

# CMUNITED-98

## RoboCup-98 Small-Robot World Champion Team

*Manuela Veloso, Michael Bowling, Sorin Achim, Kwun Han, and Peter Stone*

- The CMUNITED small-robot team became the 1998 RoboCup small-robot league champion, repeating its 1997 victory. CMUNITED-98 built on the success of cmunited-97 and involved a number of improvements. This article gives an overview of the CMUNITED-98 team, focusing on this year's improvements. It concludes with the results of the RoboCup-98 competition.

The CMUNITED-98 small-size robot team is a complete autonomous architecture composed of the physical robotic agents, a global vision-processing camera overlooking the playing field, and several clients as the minds of the small-size robot players. The global vision algorithm perceives the 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 communication. Motion is not perfectly executed because of inherent mechanical inaccuracies and unforeseen interventions from other agents. This team competed in and won the 1998 RoboCup competition in Paris (Stone, Veloso, and Riley 1999; Kitano et al. 1997).

This article gives an overview of the CMUNITED-98 robot team, focusing on improvements from the CMUNITED-97 team. These improvements include a robust low-level control algorithm, which handles a moving target with integrated obstacle avoidance, and active team collaboration at the strategic level. This article concludes with results from the RoboCup-98 competition.

### Hardware

The CMUNITED-98 robots (figure 1) are entirely our new constructions built on our experience in 1997 (Veloso, Stone, and Han 1998). The new robots represent an upgrade of our previously built CMUNITED-97 robots. Improvements were made in two major areas: (1) motors and control and (2) the mechanical chassis (including a kicking device).

CMUNITED-98 uses two high-torque, 6V DC, geared motors, which are overpowered and use a simple pulse-width-modulation control. This design is simpler than the design for our CMUNITED-97 robots, which used motor encoders for hardware feedback. Although our previous team had accurate navigation, it was not easily interruptible, which is necessary for operating in a highly dynamic environment. In the CMUNITED-98 robots, the closed-loop motion control is achieved through software using only visual feedback.

In designing the mechanical structure of the CMUNITED-98 robots, we focused on modularity and robustness. The final design includes a battery module supplying three independent power paths (for the main-board, motors, and radio modules). It also includes a single board containing all the required electronic circuitry with multiple add-on capabilities. The mobile base module includes a kicking device driven by a DC motor. This motor is hardware activated by an array of four infrared sensors, which is enabled or disabled by the software control. All this was combined in a layered design within an aluminum and plastic frame. In addition, each of the modules within this design is completely interchangeable.



Figure 1. The CMUNITED-98 Robots.

## Vision Processing

The CMUNITED-98 vision module remains largely the same as the one used in the CMUNITED-97 team (Han and Veloso 1998). The algorithm successfully detects and tracks 11 objects (5 teammates, 5 opponents, and 1 ball) at 30 frames a second. The algorithm determines the position and orientation for the robots. In addition, a Kalman-Bucy filter (Kalman and Bucy 1961) is used as a predictor of the ball's trajectory. This prediction is an integral factor in our robots' control and strategic decisions.

## Motion Control

Before developing strategic behaviors, the robots need a general control mechanism. This mechanism must reliably control the robot to a precise position on the field. The goal of our low-level motion-control mechanism is to be as fast as possible while still accurate and reliable, which is challenging because of 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 function using a reactive control mechanism that directs a differential drive robot to a target configuration. Although based on the CMUNITED-97's motion control (Veloso et al. 1998), CMUNITED-98 includes a number of major improvements. First, the target configuration for the motion planner has been extended to include both the Cartesian position and the direction that the robot is required to be facing when arriving at the target position. Second, the motion-controller algorithm drives the two-wheeled robot smoothly and includes the following three fea-

tures: (1) obstacle avoidance is integrated into the controller, (2) 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, and (3) the motion controller returns an estimate of the time that the robot will achieve the desired target configuration.

