

# **AGILO RoboCuppers 2002: Applying Cooperative Game State Estimation, Experience-based Learning, and Plan-based Control to Autonomous Robot Soccer**

Michael Beetz, Sebastian Buck, Robert Hanek,  
Andreas Hofhauser, and Thorsten Schmitt

Munich University of Technology, Germany  
agilo-teamleader@mail9.in.tum.de  
<http://www9.in.tum.de/agilo/>

**Abstract.** This paper describes the computational model underlying the AGILO autonomous robot soccer team and its implementation. The most salient aspects of the AGILO control software are that it includes (1) a cooperative probabilistic game state estimator working with a simple off-the-shelf camera system; (2) a situated action selection module that makes ample use of experience-based learning and produces coherent team behavior even if inter-robot communication is perturbed; and (3) a playbook executor that can perform preprogrammed complex soccer plays in appropriate situations by employing plan-based control techniques. The use of such sophisticated state estimation and control techniques characterizes the AGILO software. The paper discusses the computational techniques and necessary extensions based on experimental data from the 2001 robot soccer world championship.

## **1 Overview**

The AGILO RoboCup team [1] is realized using inexpensive, off the shelf, easily extendible hardware components and a standard software environment. The team consists of four Pioneer I robots. Each robot is equipped with a single onboard linux computer, a wireless ethernet for communication, and several sonar sensors for collision avoidance. A color CCD camera with an opening angle of  $90^\circ$  is mounted fix on the robot. The robot also has a guide rail and a kicking device. That enables the robot to dribble and shoot the ball.

The incoming video streams from the CCD camera are processed by the *vision-based cooperative state estimation module*, which computes the *belief state* of the robot with respect to the game situation. The action selection module then computes an abstract feature description of the estimated game state that can be used to recognize relevant game situations. There are two basic components for action selection. First, the situated action selection module selects action based on a limited horizon utility assessment. Second, a plan-based controller that enables the team to execute more complex and learned plays. The following sections detail the software design and the operation of these software components.

## **2 Vision-based, Cooperative Game State Estimation**

The game state estimators of the AGILO robots maintain a belief state that contains the respective robot's belief about the current game situation [4]. The belief state includes

the estimated positions and orientations of the robot itself, its team mates, the ball, and the opponent robots and provides the information that is necessary for selecting the appropriate actions.

Estimating the game state both accurately and reliably is very difficult because the state is to be estimated by multiple mobile inaccurate sensors with uncertain positions, the soccer field is only partly accessible for each sensor due to occlusion caused by other robots and a limited camera view, the robots change their direction and speed very abruptly, the models of the dynamic states of the robots of the other team are very crude and uncertain, and it requires the processing of vast amounts of data very quickly.

State estimation is an iterative process where each iteration is triggered by the arrival of a new piece of evidence, a captured image or a state estimate broadcasted by another robot. The state estimation subsystem consists of three interacting estimators: the self localization system [3], the ball estimator, and the opponents estimator. This decomposition of game state estimation into specialized estimation problems reduces the overall complexity of the state estimation process and enables the robots to exploit the structures and assumptions underlying the different subtasks of the complete estimation task.

The state estimation modules of different robots cooperate to increase the coverage, accuracy, and reliability of the estimation process. This is achieved by the individual robots broadcasting their own observations and their covariances to their team mates. The other robots use these broadcasted estimates as additional evidence for their own estimation processes. Of course, how to weigh the broadcasted estimates against the own perceived features is a subtle and delicate issue. For example, we are weighing the own estimates of the ball much higher than the broadcasted estimates. This implies that in behaviors such as going to the ball the control signals are much tighter coupled to the robot's own perception. The other robots' estimates of the ball are only considered as stronger evidence if the ball is far away from the robot itself or not visible at all. We are currently analyzing the advantages and disadvantages of different parameterizations using our log files of the RoboCup competition last year.

Despite crude parameterizations, the cooperation between the robots enables them to track temporarily occluded objects and to faster recover their position after they have lost track of it.

### **3 Situated Action Selection and Execution**

Throughout the game the AGILO robots have a fixed set of tasks with different priorities. The tasks are *shoot the ball into the goal*, *dribble the ball towards the goal*, *look for the ball*, *block the way to the goal*, *get the ball*, ... The situated action selection module enables the robots to select a task and to carry out the task such that in conjunction with the actions of the team mates it will advance the team's objectives the most. We consider a task to be the intention of the AGILO robot team to perform certain actions. Action selection and execution is constrained by (1) tasks being achievable only if certain conditions hold (eg, the robot has the ball) and (2) a robot being able to only execute one action at a time.

