The CMUnited-98 Champion Small-Robot Team

Manuela Veloso, Michael Bowling, Peter Stone
Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213
{veloso,mhb,pstone}@cs.cmu.edu
http://www.cs.cmu.edu/˜robosoccer

December 18, 2002

Abstract

Robotic soccer presents a large spectrum of challenging research opportunities. In this article, we present the main research and technical contributions of our champion CMUnited-98 small robot team. The team is a multiagent robotic system with global perception, and distributed cognition and action. We describe the main features of the hardware design of the physical robots, including differential drive, robust mechanical structure, and a kicking device. We introduce our new robot motion algorithm which reactively generates motion control to account for the target point, the desired robot orientation, and obstacle avoidance. Our robots exhibit successful collision-free motion in the highly dynamic robotic soccer environment. At the strategic and decision-making level, we present the role-based behaviors of the CMUnited-98 robotic agents. Team collaboration is remarkably achieved through a new algorithm that allows for team agents to anticipate possible collaboration opportunities. Robots position themselves strategically in open positions that increase passing opportunities. The article terminates with a summary of the results of the RoboCup-98 games in which the CMUnited-98 small robot team scored a total of 25 goals and suffered 6 goals in the 5 games that it played.

Keywords: Mobile robots, Real-time decision making, Reactive motion control, Multiple collaborating and competing agents.

[0]This research is sponsored in part by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement numbers F30602-97-2-0250 and F30602-98-2-0135, and in part by the Department of the Navy, Office of Naval Research under contract number N00014-95-1-0591. The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing official policies or endorsements, either expressed or implied, of the Air Force, of the Department of the Navy, Office of Naval Research or the United States Government.

1 Introduction

The CMUnited-98 small-size robot team is a complete, autonomous architecture composed of the physical robotic agents, a global vision processing camera over-looking the playing field, and several clients as the minds of the small-size robot players. Fig. 1 sketches the building blocks of the architecture.

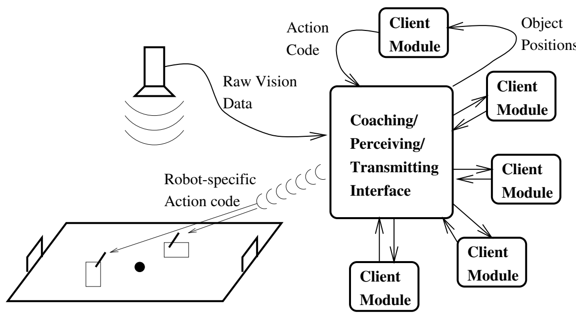

Figure 1: The CMUnited architecture with global perception and distributed reaction.

The complete system is fully autonomous consisting of a well-defined and challenging processing cycle. The global vision algorithm perceives the dynamic environment and processes the images, giving the positions of each robot and the ball. This information is sent to an off-board controller and distributed to the different agent algorithms. Each agent evaluates the world state and uses its strategic knowledge to make decisions. Actions are motion commands that are sent by the off-board controller through radio frequency communication. Commands can be broadcast or sent directly to individual agents. Each robot has an identification binary code that is used on-board to detect commands intended for that robot. Motion is not perfectly executed due to inherent mechanical inaccuracies and unforeseen interventions from other agents. The effects of the actions are therefore uncertain.

The physical robots themselves are of size $15\text{cm} \times 12\text{cm} \times 10\text{cm}$. Fig. 2 shows our robots. A differential drive mechanism is used in all of the robots. Two motors with integrated gear boxes are used for the two wheels. Differential drive was chosen due to its simplicity and due to the size constraints. The size of our robots conforms to RoboCup Competition rules[1]. Employing the differential drive mechanism means that the robot is non-holonomic, which makes the robot control problem considerably more challenging.

Although it may be possible to fit an on-board vision system onto robots of small size, in the interest of being able to quickly move on to strategic multiagent issues, the CMUnited-98 teams uses a global vision system. The fact that perception is achieved by a video camera overlooking the complete field offers an opportunity to get a global view of the world state. This setup may simplify the sharing of information among multiple agents, but it also presents a challenge for reliable and real-time processing of the movement of multiple mobile objects, namely the ball, five robots on our team, and five robots on the opponent's team [3, 4, 1].

