is a multiple-objective function. To solve this optimization problem, we restate this function into a single-objective function.

As each term in the multiple-objective function may have a different relevance (e.g., staying close to the goal may be more important than staying away from opponents), we want to consider different functions of each term. In our CMUnited-98 teams, we weight the terms differently, namely w_{O_i}, w_{T_i}, w_B, and w_G, for the weights for opponents, teammates, the ball, and the goal, respectively. For CMUnited-98, these weights were hand tuned to create a proper balance. This gives us a weighted single-objective function:

$$\max\left(\sum_{i=1}^{n} w_{O_i} dist(P, O_i) + \sum_{i=1}^{n} w_{T_i} dist(P, T_i) - w_B\, dist(P, B) - w_G\, dist(P, G) \right)$$

This optimization problem is then solved numerically under constraints which are specific to each team environment. We now present the set of constraints for both the simulator and the small-robot teams. We used the SPAR algorithm in both of these platforms.

4.4. Constraints in the Simulator Team

One constraint in the simulator team relates to the position, or role, that the passive agent is playing relative to the position of the ball. The agent only considers positions that are within a rectangle whose corner is on the ball that is closest to the position home of the position that it is currently playing. This constraint helps ensure that the player with the ball will have several different passing options in different parts of the field. In addition, players don't need to consider moving too far from their positions to support the ball.

In addition to this first constraint, the agents observe three additional constraints. In total, the constraints in the simulator team are:

- Stay in an area near home position
- Stay within the field boundaries
- Avoid being in an off-sides position
- Stay in a position in which it would be possible to receive a pass.

This last constraint is evaluated by checking that there are no opponents in a cone with vertex at the ball and extending to the point in consideration.

4.5. Constraints in the Real Robot Team

The real robotic environment is different from the simulator in several aspects. The main differences result from the fact that the agents here are real physical artifacts that occupy space, do not control the ball carefully, and have rather unreliable accurate motion. The robots' motion is quite crude manipulating the ball, as they cannot hold it, but act on the ball simply by pushing it with their mostly flat sides.

The constraints in the real robot team therefore involve two main aspects on the position of the passive agent:

- Do not block possible direct shot from active teammate to the goal.
- Do not stand behind other robots, because these are difficult positions to receive passes in case the active team decides to do so.

Figures 2(a) and (b) illustrate these two constraints and Figure 2(c) shows the combination of these two constraints and the resulting position returned by our algorithm for the anticipating passive teammate.

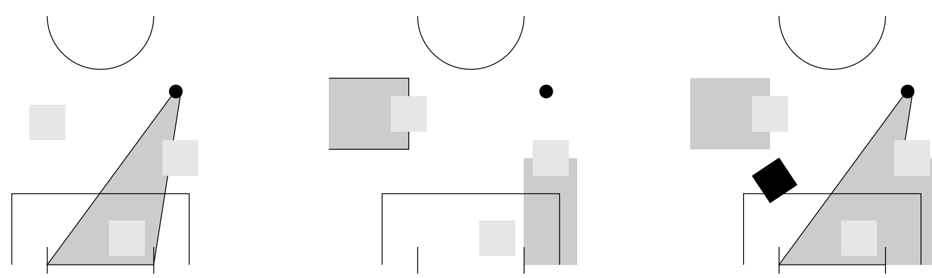

(a) Do not block goal shot (b) Avoid difficult collaboration (c) Result

Figure 2. Constraints for the anticipation algorithm for the CMUnited-98 small-robot team: (a) and (b) show three opponents robots, and the position of the ball corresponding also to the position of the active teammate; (c) shows the resulting position of the passive agent, dark square, as returned by our anticipation algorithm, using the constraints in (a) and (b).

Using this anticipation algorithm in competition, the attacking team agents behaved in an exemplary collaborative fashion. Their motion on the field was a beautiful response to the dynamically changing adversarial environment. The active and passive agents moved in coordination using the anticipation algorithm increasing very significantly successful collaboration.

4.6. Results

In this paper, we provide the results of all of our real RoboCup-98 games as anecdotal evidence of the performance of our algorithm. Note that the environments are highly adversarial and that our agents had never seen the opponent teams before. We also report on the results of extensive, controlled empirical tests run in simulation.

Tables 2 and 3 show the scores of the games of CMUnited-98 simulator and real robot teams. Both the simulator and small-robot CMUnited-98 teams won their respective leagues at RoboCup-98. The simulator and small-robot teams placed first out of 34 and 11 teams, respectively.

Opponent Name	Affiliation	Score
UU	Utrecht University	22–0
TUM / TUMSA	Technical University Munich	2–0
Kasuga-Bitos II	Chubu University	5–0
Andhill'98	NEC	8–0
ISIS	Information Sciences Institute	12–0
Rolling Brains	Johannes Gutenberg-University	13–0
Windmill Wanderers	University of Amsterdam	1–0
AT-Humboldt'98	Humboldt University of Berlin	3–0
TOTAL		66–0

Table 2. The scores of CMUnited-98's games in the simulator league of RoboCup-98.

Opponent Name	Affiliation	Score
iXS	iXs Inc.	16–2
5DPO	University of Porto	0–3
Paris-8	University of Paris-8	3–0
Cambridge	University of Cambridge	3–0
Roboroos	University of Queensland	3–1
TOTAL		25–6

Table 3. The scores of CMUnited-98's games in the small-robot league of RoboCup-98.