## Differential Drive Control

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

$$\Delta = \theta - \phi$$

$$(t, r) = \begin{pmatrix} \cos^2 \Delta \cdot \text{sgn}(\cos \Delta), \\ \text{sgn}^2 \Delta \cdot \text{sgn}(\sin \Delta) \end{pmatrix}$$

$$v_l = v(t - r)$$

$$v_r = v(t + r) \quad (1)$$

where  $\theta$  is the direction of the target point  $(x^*, y^*)$ ,  $\phi$  is the robot's orientation, and  $v$  is the desired speed (figure 2a).<sup>1</sup>

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

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

where  $\theta'$  is the new target direction,  $\alpha$  is the difference between  $\theta$  and  $\phi^*$ ,  $d$  is the distance to the target point, and  $c$  is a clearance parameter (figure 2a). This adjustment will keep the robot a distance  $c$  from the target point while it is circling to line up with the target direction,  $\phi^*$ . This new target direction,  $\theta'$ , 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 hand-tuned linear function of  $d$ ,  $\alpha$ , and  $\Delta$ .

## Obstacle Avoidance

Obstacle avoidance was also integrated into the motion control by adjusting the target direction of the robot based on any immediate obstacles in its path. This adjustment can be seen in figure 2b. If a target direction passes too

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

### 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. Figure 3a illustrates the control path to reach a stationary target with a specific direction, using the control mechanism described previously. Our extension allows for the target configuration to be given as a function of time,

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

where  $t = 0$  corresponds to the present. 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

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

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

Unfortunately,  $t^*$  cannot easily be 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 constraint. An illustration of this procedure is shown in figure 3b, where the goal is to hit the moving ball. 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.

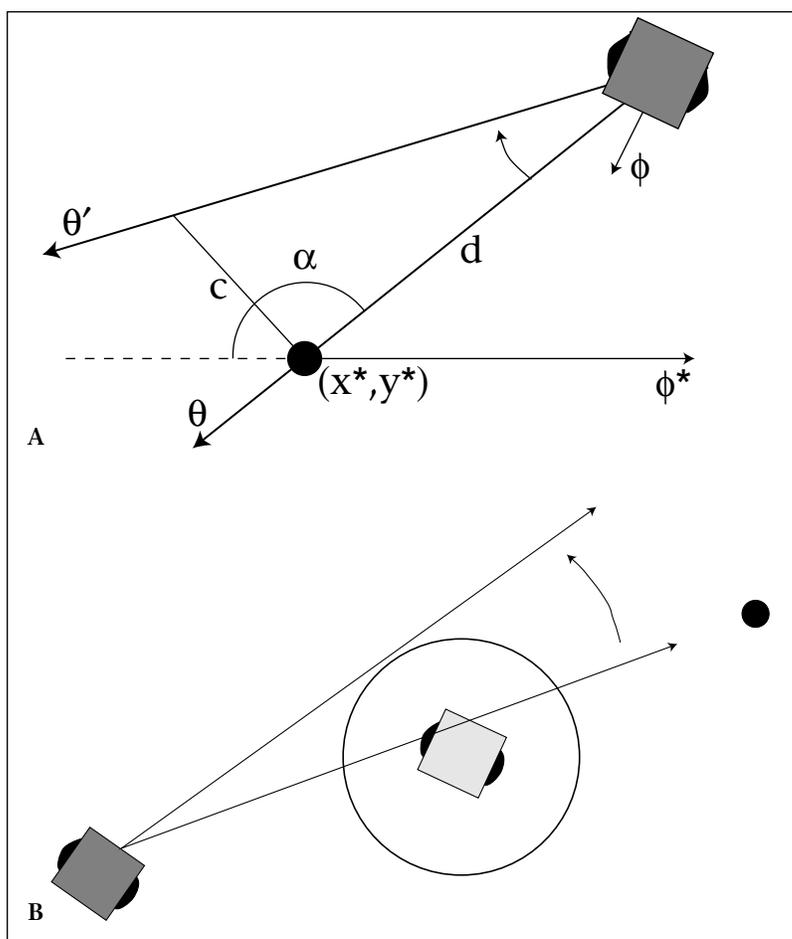


Figure 2. Motion to Position and Orientation.

