RoboCup: A Challenge Problem for AI and Robotics

Hiroaki Kitano[1], Minoru Asada[2], Yasuo Kuniyoshi[3],
Itsuki Noda[3], Eiichi Osawa[1], Hitoshi Matsubara[3]

[1] Sony Computer Science Laboratory,
3-14-13 Higashi-Gotanda, Shinagawa, Tokyo 141 Japan
[2] Dept. of Mechanical Engineering, Osaka University,
Suita, Osaka 565 Japan
[3] Electrotechnical laboratory,
1-1-4 Umezono, Tsukuba, 305 Japan

Abstract. RoboCup is an attempt to foster AI and intelligent robotics research by providing a standard problem where wide range of technologies can be integrated and examined. The first **RoboCup** competition was held at IJCAI-97, Nagoya. In order for a robot team to actually perform a soccer game, various technologies must be incorporated including: design principles of autonomous agents, multi-agent collaboration, strategy acquisition, real-time reasoning, robotics, and sensor-fusion. RoboCup is a task for a team of multiple fast-moving robots under a dynamic environment. Although RoboCup's final target is a world cup with real robots, RoboCup offers a software platform for research on the software aspects of RoboCup. This paper describes technical challenges involved in RoboCup, rules, and simulation environment.

1 Introduction

RoboCup (The World Cup Robot Soccer) is an attempt to promote AI and robotics research by providing a common task for evaluation of various theories, algorithms, and agent architectures. In order for the robot (physical robot and software agent) to play a soccer game reasonably well, wide range of technologies need to be integrated and numbers of technical breakthrough must be accomplished. The range of technologies spans both AI and robotics research, such as design principles of autonomous agents, multi-agent collaboration, strategy acquisition, real-time reasoning and planning, intelligent robotics, sensor-fusion, and so forth. The first RoboCup, RoboCup-97, was held during IJCAI-97 at Nagoya, Japan, as a part of IJCAI-97's special program. Series of competitions are planned afterwards, just like the Formula One Championship. RoboCup consists of three competition tracks:

Real Robot League: Using physical robots to play soccer games. For RoboCup-97, there are two categories: small-size and middle size. A team for each category of league consists of up to five robots. The size of a robot for small size shall be within 180 cm^2 floor area and the maximum length of longer dimension shall be within 18cm. A middle size robot shall be witin 2000 cm^2. New categories will be created with technical needs and progress. Legged robot league and humanoid robot league, as well as wheel-based robots with 11 robots a team, are planned for future competitions.

Software Robot League: Using software agents to play soccer games on an official soccer server over the network.

Expert Robot Competition: Competition of robots which has special skills, but not able to play a game.

Although RoboCup's primary objective is a world cup with real robots, RoboCup offers a software platform for research on the software aspects of RoboCup. Software robot league, also called the simulator league, enables wider range of researchers to take part in this program. It also promotes research on network-based multi-agent interactions, computer graphics, and physically realistic animations — a set of technologies potentially promotes advanced use of internet. In addition, we intend to create an award for an expert robot, which demonstrates a high-level of competence for a specific task, such as shooting, intercepting, etc.

While it is so obvious that building robot to play soccer game is an immense challenge, readers might wonder why we propose RoboCup. It is our intentions to use RoboCup as a vehicle to revitalize AI research, by offering publicly appealing, but formidable challenge. One of the effective ways to promote engineering research, part from specific application developments, is to set a significant long term goal. When the accomplishment of such a goal has significant social impact, it is called the grand challenge project [Kitano et. al, 93]. Building a robot to play soccer game itself do not generate significant social and economic impact, but the accomplishment will certainly considered as a major achievement of the field. We call this kind of project as a landmark project. RoboCup is a landmark project as well as a standard problem.

The successful landmark project claims to accomplish a very attractive and broadly appealing goals. The most successful example is the Apollo space program. In case of the Apollo project, the U.S. committed the goal of "landing a man on the moon and returning him safely to earth." [Kennedy 61] The accomplishment of the goal itself marks the history of the mankind. Although the direct economic impact of having someone landed on the moon is slim[4], technologies developed to achieve this goal was so significant that it formed the powerful technological and human foundations to the American industries. The important issue for the landmark project is to set the goal high enough so that a series of technical breakthrough is necessary to accomplish the task, and the goal need to be widely appealing and exciting. In addition, a set of technologies necessary to accomplish the goal must be the technologies which can form the foundation of the next generation industries.