We find that the RoboCup games offered a fully unbiased evaluation setup, as the dynamics of the environment was unknown and not modeled. Our anticipation algorithm clearly provided a significant advantage over the other teams. Video clips of our games show the dynamic movement of our robots as a function of the position of the other robots and of the ball. Using our SPAR algorithm, the passive agents intelligently move to open areas, hence responding in real-time to the highly dynamic character of the robotic soccer domain. In the games, robots position themselves leading to successful collaboration opportunities as ball passes between robots towards goal scoring.

As a controlled experiment, we tested teams using single-agent (ball-dependent) and multi-agent (SPAR) anticipation against a common opponent in an offense-defense scenario: 6 attackers against 4 defenders with the ball being reset whenever the attackers scored or the defenders cleared the ball. As shown in Table 4, when the attackers used SPAR, they scored significantly more frequently than when they used ball-dependent positioning. Each experiment was run for 200,000 simulator cycles or 333.3 minutes of real time (1 cycle = 100 msec).

Anticipation	Goals	Mean cycles between goals	Std. Dev.
Single-agent	191	1048.45	$\pm$ 71.2
Multi-agent	272	739.64	$\pm$ 45.0

Table 4. A comparison of single-agent (ball-dependent) and multi-agent (SPAR) anticipation strategies against a common opponent over 200,000 simulator cycles.

5. CONCLUSION

In this paper we have presented our work investigating and introducing the concept of anticipation to increase collaboration in teams of agents. We have been working in the concrete domain of robotic soccer.

We discussed our basic flexible team architecture in which individual agents flexibly move within their role as a function of the position of the ball to try to optimize their individual actions towards the ball.

We then introduced a novel anticipation algorithm that allows for team agents to strategically position themselves in anticipation of possible collaboration needs from other teammates. In simulation, we demonstrated empirically that our novel multi-agent anticipation strategy significant improves over our previous single-agent anticipation strategy. We used this new multi-agent anticipation algorithm in our RoboCup-98 teams achieving very successful results.

Given that we have developed algorithms for individual action and team collaboration through anticipation, the current on-going step on our research agenda is to develop algorithms to predict the actions of the opponent agents and proactively counter them.

Acknowledgments

This research was sponsored in part by the Defense Advanced Research Projects Agency (DARPA), and Rome Laboratory, Air Force Materiel Command, USAF, under agreement numbers F30602-97-2-0250 and F30602-98-2-0135. Views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Defense Advanced Research Projects Agency (DARPA), the Air Force Research Laboratory (AFRL) or the U.S. Government.

REFERENCES

1. B. Grosz and S. Kraus, "Collaborative plans for complex group actions," *Artificial Intelligence* **86**, pp. 269–368, 1996.
2. M. Tambe, "Towards flexible teamwork," *Journal of Artificial Intelligence Research* **7**, pp. 81–124, 1997.
3. C. Goldman and J. Rosenschein, "Emergent coordination through the use of cooperative state-changing rules," in *Proceedings of the Twelfth National Conference on Artificial Intelligence*, pp. 408–413, Morgan Kaufman, (Philadelphia, PA), 1994.
4. P. Stone and M. Veloso, "Task decomposition, dynamic role assignment, and low-bandwidth communication for real-time strategic teamwork," *Artificial Intelligence* **110**, pp. 241–273, June 1999.
5. M. Veloso, P. Stone, and K. Han, "CMUnited-97: RoboCup-97 small-robot world champion team," *AI Magazine* **19**(3), pp. 61–69, 1998.
6. P. Stone and M. Veloso, "The CMUnited-97 simulator team," in *RoboCup-97: Robot Soccer World Cup I*, H. Kitano, ed., Springer Verlag, Berlin, 1998.
7. H. Kitano, Y. Kuniyoshi, I. Noda, M. Asada, H. Matsubara, and E. Osawa, "RoboCup: A challenge problem for AI," *AI Magazine* **18**, pp. 73–85, Spring 1997.
8. A. K. Mackworth, "On seeing robots," in *Computer Vision: Systems, Theory, and Applications*, A. Basu and X. Li, eds., pp. 1–13, World Scientific Press, Singapore, 1993.
9. I. Noda, H. Matsubara, K. Hiraki, and I. Frank, "Soccer server: A tool for research on multiagent systems," *Applied Artificial Intelligence* **12**, pp. 233–250, 1998.
10. H. Kitano, M. Tambe, P. Stone, M. Veloso, S. Coradeschi, E. Osawa, H. Matsubara, I. Noda, and M. Asada, "The RoboCup synthetic agent challenge 97," in *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence*, pp. 24–29, Morgan Kaufmann, (San Francisco, CA), 1997.
11. P. Stone, *Layered Learning in Multi-Agent Systems*. PhD thesis, Computer Science Department, Carnegie Mellon University, Pittsburgh, PA, December 1998. Available as technical report CMU-CS-98-187.
12. M. Veloso, M. Bowling, S. Achim, K. Han, and P. Stone, "The CMUnited-98 champion small robot team," in *RoboCup-98: Robot Soccer World Cup II*, M. Asada and H. Kitano, eds., Springer Verlag, Berlin, 1999.
13. K. Han and M. Veloso, "Reactive visual control of multiple non-holonomic robotic agents," in *Proceedings of the International Conference on Robotics and Automation*, (Leuven,Belgium), May 1998.
14. M. J. Mataric, "Interaction and intelligent behavior," MIT EECS PhD Thesis AITR–1495, MIT AI Lab, August 1994.
15. M. Bowling and M. Veloso, "Motion control in CMUnited-98," in *RoboCup-99: Robot Soccer World Cup III*, M. Veloso, E. Pagello, and H. Kitano, eds., Springer Verlag, Berlin, 2000. To appear.
16. J.-C. Latombe, *Robot Motion Planning*, Kluwer Academic Publishers, Boston, 1991.