This article presents the main technical contributions of our CMUnited-98 small robot team. It focuses on the problems of motion control and the robots' strategy. The vision processing for the

[1] see http://www.robocup.org/RoboCup/

Figure 2: The CMUnited-98 robots.

team is the same system used in the CMUnited-97 team [8], and so is not described here. Section 2 presents the motion planning approach for our robots including path planning to intercept moving targets and obstacle avoidance. Section 3 introduces the individual and team behaviors of the CMUnited-98 robots. We introduce the novel concept of *anticipation* which allows for the robots to effectively receive passes from teammates. Section 4 summarizes the results of the RoboCup-98 games and Section 5 draws conclusions.

2 Motion Control

The goal of our low level motion control is to be as fast as possible while remaining accurate and reliable. This is challenging due to the lack of feedback from the motors, forcing all control to be done using only visual feedback. Our motion control algorithm is robust. It addresses stationary and moving targets with integrated obstacle avoidance. The algorithm makes effective use of the prediction of the ball's trajectory provided by the Kalman-Bucy filter.

We achieve this motion control functionality by a reactive control mechanism that directs a differential drive robot to a target configuration. Though based on the CMUnited-97's motion control [8], CMUnited-98 includes a number of major improvements. The target configuration for the motion planner has been extended. The target configuration includes: (i) the *Cartesian position*; and (ii) the *direction* that the robot is required to be facing when arriving at the target position. Obstacle avoidance is integrated into this controller. Also, the target configuration can be given as a function of time to allow for the controller to reason about intercepting the trajectory of a moving target.

2.1 Differential Drive Control for Position and Direction

CMUnited- 98's basic control rules were improved from those used in CMUnited- 97. The rules are a set of reactive equations for deriving the left and right wheel velocities, v_l and v_r, in order to reach a target position, (x^*, y^*):

$$\begin{aligned} \Delta &= \theta - \phi \qquad (1) \\ (t, r) &= (\cos^2 \Delta \cdot \text{sgn}(\cos \Delta), \sin^2 \Delta \cdot \text{sgn}(\sin \Delta)) \end{aligned}$$

$$
\begin{aligned}
v_l &= v(t-r) \\
v_r &= v(t+r),
\end{aligned}
$$

where θ is the direction of the target point (x^*, y^*), ϕ is the robot's orientation, and v is the desired speed (see Fig. 3(a))[2]. A few aspects of these equations deserve explanation. The use of $\sin^2$ and $\cos^2$ restricts the values $(t \pm r)$ to the interval $[0, 1]$, which bounds the magnitude of the computed wheel velocities by v. These equations also do not necessarily drive the robot forward, possibly driving the robot backwards towards the target.

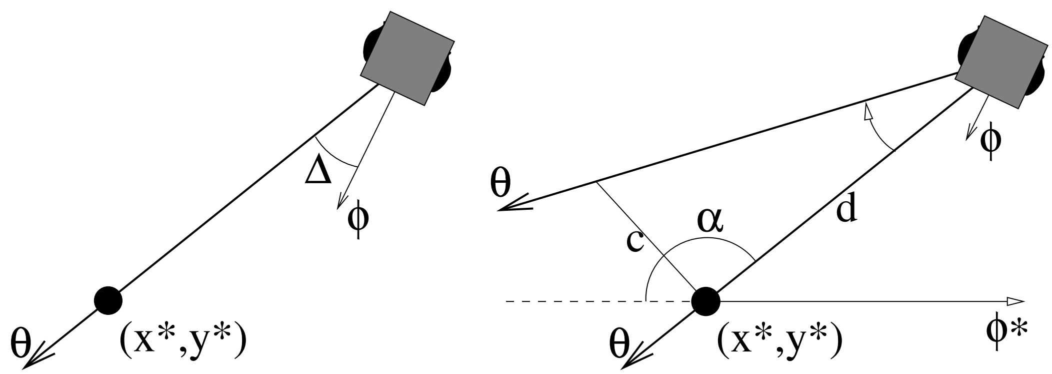

(a) (b)

Figure 3: (a) The parameters used to reach a target configuration (x^*, y^*), without a specified target orientation. (b) The adjustment of θ to θ' to reach a target configuration of the form (x^*, y^*, ϕ^*).

