Philips CFT RoboCup Team description

Ton Peijnenburg[1], (ed)

[1] Philips CFT North America, 811 East Arques Avenue, Sunnyvale, CA 94086
a.t.a.peijnenburg@philips.com
https://www.quick.philips.com/robocup

Abstract. At the start of 2000, an enthusiastic team of colleagues from the Philips Centre for Industrial Technology (CFT) has taken up the challenge to participate in the RoboCup World Championships [1]. The first time appearance was at the 2002 German Open, where they made it 1st. In the 2002 RoboCup World Championships in Japan, the team got kicked out in the quarterfinals. The Philips project is set up as a hobby project, where a team of Philips employees designs and builds autonomous robots in their spare time. Facilities autonomous soccer playing robots is being designed and built. A brief description of the major design considerations for the CFT robots as well as some thoughts on intelligent system development are collected in this paper.

Introduction

A more elaborate description of the design of our robots can be found in our 2002 team description paper [2]. The main characteristics can be summarized as follows.

- Four wheel omni-directional drive unit capable of speeds up to 4 m/s and accelerations of 10 m/s^2.
- Ball kicking device with mechanical spring for energy storage, capable of kicking the ball with speeds up to 8 m/s.
- Ball handling device with omni-directional wheel that applies a unidirectional torque.
- Color CCD camera with Philips image processor built around a TriMedia processor.
- Basic coordination between robots: keeper and field player roles together with a ball claim mechanism for field players.

The development paradigm that the Philips CFT RoboCup team has used so far is best described as an “engineering” paradigm. It has developed dedicated hardware for the soccer robot, unlike other teams who have decided to use commercial platforms with slight modifications. Being an engineering center of excellence, an engineering approach provided the best fit to the capabilities at Philips CFT. Not withstanding this however, exploring intelligent control of cooperative autonomous robot units is definitely the ultimate goal. It will provide promising extensions to the existing technological base.

In our first year, we have found that it could be beneficial to use dedicated hardware and robot control with fairly limited intelligence. Our much-debated victory over CS Freiburg during the 2002 German Open supports such a statement. Simple control allowed for few failure mechanisms in the high-level control software and predictable behavior as a result of that. However, in Japan it soon became clear that our simplistic approach was no longer sufficient. Notwithstanding the fact that our early exit can be largely contributed to failure of the simple inter-robot coordination mechanism, our limited capabilities in both absolute positioning and recovery from less effective team settings also played an important role.

After the tournaments, we have identified a number of problem areas that needed immediate attention. A number of these areas will be described in more detail in this paper, together with the solutions we pursue for those areas.

Motion control

The 8 degree-of-freedom drive unit [3] is controlled in (x,y,phi) Cartesian coordinates. Orientations and speeds of the four independent wheels are calculated based on the desired motion of the robot's center. This desired motion was converted to a desired Center Of Rotation (COR). Based on this COR, the desired wheel orientation was determined and the desired wheel travel. Originally there was no odometry software that tracked the robot's position based on the encoder data. The encoder data was only used to close the low level servo loops. This worked sufficiently well because vision was used to guide the robot to the ball and towards the opponent's goal. However for self localization it is advantageous to have the odometry available at Cartesian level. Ideally we would like close the loop between the sensed odometry data and the motion command. This implied that we wanted to close the servo loop around the kinematics. This is non-trivial given that this introduced a number of non-linearities in the control loops.

This new approach has been studied and implemented. In addition, for optimal performance, the controller switches to non-linear control for large motions and for energy consumption reasons the controller switches to more classical linear control methods for small adjustments.

Vision system

Images from the 60 fps color camera are currently processed at a rate of 15 fps. We are clearly processing limited in this case. The speed of the camera may be explored further by introducing new processing hardware (parallel processing chips with up to 320 parallel processing units) and changes in algorithms. To use ball position data from the camera for closed loop control, both the update rate and the sensor delay are of importance. Both issues are given attention in the new designs.

Color segmentation is currently based on regions in the UV plane. Robustness for variations in lighting conditions is far from excellent. Therefore, the color selection regions will be extended from 2D to 3D regions that incorporate the luminance

A new 240 fps B/W camera will be introduced on the system. With additional parallel processing hardware, region of interest processing and new algorithms, we will try to exploit the camera speed. Currently, 25 fps processing is possible.

Ball handling and kicking

The Philips CFT kicker has been widely discussed ever since the first public shot was kicked in Germany. The mechanical spring is the most efficient way to store energy, and a substantial amount of energy is stored in our design: 19.3 Joule. The ball in our case will have a starting speed of 8 m/s. At this moment, we operate the kicking device at 85% of its maximum energy. The powerful kicking device in combination with the ball handler make it possible to accurately shoot the ball across the field in the opponent's goal.

The ball handling mechanism, although specially designed for this purpose, appeared not to let the ball roll freely during a dribble. The applied torque and available friction on the robot were too large. A redesign of the ball handling mechanism has been made and will be implemented on the robots.