We define that a task assignment  $a_1$  is better than  $a_2$  if there exists a task in  $a_2$  that has lower priority than all the ones in  $a_1$  or if they achieve the same tasks but there exists a task  $t$  in  $a_1$  such that all tasks with higher priority are performed at least as fast as in  $a_2$  and  $t$  is achieved faster by  $a_1$  than by  $a_2$ . This performance criterion implies that if an AGILO robot can shoot a goal it always will try because this is the task with the highest priority. Also, if the AGILO team can get to the ball it tries to get there with the robot that can reach the ball the fastest. This strategy might not yield optimal assignments but guarantees that the highest priority tasks are achieved as quickly as possible.

To achieve a high degree of autonomy the AGILO robots perform the task assignment and execution distributedly on the individual robots. This makes the task assignment more robust against problems in inter robot communication. These problems can be caused by robots being sent off the field, computers being crashed after heavy collisions, and communication being corrupted due to interferences with other communication channels.

The most salient features of the situated action selection are the following ones. First, to realize a competent and fast task assignment and execution mechanism the AGILO controllers make ample use of automatic learning mechanisms. Second, the task assignment mechanism works distributedly on the individual robots and are robust against communication corruptions. Finally, the task assignment and execution mechanism always produces purposeful behavior and always aims at the achievement of high priority tasks.

An important means for developing competent robot soccer skills is a robot simulator that allows for realistic, controllable, and repeatable experiments. For this reason we have developed a robot simulator [2] that accurately simulates how the dynamic state of the robot changes as the robot's control system issues new driving commands such as setting the target translational and rotational velocities.

## 4 Plan-based Control

While our situated action selection aims at choosing actions that have the highest expected utility in the respective situation it does not take into account a *strategic* assessment of the alternative actions and the respective *intentions* of the team mates. This is the task of the plan-based action control. While situated action selection achieves an impressive level of performance it is still hampered by the requirement for small action and state spaces, a limited temporal horizon, and without explicitly taking the intentions of the team mates into account.

The goal of plan-based control in robotic soccer is therefore to improve the performance of the robot soccer team by adding the capability of learning and execute soccer plays. Soccer plays are properly synchronized, cooperative macro actions that can be executed in certain game contexts and have, in these contexts, a high success rate. Plans for soccer plays specify how the individual players of a team should respond to changing game situations in order to perform the play successfully.

The integration of soccer plays into the game strategies enables robot teams to consider play specific state spaces for action selection, parameterization, and synchronization. In addition, the state space can reflect the intentions of the other robots. An action

that is typically bad might be very good if I know that my team mate intends to make a particular move. Further, action selection can consider a wider time horizon, and the robots can employ play specific routines for recognizing relevant game situations.

In order to realize an action assessment based on strategic consideration and on considerations of the intentions of the teammates, we develop a robot soccer playbook, a library of plan schemata that specify how to perform individual team plays. The plans, or better plays, are triggered by opportunities, for example, the opponent team leaving one side open. The plays themselves specify highly reactive, conditional, and properly synchronized behavior for the individual players of the team.

The high-level controller of each soccer robot is realized as a structured reactive controller (SRC) and implemented in an extended RPL plan language. The high-level controller works as follows. It executes as the default strategy the situated action selection. At the same time, the controller continually monitors the estimated game situation in order to detect opportunities for making plays. If an opportunity is detected, the controller decides based on circumstances including the score and the estimated success probability of the intended play whether or not to perform the play.

## 5 Conclusion

This paper has shortly described and discussed the control software of the AGILO autonomous robot soccer team. Similar to advanced autonomous robotic agents acting in human working environments the AGILO team employs sophisticated state estimation and control techniques, including experience-based learning and plan-based control mechanisms.

We have shown that the application of probabilistic state estimation techniques together with information exchange between the robots results in game state estimators that are capable of estimating complete states including robots with surprising accuracy and robustness even with restrictive and noisy camera systems. We have also seen that the ample use of experience-based learning has resulted in powerful control mechanisms, including competent coordination, with little runtime computational cost. Finally, we have explained how plan-based control mechanisms will enhance the robot's playing skills by enabling the robots to perform complex soccer plays. The results of the 2001 robot soccer world championship have shown that these techniques allow for competitive soccer play despite an inferior hardware equipment.

## References

1. M. Beetz, S. Buck, R. Hanek, T. Schmitt, and B. Radig: *The Agilo Autonomous Robot Soccer Team: Computational Principles, Experiences, and Perspectives*. International Joint Conference on Autonomous Agents and Multi Agent Systems (AAMAS), 2002.
2. S. Buck, U. Weber, M. Beetz, and T. Schmitt: *Multi Robot Path Planning for Dynamic Environments: A Case Study*. Proceedings of the IEEE International Conference on Intelligent Robots and Systems, 2001.
3. R. Hanek and T. Schmitt: *Vision-Based Localization and Data Fusion in a System of Cooperating Mobile Robots*. Proceedings of the IEEE International Conference on Intelligent Robots and Systems, 2000.
4. T. Schmitt, R. Hanek, S. Buck, and M. Beetz: *Cooperative Probabilistic State Estimation for Vision-based Autonomous Mobile Robots*. Proceedings of the IEEE International Conference on Intelligent Robots and Systems, 2001.