We extend these equations for target configurations of the form (x^*, y^*, ϕ^*), where the goal is for the robot to reach the specified target point (x^*, y^*) while facing the direction ϕ^*. This is achieved with the following adjustment:

$$
\theta' = \theta + \min\left(\alpha, \tan^{-1}\left(\frac{c}{d}\right)\right),
$$

where θ' is the new target direction, α is the difference between our angle to the target point and ϕ^*, d is the distance to the target point, and c is a clearance parameter (see Fig. 3(b).) This will keep the robot a distance c from the target point while it is circling to line up with the target direction, ϕ^*. This new target direction, θ', is now substituted into equation 1 to derive wheel velocities.

In addition to our motion controller computing the desired wheel velocities, it also returns an estimate of the time to reach the target configuration, $\hat{T}(x^*, y^*, \phi^*)$. This estimate is a crucial component in our robot's strategy. It is used both in high-level decision making, and for low-level ball interception, which is described later in this section. For CMUnited-98, $\hat{T}(x^*, y^*, \phi^*)$ is computed using a very simple linear function of d, α, and Δ:

$$
\hat{T}(x^*, y^*, \phi^*) = w_d d + w_\alpha \alpha + w_\Delta \Delta.
$$

The weights were set by simple empirical measurements. w_d is the inverse of the robot's translational speed; w_Δ is the inverse of the robot's rotational speed; and w_α is the inverse of the speed

[2]All angles are measured with respect to a fixed coordinate system.

of the robot when traversing a circle of radius, c. It is interesting to note that even this crude time estimate can be incredibly useful for building more complex behaviors, which are discussed later in this article.

2.2 Obstacle Avoidance

Obstacle avoidance was also integrated into the motion control. This is done by adjusting the target direction of the robot based on any immediate obstacles in its path. This adjustment can be seen in Fig. 4(b).

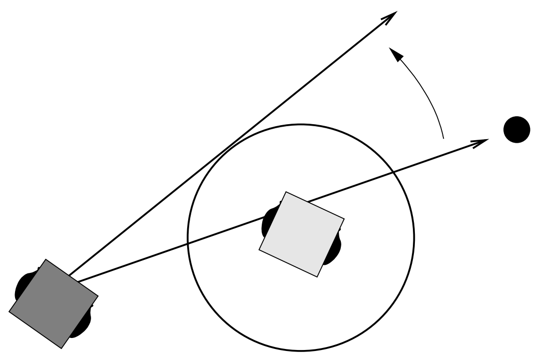

Figure 4: The adjustment of θ' to θ'' to avoid immediate obstacles.

If a target direction passes too close to an obstacle, the direction is adjusted to run tangent to the a preset allowed clearance for obstacles. Since the motion control mechanism is running continuously, the obstacle analysis is constantly replanning obstacle-free paths. This continuous replanning allows for the robot to handle the highly dynamic environment and immediately take advantage of short lived opportunities.

2.3 Moving Targets

One of the real challenges in robotic soccer is to be able to control the robots to intercept a moving ball. This capability is essential for a high-level ball passing behavior. CMUnited-98's robots successfully intercept a moving ball and several of their goals in RoboCup-98 were scored using this capability.

This interception capability is achieved as an extension of the control algorithm to aim at a stationary target. Fig. 5(a) illustrates the control path to reach a stationary target with a specific direction, using the control mechanism described above. Our extension allows for the target configuration to be given as a function of time, where $t = 0$ corresponds to the present,

$$f(t) = (x^*, y^*, \phi^*).$$

At some point in the future, t_0, we can compute the target configuration, $f(t_0)$. We can also use our control rules for a stationary point to find the wheel velocities and estimated time to reach this hypothetical target as if it were stationary. The time estimate to reach the target then informs us whether it is possible to reach it within the allotted time. Our goal is to find the nearest point in the future where the target can be reached. Formally, we want to find,

$$t^* = \min\{t > 0 : \hat{T}(f(t)) \le t\}.$$

After finding t^*, we can use our stationary control rules to reach $f(t^*)$. In addition we scale the robot speed so to cross the target point at exactly t^*.

Unfortunately, t^*, cannot be easily computed within a reasonable time-frame. We approximate this value, t^*, by discretizing time with a small time-step. We then find the smallest of these discretized time points that satisfies our estimate constraint. An example of this is shown in Fig. 5(b), where the goal is to hit the moving ball.