Keeper sensing

Our keeper is not designed differently from the other players. We found that relying on odometer data only is not sufficient for the keeper to remain aligned well enough during a game. Ultrasound sensors that sense the back plane of the goal will be used to determine the robot's alignment and positioning. A motor current sensor detects collisions and triggers realignment. White line sensors under the robot will indicate crossing of the goal area lines.

At the German Open of 2002 we were initially hit with many red cards because our collision avoidance was lacking. During the match we added software to compute the difference between the estimated current and the actual current consumption of the robot. Discrepancies between these values were used as a sign of a collision. This significantly reduced our red-card rate. After the matches we optimized this algorithm further because it limited us in speed due to the increasing rate of false detections at higher velocities.

To see the ball to the sides of the robot is currently not possible with our one camera design. An omni-vision camera with double range mirror (both close and long range sections) has been developed and will be implemented and tested on the keeper first.

Other special hardware (i.e. a dedicated robot design) is currently not foreseen.

Robot coordination

With absolute field navigation abilities, it will be possible to add team play skills to our robot soccer team. While earlier inter-robot communication was limited to let just one robot claim the ball, now also absolute robot positions can be exchanged. The opens our game for adding team based strategy with different (dynamic) roles to our game, like have one robot to defend while others attack. For these tactical communication issues, a C++ communication library is being developed, allowing programmers to easily select a communication service on some preferred properties, i.e. the reliability of the channel may be traded against communication delay, peer to peer as well as broadcast modes can be selected and a priority mechanism for messages will be available.

Learning behavior selection

As mentioned previously, the robot behavior currently exists of a simple set of executable skills, residing in a finite state machine with fixed transitions between states. Alternative FSMs may be defined using the 'stactics' formal language and accompanying machine. It is quite similar to Maes' [4] action selection dynamics (Maes 1991, 'A bottom-up mechanism for behavior selection in a artificial creature'). The FSM decides when and which transition is made. The appropriate skill for a new state is selected and executed.

One of the problems with such a system is that there is a rather big chance that another implementation of that skill is better, but it is very hard to convince your co-developers that your implementation is better especially due to the very limited sample size (lack of real matches) to prove (or disprove) this. Another disadvantage is that the system will play the same against different opponents while a change in the selection of skills might give better results. Again this is hard to do by hand. Different people will judge the potential effectiveness of skills differently and changing skills during a game might prove to be disastrous, as some teams have found out the hard way in the past. Also the entire system evolves continuously and new skills, which use the new functionality, have to be compared with the old implementations to see whether their effectiveness has improved or not. Our learning system addresses these problems.

The approach uses the Interval Estimation Algorithm (IE) as described in (Kaelbling 1993, 'Learning in Embedded Systems'). For some states, instead of having the possibility to execute only one skill, more skills are implemented and one is chosen according to the IE algorithm. Some changes have been made to the original algorithm to make it work for non-stationary environments and especially RoboCup. The fitness function mostly addresses only time or amount of hits/(shots at the goal). Because the system is 'reset' at the beginning of every game with the average of all the previous experiences it can switch between skills (within a state) due to a low bias. This means that if the skill which proved to be the best during the last games, is not the best anymore the system will figure it out within a couple of minutes and will choose different skills. So effectively it changes tactics during a game to improve its possibilities against a particular opponent.

Speech

A speech recognition system is currently under development. Its goal is to make the robots more human like and more fun to watch and work with. Moreover, we will be able to assess the role of speech in the interaction with human coaches and (in the future) players.

A robot operator will use a headset wireless microphone when talking to the robot. The signal will be processed and recognized at the sideline computer, thereafter commands will be sent to the robots. The robot will generate a speech response (using Festival speech) to provide feedback to the operator.

Philips Research in Aachen, Germany has developed the recognition software. Dialogue, grammar and acoustic models have been designed for our robots and implemented. The system is largely speaker independent, so that most of the team members can use it. Vocabulary is limited and the dialogue is rather strict in the sense that one has to follow certain interactive steps to describe a command. The operator however will still be able to use natural sentences when talking to the robot.

We are curious to see how well the system will perform. Most likely we will have problems getting accurate recognition in a tournament environment, with many people talking, applauding etc.

Conclusion

Based on the experiences from its first year of competition, many adaptations will be implemented by the Philips team. Most of these are hardware modifications (i.e. sensors and actuators) together with the related processing software. One of the changes however will introduce the first intelligence with learning capabilities in the existing finite state machine controller. For the Philips CFT team, it will mark the introduction of AI in its current “engineering” approach of the robotic soccer problem.

References

1. https://www.quick.philips.com/robocup
2. A.T.A. Peijnenburg, T.P.H. Warmerdam et.al.: Philips CFT RoboCup Team Description. In: preliminary proceedings 2002 RoboCup conference, July 2002
3. Kein L. Moore and N.S. Flann: A Six-Wheeled Omnidirectional Autonomous Mobile Robot. In: IEEE Control Systems Magazine, December 2000, p. 53.

4. P. Maes: Artificial Life Meets Entertainment: Interacting with Lifelike Autonomous Agents. In: Comm. ACM, Vol. 38, No. 11, Nov. 1995, pp. 108-114.