[4] To be fair, the Apollo mission was planned to gain the "National Prestige" and to demonstrate technical superiority over the former Soviet Union. Even in this, aspect, no direct military advantage was gained by having few astronauts on the moon.

In case of RoboCup, the ultimate goal is to "develop a robot soccer team which beats Brazil world cup team." (a more modest goal is "to develop a robot soccer team which play like a human players.") Needless to say, the accomplishment of the ultimate goal will take decades of efforts, if not centuries. It is not feasible, with the current technologies to accomplish this goal in any near term. However, this goal can easily create a series of well directed subgoals. Such an approach is common is any ambitious, or overly ambitious, project. In case of the american space program, the Mercury project and the Gemini project, which manned orbital mission, were two precursors to the Apollo mission. The first subgoal to be accomplished in RoboCup is "to build a real and software robot soccer teams which plays reasonably well with modified rules." Even to accomplish this goal will undoubtfully generates technologies which impacts broad range of industries.

One other aspect of RoboCup is a view that RoboCup is a standard problem so that various theories, algorithms, and architectures can be evaluated. Computer chess is a typical example of the standard problem. Various search algorithms were evaluated and developed using this domain. With the recent accomplishment by the Deep Blue team, which beated Kasparov, a human grand master, using the official rule, computer chess challenge is close to the finale. One of the major reasons for the success of computer chess as a standard problem is that the evaluation of the progress was clearly defined. The progress of the research can be evaluated as a strength of the system, which was indicated as the rating. However, as computer chess is about to complete its original goal, we need a new challenge. The challenge need to foster a set of technologies for the next generation industries. We consider that RoboCup fulfill such a demand. Table 1 illustrate difference of domain characteristics between computer chess and RoboCup.

	Chess	RoboCup
Environment	Static	Dynamic
State Change	Turn taking	Real time
Info. accessibility	Complete	Incomplete
Sensor Readings	Symbolic	Non-symbolic
Control	Central	Distributed

context recognition, vision, strategic decision-making, motor control, intelligent robot control, and many more.

2 Research Issues of RoboCup

In this section, we discuss several research issues involved in the development of real robots and software agents for RoboCup. One of the major reasons, why RoboCup attract so many researchers is that it requires integration of broad range of technologies into a team of complete agents, as opposed to a task-specific functional module. Following is a partial list of research areas involved in RoboCup:

- Agent Architecture in general
- Combining reactive approach and modeling/planning approach
- Real-time recognition, planning, and reasoning
- Reasoning and action in dynamics environment
- Sensor fusion
- Multi-agent systems in general
- Behavior learning for complex tasks
- Strategy acquisition
- Cognitive modeling in general

In addition to these technologies, providing network-based soccer server with high quality 3D graphics capability requires advancement of technologies for the real time animation of simulated soccer players and network-based interactive multi-user server system. In addition, numbers of natural language researchers are using RoboCup for the target domain of their automatic commentary generation systems, as seen in DFKI's Roccoco system. These are key technologies for network-based services in coming years. In this paper, we will analyse briefly on some of these issues.

2.1 Agent Architecture

The existing robot players have been designed to perform almost single behavior such as pushing/dribliling/rolling [Connel and Mahadevan 93a; Asada et. al, 95; Sahota 94], juggling [Rizzi and Koditschek 93; Schaal and Atkeson 94], or hitting [Watanabe et al. 94]. A RoboCup player should be designed so that it can perform multiple subtasks such as shooting (including kicking), dribbling (pushing), passing, heading, and throwing a ball which often involve a common behavior, avoiding the opponents. Roughly speaking, there are two ways to build up a RoboCup player. Design each component which is specialized for a single behavior and assemble them into one. The other approach is to design one or two components that can perform multiple subtasks. The former seems easier to design but difficult to assemble and *vice versa*. In addition, the problem of how to combine reactive approach and deliberative approach will be a major research issue. To quickly react against the ball and move around the field, use of subsumption architecture [Brooks 86], or other reactive approach may be effective. However, soccer players need to have global strategy as well as local tactics. These cannot be accomplished by mere reactive systems. One the other hand, deliberation-based approach, which involves planning and reasoning, may be too slow to react a quickly moving ball and to cope with a dynamically changing environment. The agent architecture for RoboCup players need to address the issue of how to combine these approaches.