(a) (b)

Figure 5: (a) Control for stationary target. (b) Control for moving target.

The target configuration as a function of time is computed using the ball's predicted trajectory. Our control algorithm for stationary points is then used to find a path and time estimate for each discretized point along this trajectory, and the appropriate target point is selected.

3 Strategy

The main focus of our research is on developing algorithms for collaboration between agents in a team. An agent, as a member of the team, needs to be capable of individual autonomous decisions while, at the same time, its decisions must contribute towards the team goals.

CMUnited-97 introduced a flexible team architecture in which agents are organized in *formations* and *units*. Each agent plays a *role* in a unit and in a formation [5, 8]. CMUnited-98 builds upon this team architecture by defining a set of roles for the agents. It also introduces improvements within this architecture to help address the highly dynamic environment.

CMUnited-98 uses the following roles: goalkeeper, defender, and attacker. The formation used throughout RoboCup-98 involved a single goalkeeper and defender, and three attackers.

3.1 Goalkeeper

The ideal goalie behavior is to reach the expected entry point of the ball in the goal *before* the ball reaches it. Assuming that the prediction of the ball trajectory is correct and the robot has a uniform movement, we can state the ideal goalie behavior: given the predicted v_g and v_b as the velocities of the goalie and of the ball respectively, and d_g and d_b as the distances from the goalie and the ball to the predicted entry point, then, we want $\frac{d_g}{v_g} = \frac{d_b}{v_b} - \epsilon$, where ϵ is a small positive value to account for the goalie reaching the entry point slightly before the ball.

Unfortunately, the ball easily changes velocity and the movement of the robot is not uniform and is uncertain. Therefore we have followed a switching behavior for the goalie based on a threshold of the ball's estimated trajectory.

If the ball's estimated speed is higher than a preset threshold, the goalie moves directly to the ball's predicted entry goal point. Otherwise, the goalie selects the position that minimizes the largest portion of unobstructed goal area. This is done by finding the location that bisects the angles of the ball and the goal posts as is illustrated in Fig. 6.

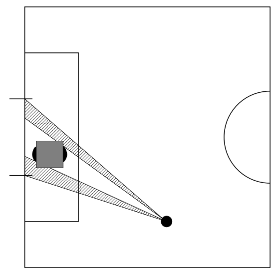

Figure 6: The goalkeeper positions itself to minimize the unobstructed goal area.

The use of the predicted ball's velocity for the goalie's behavior was shown to be very effective in the RoboCup-98 games. It was particularly appropriate for defending a penalty shot, due to the accuracy of the predicted ball's trajectory when only one robot is pushing the ball.

3.2 Defender

The CMUnited-97's team did not have a well-specified defender's role, but our experience at RoboCup-97 made us understand that the purpose of a defending behavior is two-fold:

1. to stop the opponents from scoring in our goal; and
2. to not endanger our own goal.

The first goal is clearly a defender's role. The second goal comes as the result of the uncertain ball handling by the robots. The robots can easily push (or touch) the ball unexpectedly in the wrong direction when performing a difficult maneuver.

To achieve the two goals, we implemented three behaviors for the defender. *Blocking*, illustrated in Fig. 7(a), is similar to the goalkeeper's behavior except that the defender positions itself further away from the goal line. *Clearing*, illustrated in Fig. 7(b), pushes the ball out of the defending area. It does this by finding the largest angular direction free of obstacles (opponents and teammates) that the robot can push the ball towards. *Annoying*, illustrated in Fig. 7(c), is somewhat similar to the goalkeeping behavior except that the robot tries to position itself between the ball and the opponent nearest to it. This is an effort to keep the opponent from reaching the ball.