A. The adjustment of  $\theta$  to  $\theta'$  to reach a target configuration of the form  $(x^*, y^*, \phi^*)$ . B. The adjustment to avoid immediate obstacles.

### 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, but at the same time, its decisions must contribute toward 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 a formation (Stone and Veloso 1998; Veloso, Stone, and Han 1998). CMUNITED-98 builds on 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: goal keeper, defender, and attacker. The formation used throughout RoboCup-98 involved a single goal keeper and defender and three attackers. The goal tender's behavior is similar to CMUNITED-97's and is described in Veloso et al. (1999a). This article describes the defender's behavior

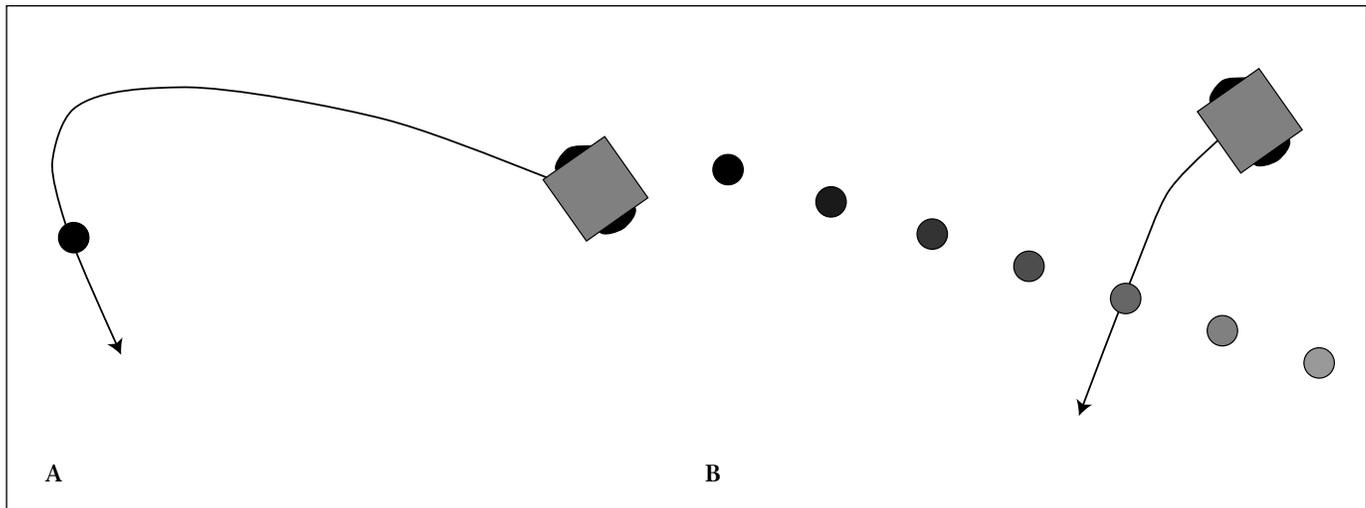


Figure 3. Motion for Intersection of Target.

A. Control for stationary target. B. Control for moving target.

and the collaborative behaviors developed for the attackers.

### 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 twofold: (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 a result of the uncertain ball handling by the robots. The robots can easily push the ball unexpectedly in the wrong direction when performing a difficult maneuver.

To achieve the two goals, we implemented three behaviors for the defender. First, *blocking*, illustrated in figure 4a, is similar to the goal keeper's behavior except that the defender positions itself further away from the goal line. Second, *clearing*, illustrated in figure 4b, pushes the ball out of the defending area by finding the largest angular direction free of obstacles (opponents and teammates) that the robot can push the ball toward. Third, *annoying*, illustrated in figure 4c, is somewhat similar to the goal-keeping behavior except that the robot tries to position itself between the ball and the opponent nearest to it in an effort to keep the opponent from reaching the ball.