2.2 Physical Components

Since the RoboCup player should move around quickly it should be compact, therefore the development of the integrated multi-functional module should be a new target of mechanical design for the RoboCup player. We need compact and powerful actuators with wide dynamic ranges. Also, we have to develop sophisticated control techniques to realize multiple behaviors by components as few as possible with low energy consumption.

The ultimate goal of a RoboCup player is like a humanoid type that can run and kick or pass a ball by its legs and feet, can throw a ball by its arms and hands, and can do a heading by its head. Since to build up a team of the humanoid types seems impossible within the current technology, this is just for a demonstration track for now. However, we expect that sometime in future participants of RoboCup overcome technical difficulties and participate with humanoid robots.

In addition, an attempt is being made to provide a standard physical components for robots. We are currently discussing about a possibility of making a standard for autonomous robot. The standard OpenR is not necessary designed for RoboCup, but RoboCup is one of its significant application areas [Fujita and Kageyama 97].

2.3 Vision and sensor fusion

The visual information is the richest source of information to perceive not only the external world but the effects of the robot actions as well. The Computer Vision researchers have been seeking for the accurate 3-D geometry reconstructed from 2-D visual information believing in that the 3-D geometry is the most powerful and general representation to be used in many applications such as view generation for video database and robot manipulation and navigation. However, the time-consuming 3-D reconstruction might not be necessary nor optimally encoded for the task given to the RoboCup player. In order to react to the situation in real time, the RoboCup player needs the information which behavior to select against which situation. This does not mean to build up a special-purpose vision system but to claim that vision is a part of complex system that interacts in specific ways with world [Aloimonos 94]. The RoboCup is one of such worlds which make clear the role of vision and evaluate the performance of the image processing that have been left ambiguous in the computer vision field.

In addition to vision, the RoboCup player might need other sensing such as sonar, touch, and force/torque to discriminate the situations that cannot be discriminated from only the visual information nor covered by the visual information. Again, the RoboCup player needs the real time processing for multi-sensor fusion and integration. Therefore, the deliberative approaches to obtain the robust estimation by multi-sensor system does not seem suitable. We should develop a method of sensor fusion/integration for the RoboCup.

2.4 Learning behaviors

The individual players has to perform several behaviors one of which is selected depending on the current situation. Since programming the robot behaviors against the all situations considering the uncertainties in sensory data processing and action execution is infeasible, robot learning methods seem promising. As a method for robot learning, reinforcement learning has recently been receiving increased attention with little or no *a priori* knowledge and higher capability of reactive and adaptive behaviors [Connel and Mahadevan 93b]. However, almost of the all existing applications have been done only with computer simulations in toy world, and real robot applications are very few [Asada et. al, 95; Connel and Mahadevan 93a]. Since the prominence of the role of the reinforcement learning is largely determined by the extent to which it can be scaled to larger and complex robot learning tasks, the RoboCup seems a very good platform.

One example of research on this issue, among other research such as [Stone and Veloso 96], is a project at Asada Lab at Osaka University. Here, we will only show some photo's on his robots in action, so that interested readers can access his papers for detail.

Figure 1 shows how a real robot shoots a ball into a goal by using the state and action map obtained by the method [Asada et. al, 96]. 16 images are shown in raster order from the top left to the bottom right in every 1.5 seconds, in which the robot tried to shoot a ball, but failed, then moved backward so as to find a position to shoot a ball, finally succeeded in shooting. Figure 2 shows a sequence of images taken by the robot during the task execution shown in Figure 1. Note that the backward motion for retry is just the result of learning and not hand-coded. The method used here is an off-line learning one. Currently, they used an on-line learning method [Uchibe et. al, 96a].

At the primary stage of the RoboCup tournament, one to one competition seems feasible. Since the player has to take the opponent motions into consideration, the complexity of the problem is much higher than that of simple shooting without an opponent. To reduce the complexity, the task decomposition is often used. Asada et al. [Asada et al 94b] proposed a method of learning a shooting behavior avoiding a goal keeper. The shooting and avoiding behaviors are independently acquired and they are coordinated through the learning. Their method still suffers from the huge state space and the perceptual aliasing problem [Whitehead and Ballard 90] due to the limited visual field.