Selecting when each of these behaviors is used is very important to the effectiveness of the defender. For example, clearing the ball when it is close to our own goal or when it can bounce

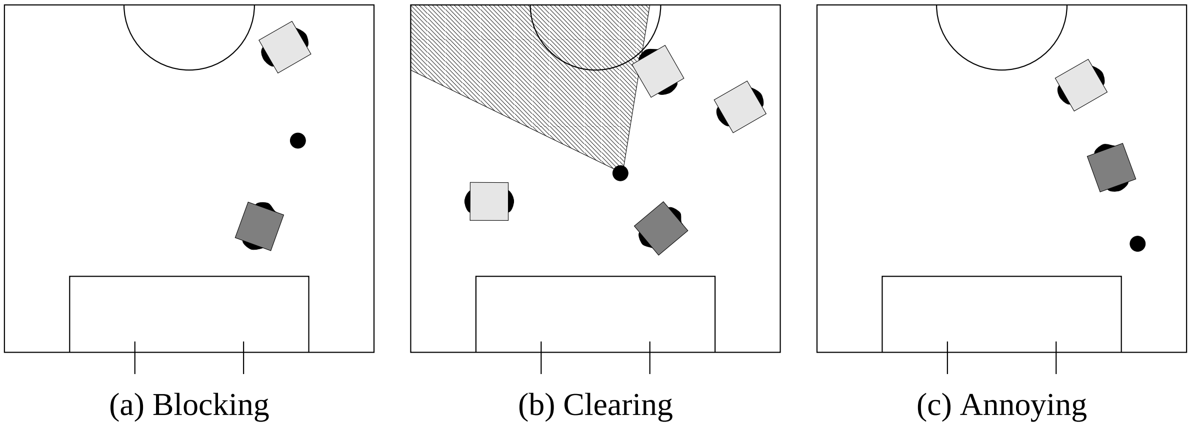

Figure 7: The defender's behaviors. The dark and light robots represent the defender and the opponents respectively.

back off another robot, can lead to scoring in our own goal. We used the decision tree in Fig. 8 to select which action to perform based on the current state.

The two attributes in the tree, namely *Ball Upfield* and *Safe to Clear*, are binary. *Ball Upfield* tests whether the ball is upfield (towards the opponent's goal) of the defender. *Safe to Clear* tests whether the open area is larger than a preset angle threshold. If *Ball Upfield* is false then the ball is closer to the goal than the defender and the robot *annoys* the attacking robot. The CMUnited-98's annoying behavior needs to select one particular opponent robot to annoy. For example, when two opponent robots attack simultaneously, the current annoying behavior is able to annoy only one of them. We are planning on further improving this behavior for RoboCup-99.

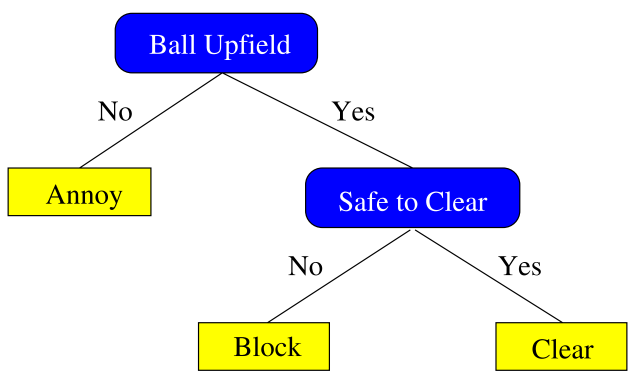

Figure 8: The decision tree heuristic used by the defender to select its behavior.

If *Ball Upfield* is true, the defender clears or blocks, depending on the value of *Safe to Clear*. Clearing was shown to be very useful at RoboCup-98, with even a couple of our goals scored directly by a clearing action of the defender.

3.3 Attackers - Active Teammate and Anticipation

Attacking involves one of the best opportunities for collaboration, and much of the innovation of CMUnited-98 has been developing techniques for finding and exploiting these opportunities.

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 another task or assist the active agent(s). This simplistic distinction between active and passive agents to capture teamwork was realized in CMUnited-97. The agent that goes to the ball is viewed as the active agent, while the other teammates are passive.

CMUnited-98 significantly extends this simplistic view in two ways: (i) 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. In CMUnited-98 this collaboration is built on robust individual behaviors.

3.3.1 Individual Behaviors.

We first developed individual behaviors for passing and shooting. Passing and shooting in CMUnited-98 is handled effectively by the motion controller. 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.

3.3.2 Decision Theoretic Action Selection.

Given the individual behaviors, we must select an active agent and appropriate behavior. This is done by 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.

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.