Selecting when each of these behaviors is used is 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 back off another robot can lead to scoring in our own goal. We used the decision tree in figure 5 to select which action to perform

based on the current state. The two attributes in the tree, namely, (1) Ball Upfield and (2) Safe to Clear, are binary. Ball Upfield tests whether the ball is upfield (toward 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; otherwise, it either clears or blocks depending on the value of Safe to Clear.

### Active and Anticipating Attackers

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 multiagent 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, and the other teammates are passive. CMUNITED-98 significantly extends this simplistic view in two ways: (1) we use a decision-theoretic algorithm to select the active agent and (2) we use a technique for passive agents to anticipate future collaboration.

#### Individual Behaviors

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

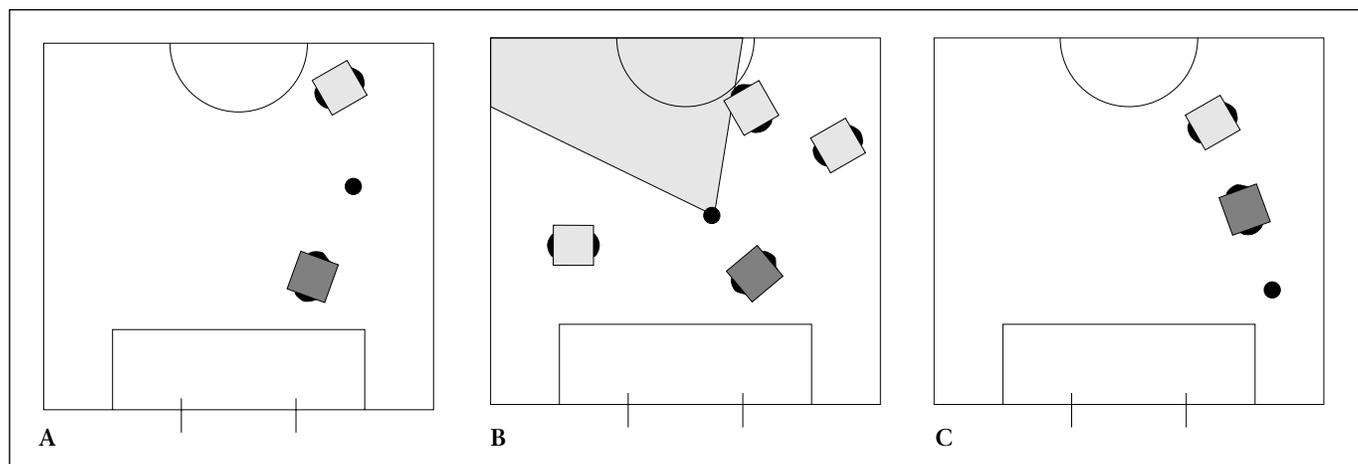


Figure 4. The Defender's Behaviors.

The dark and light robots represent the defender and the opponents, respectively. A. Blocking. B. Clearing. C. Annoying.

*motion controller*. The target configuration is specified as the ball (using its estimated trajectory), and the target direction is either toward the goal or toward another teammate, giving us robust and accurate individual behaviors that can handle obstacles as well as intercept a moving ball.

#### Decision-Theoretic Action Selection

Given the individual behaviors, we must select an active agent and appropriate behavior by using decision-theoretic analysis with 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 the 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, that is, approximately 30 times a second. The probabilities of success, estimates of time, and values of actions are continuously being recomputed, allowing for quick changes of actions if shooting opportunities become available, or collaboration with another agent appears more useful.

#### Dynamic Positioning

The selected action determines the behavior

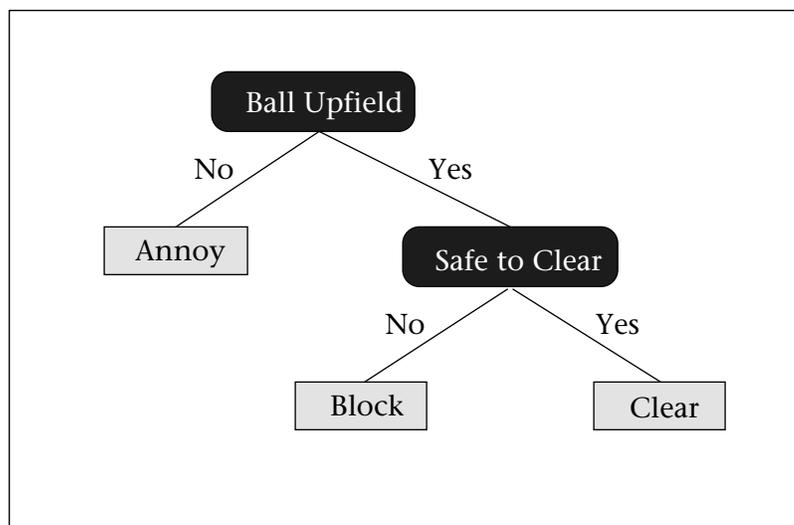


Figure 5. The Decision Tree Heuristic Used by the Defender to Select Its Behavior.

for the active agent, but it is unclear what the passive agents should be doing. CMUNITED-98 introduced a new technique for the passive agents to strategically position themselves to anticipate future opportunities for collaboration. The algorithm for this positioning is called *strategic positioning with attraction and repulsion* (SPAR). This algorithm was also used successfully in the CMUNITED-98 simulator team (Stone, Veloso, and Riley 1999).

This strategic position takes into account the position of the other robots (teammates and

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.

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 position of opponent  $i = 1, \dots, n$ ;  $T_i$ , the position of teammate  $i = 1, \dots, n$ ;  $B$ , the position of the active teammate and ball;  $G$ , the position of the opponent’s goal; and  $P$ , 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, SPAR, which extends similar approaches using potential fields (Latombe 1991), to our highly dynamic, multiagent 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 (1) repulsion from opponents, maximize the distance to each opponent:  $\forall_i \max dist(P, O_i)$ ; (2) repulsion from teammates, maximize the distance to other passive teammates:  $\forall_i \max dist(P, T_i)$ ; (3) attraction to the ball,  $\min dist(P, B)$ ; (4) 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 in a single-objective function.

Because each term in the multiple-objective function can have a different relevance (for example, staying close to the goal might be more important than staying away from opponents), we want to consider different functions for 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, giving us a weighted single-objective function:

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

This optimization problem is then solved under a set of constraints: First, do not block a possible direct shot from the active teammate. Second, do not stand behind other robots because these are difficult positions to receive a pass from the active teammate.

The solution to this optimization problem with constraints gives us a target location for the passive agent. Figures 6a and 6b illustrate these two sets of constraints, and figure 6c shows the combination of these constraints and the resulting position of the anticipating passive teammate.

## 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 2 games, followed by the 8 advancing teams playing a 3-round playoff. CMUNITED-98 won four of its five games, sweeping the playoff competition with 25 goals scored and only 6 suffered. The individual results of these games are shown in table 2.

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 SDPO, where our robots were often responding to outside commands and just spinning in circles; thus, we had to forfeit at half-time and suffered a clear loss against SDPO, a good team that 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 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 motion control with obstacle avoidance still managed to find paths and create scoring chances around their defenders. The final two games were close against good opponents. Our interception was tested against CAMBRIDGE and included blocking a powerful shot by its goal tender, which was deflected back into its goal. The final game against ROBOROOS demonstrated the dynamic positioning, especially during the final goal,