3.3.3 Dynamic Positioning (SPAR).

Although there is a clear action to be taken by the active agent, it is unclear what the passive agents should be doing. Although, in a team multiagent system such as robotic soccer, success and goal

Attacker	Action	Probability of Success 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.

achievement often depends upon collaboration; so, we introduce in CMUnited-98, the concept that team agents should not actually be "passive."

CMUnited-97's team architecture allowed for the passive agents to flexibly vary their positions within their role only as a function of the position of the ball. In so doing, their goal was to *anticipate* where they would be most likely to find the ball in the near future. This is a first-level of single-agent anticipation towards a better individual goal achievement [7].

However, for CMUnited-98, 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 position takes into account the position of the other robots (teammates and opponents), the ball, and the opponent's goal. The position is found as the solution to a multiple-objective function with repulsion and attraction points. Let's 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.

Given these defined variables, we can then formalize our algorithm for strategic position, which we call SPAR for *Strategic Positioning with Attraction and Repulsion.* This extends similar approaches using potential fields [2], 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 robots and minimizes attraction to the ball and to the goal, namely:

- *Repulsion* from opponents. Maximize the distance to each opponent: $\forall i, \max dist(P, O_i)$.

- *Repulsion* from teammates. Maximize the distance to other passive teammates: $\forall i, \max dist(P, T_i)$.

- *Attraction* to the ball: $\min dist(P, B)$.

- *Attraction* to the opponent's goal: $\min dist(P, G)$.

This is a multiple-objective function. To solve this optimization problem, we restate this function into a single-objective function. This approach has also been applied to the CMUnited-98 simulator team [6].

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 team, 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 under a set of constraints:

- Do not block a possible direct shot from active teammate.

- Do not stand behind other robots, because these are difficult positions to receive passes from the active teammate.

The solution to this optimization problem under constraints gives us a target location for the "passive" agent. Fig. 9(a) and (b) illustrate these two sets of constraints and Fig. 9(c) shows the combination of these constraints and the resulting position of the anticipating passive teammate.

This positioning was very effective for CMUnited-98. The attacking robots very effectively and dynamically adapted to the positioning of the other robots. The SPAR anticipation algorithm created a number of opportunities for passes and rebounds that often led to goals and other scoring chances.

In general, we believe that our approach represents a major step in team multiagent systems in terms of incorporating *anticipation* as a key aspect of teamwork.

4 Results

CMUnited-98 successfully defended our title of the Small Robot Champion at RoboCup-98 in Paris. The competition involved 11 teams from 7 different countries. It consisted of a preliminary round of two games, followed by the 8 advancing teams playing a 3-round playoff. CMUnited-98 won four of five games, sweeping the playoff competition, scoring a total of 25 goals scored and only 6 suffered. The individual results of these games are in Table 2.

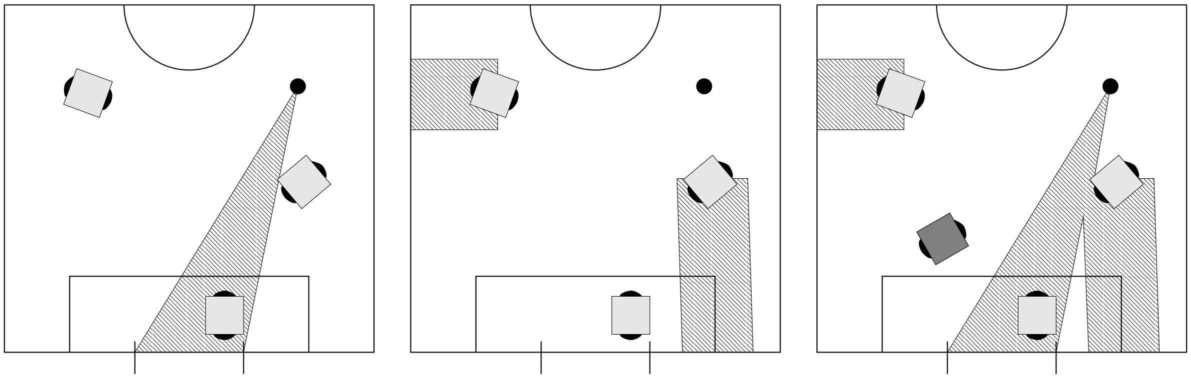

(a) Don't block goal shot.

(b) Avoid difficult collaboration.

(c) Anticipate optimal position for collaboration.

Figure 9: Constraints for the dynamic anticipation algorithm are represented as shaded regions; (a) and (b) show three opponents and the current position of the ball; (c) illustrates the position of the passive agent - dark square - as returned by SPAR.

Phase	Opponent	Affiliation	Score (CMU - Opp.)
round-robin	iXS	iXs Inc., Japan	16 – 2
round-robin	5DPO*	University of Porto, Portugal	0 – 3
quarter-final	Paris-8*	University of Paris-8, France	3 – 0
semi-final	Cambridge	University of Cambridge, UK	3 – 0
final	Roboroos	University of Queensland, Australia	3 – 1

Table 2: The scores of CMUnited-98's games at RoboCup-98. The games marked with an * were forfeited at half time.

There were a number of technical problems during the preliminary rounds, including outside interference with our radio communication. This problem was the worst during our game against 5DPO, in which our robots were often responding to outside commands just spinning in circles. This led to our forfeit at half time and a clear loss against 5DPO, a very good team which ended in third place at RoboCup-98. Fortunately, the communication problems were isolated and dealt with prior to the playoff rounds.

The three playoff games were very competitive and showcased the strengths of our team. Paris-8 had a strong defense with a lot of traffic in front of the goal. Our team's obstacle avoidance still managed to find paths and to create scoring chances around their defenders. The final two games were very close against very good opponents. Our interception was tested against Cambridge, and included blocking a powerful shot by their goalie, which was deflected back into their goal. The final game against Roboroos demonstrated the dynamic positioning, especially during the final goal, which involved a pass to a strategically positioned teammate.

5 Conclusion

The success of CMUnited-98 at RoboCup-98 was due to several technical innovations, including robust hardware design, effective vision processing, reliable time-prediction based robot motion with obstacle avoidance, and a dynamic role-based team approach. The CMUnited-98 team demonstrated in many occasions its collaboration capabilities which resulted from the robots' behaviors. Most remarkably, CMUnited-98 introduces the concept of *anticipation*, in which passive robots (not going to the ball) strategically position themselves using attraction and repulsion (SPAR) to maximize the chances of a successful pass.

The CMUnited-98 team represents an integrated effort to combine solid research approaches to hardware design, vision processing, and individual and team robot behaviors.

Acknowledgements

We would like to thank the other two members of the CMUnited-98 team for contributing to its success. Sorin Achim and Kwun Han were responsible for designing and building the robots. In addition, Kwun Han was responsible for developing the vision system.

References

[1] Minoru Asada, Shoichi Noda, Sukoya Tawaratumida, and Koh Hosoda. Purposive behavior acquisition for a real robot by vision-based reinforcement learning. *Machine Learning*, 23:279–303, 1996.

[2] Jean-Claude Latombe. *Robot Motion Planning*. Kluwer Academic Publishers, Boston, 1991.

[3] Michael K. Sahota, Alan K. Mackworth, Rod A. Barman, and Stewart J. Kingdon. Real-time control of soccer-playing robots using off-board vision: the Dynamite testbed. In *IEEE International Conference on Systems, Man, and Cybernetics*, pages 3690–3663, 1995.

[4] Randy Sargent, Bill Bailey, Carl Witty, and Anne Wright. Dynamic object capture using fast vision tracking. *AI Magazine*, 18(1):65–72, Spring 1997.

[5] Peter Stone and Manuela Veloso. The CMUnited-97 simulator team. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*. Springer Verlag, Berlin, 1998.

[6] Peter Stone, Manuela Veloso, and Patrick Riley. The CMUnited-98 champion simulator team. In Minoru Asada and Hiroaki Kitano, editors, *RoboCup-98: Robot Soccer World Cup II*. Springer Verlag, Berlin, 1999.

[7] Manuela Veloso, Peter Stone, and Kwun Han. CMUnited-97: RoboCup-97 small-robot world champion team. *AI Magazine*, 19(3):61–69, 1998.

[8] Manuela Veloso, Peter Stone, Kwun Han, and Sorin Achim. The CMUnited-97 small-robot team. In Hiroaki Kitano, editor, *RoboCup-97: Robot Soccer World Cup I*, pages 242–256. Springer Verlag, Berlin, 1998.