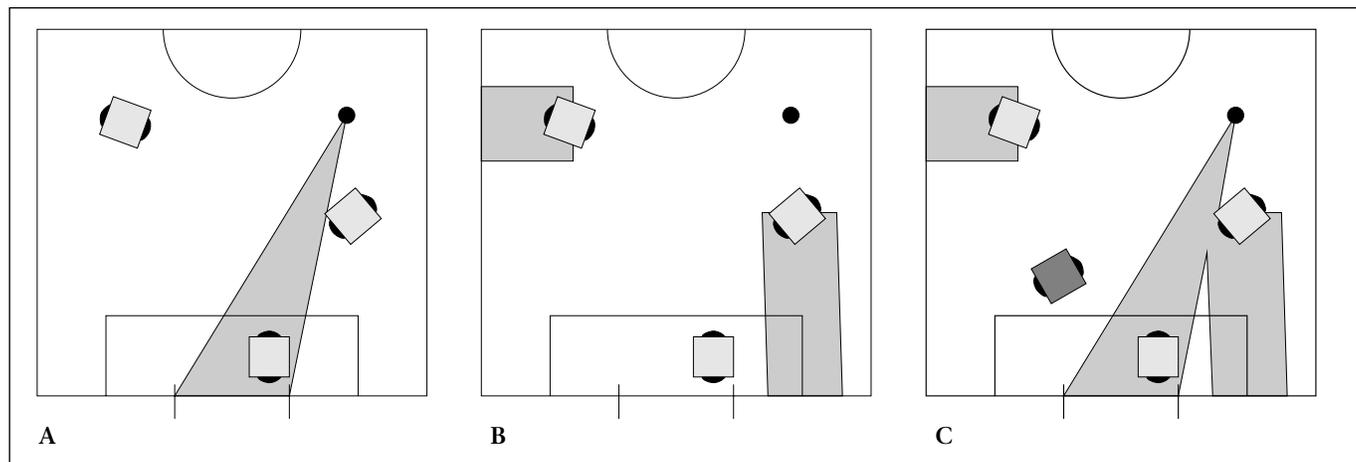


Figure 6. Constraints for the Dynamic Anticipation Algorithm Are Represented as Shaded Regions.

A. Don't block goal shot. B. Avoid difficult collaboration. C. Anticipate optimal position for collaboration. 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 (strategic positioning with attraction and repulsion).

Phase	Opponent	Affiliation	Score (CMU <sup>1</sup> -Opp.)
Round robin	IXS	iXs Inc., Japan	16 – 2
Round robin	SDPO	University of Porto, Portugal	0 – 3
Quarter final	PARIS-8	University of Paris-8, France	3 – 0
Semifinal	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.

<sup>1</sup>CMU = Carnegie Mellon University.

which involved a pass to a strategically positioned teammate.

## Conclusion

The success of CMUNITED-98 at RoboCup-98 was the result of several technical innovations, including robust hardware design; effective vision processing; reliable time-prediction-based robot motion with obstacle avoidance; a role-based team strategy; and in particular, an anticipation algorithm to effectively respond to the dynamic environment by increasing the opportunities for team collaboration. The CMUNITED-98 team demonstrated on many occasions its robust motion control and teamwork capabilities. The CMUNITED-98 team represents an integrated effort to combine solid research approaches to hardware design, vision process-

ing, and individual and team robot behaviors. Our ongoing research includes action policy learning from a crude robot simulator to the real robots, online robot recognition of the opponents' team strategy, and dynamic role and formation switching as a function of the opponent team.

## Acknowledgments

An extended version of this article appears in Veloso et al. (1999a and 1999b). This research is sponsored in part by the Defense Advanced Research Projects Agency and the Air Force Research Laboratory under agreements F30602-97-2-0250 and F30602-98-2-0135 and in part by the Department of the Navy, Office of Naval Research, under contract 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 United States Air Force, the Department of the Navy, Office of Naval Research, or the United States government.

### Note

<sup>1</sup> All angles are measured with respect to a fixed coordinate system.

### References

- Han, K., and Veloso, M. 1998. Reactive Visual Control of Multiple Nonholonomic Robotic Agents. Paper presented at the International Conference on Robotics and Automation, 16–21 May, Leuven, Belgium.
- Kalman, R. E., and Bucy, R. S. 1961. New Results in Linear Filter and Prediction Theory. *Journal of Basic Engineering* 83(4): 95–108.
- Kitano, H.; Asada, M.; Kuniyoshi, Y.; Noda, I.; and Osawa, E. 1997. RoboCup: The Robot World Cup Initiative. In Proceedings of the First International Conference on Autonomous Agents, 340–347. Menlo Park, Calif.: AAAI Press.
- Latombe, J.-C. 1991. *Robot Motion Planning*. Boston: Kluwer Academic.
- Stone, P., and Veloso, M. 1998. The CMUNITED-97 Simulator Team. In *RoboCup-97: Robot Soccer World Cup I*, ed. H. Kitano, 242–256. Berlin: Springer Verlag.
- Stone, P.; Veloso, M.; and Riley, P. 1999. The CMUNITED-98 Champion Simulator Team. In *RoboCup-98: Robot Soccer World Cup II*, eds. M. Asada and H. Kitano, 61–76. Berlin: Springer Verlag.
- Veloso, M.; Stone, P.; and Han, K. 1998. CMUNITED-97: RoboCup-97 Small-Robot World Champion Team. *AI Magazine* 19(3): 61–69.
- Veloso, M.; Bowling, M.; Achim, S.; Han, K.; and Stone, P. 1999a. The CMUNITED-98 Champion Small Robot Team. In *RoboCup-98: Robot Soccer World Cup II*, eds. M. Asada and H. Kitano. Berlin: Springer Verlag.
- Veloso, M.; Bowling, M.; Achim, S.; Han, K.; and Stone, P. 1999b. CMUNITED-98: A Team of Robotic Soccer Agents. In Proceedings of the 1999 Innovative Applications of AI. Menlo Park, Calif.: American Association for Artificial Intelligence.



**Manuela M. Veloso** is associate professor of computer science at Carnegie Mellon University (CMU). She received her Ph.D. in computer science from CMU in 1992. She received a B.S. in electrical engineering in 1980 and an M.Sc. in electrical and computer engineering in 1984 from the

Instituto Superior Tecnico in Lisbon. Veloso's long-term research goal is the effective construction of teams of intelligent agents where cognition, perception, and action are combined to address planning, execution, and learning tasks, in particular, in uncertain, dynamic, and adversarial environments. Veloso

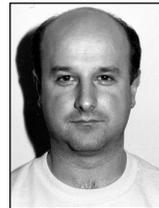
has developed teams of robotic soccer agents in three different leagues that are RoboCup world champions: simulation (1998), CMU-built small-wheeled robots (1997 and 1998), and Sony four-legged dog robots (1998). Veloso was awarded a National Science Foundation Career Award in 1995 and the Allen Newell Medal for Excellence in Research in 1997. Her e-mail address is veloso@cs.cmu.edu.



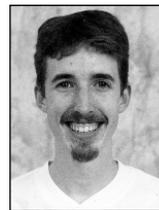
**Peter Stone** is a technical staff member in the Artificial Intelligence Department at AT&T Labs Research. He received his Ph.D. in 1998 and his M.S. in 1995 from Carnegie Mellon University, both in computer science. He received his B.S. in mathematics from the University of Chicago in 1993.

Stones research interests include planning and machine learning, particularly in multiagent systems. His doctoral thesis research contributed a flexible multiagent team structure and multiagent machine-learning techniques for teams operating in real-time noisy environments in the presence of both teammates and adversaries. His e-mail address is stone@cs.cmu.edu.

**Kwun Han** received a B.Sc. with honors in computer science from Brown University in 1996. He is currently pursuing a Ph.D. in computer science in the Computer Science Department at Carnegie Mellon University. His current research interests include robot behavior recognition, multiagent systems, machine learning, robotic soccer, and computer vision. Han was a member of the CMUNITED-97 and CMUNITED-98 small robot RoboCup champion teams. His e-mail address is kwunh@cs.cmu.edu.



**Sorim Achim** was a research engineer for the Romanian Department of Defense before joining Carnegie Mellon University in 1996. He holds an MSEE degree and is currently developing a line of educational robotics systems in Boston, Massachusetts. His e-mail address is sorin@cs.cmu.edu.



**Michael Bowling** is currently a Ph.D. candidate in computer science at Carnegie Mellon University (CMU). He received a B.Sc. from CMU in 1997. His research interests include mobile robots, machine learning, and multiagent systems. His e-mail address is mhb@cs.cmu.edu.