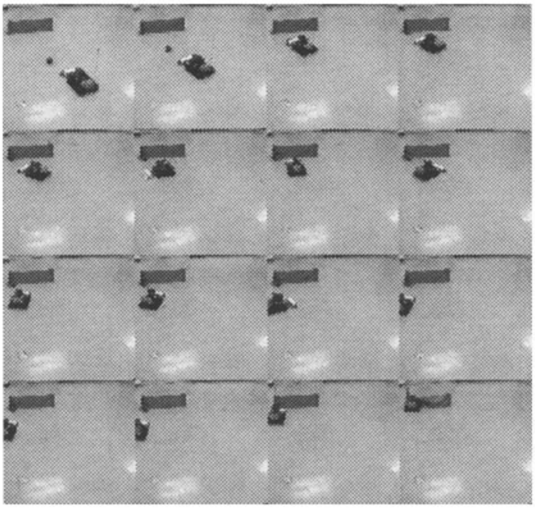

Fig. 1. A sequence of robot motion shooting a ball

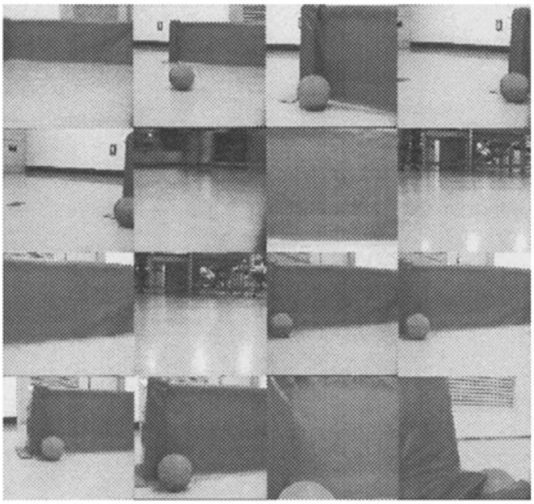

Fig. 2. A sequene of on-board camera images

Figure 3 shows a sequence of images in which the robot shoots a ball into a goal avoiding the opponent (a goal keeper) [Asada et al 94b].

2.5 Multi-Agent Collaboration

A soccer game is a specific but very attractive real-time multi-agent environment from the viewpoint of distributed artificial intelligence and multi-agent research. In a game, we have two competing teams. Each team has a team-wide common

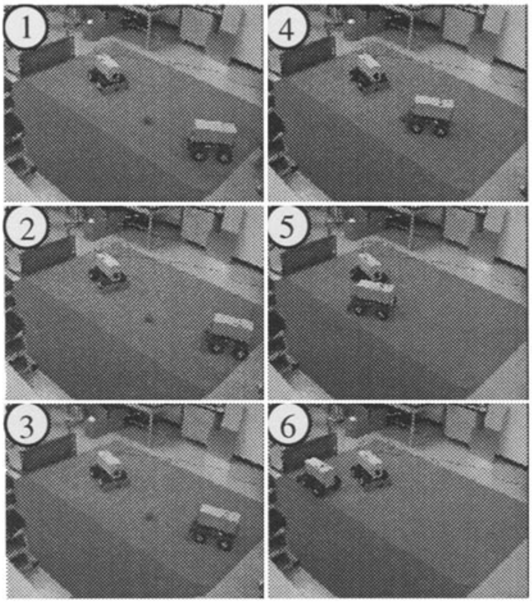

Fig. 3. A robot shooting a ball while avoiding an opponent robot

goal, namely to win the game. The goals of the two teams are incompatible. The opponent team can be seen as a dynamic and obstructive environment, which might disturb the achievement of the common team goal. To fulfill the common goal, each team needs to score, which can be seen as a subgoal. To achieve this subgoal, each team member is required to behave quickly, flexibly, and cooperatively; by taking local and global situations into account.

The team might have some sorts of global (team-wide) strategies to fulfill the common goal, and both local and global tactics to achieve subgoals. However, consider the following challenges:

1. the game environment, i.e. the movement of the team members and the opponent team, is highly dynamic.
2. the perception of each player could be locally limited.
3. the role of each player can be different.
4. communication among players is limited, therefore, each agent is required to behave very flexibly and autonomously in real-time under the resource bounded situation.

These restrictions are realistic and provides an interesting avenue of research for multi-agent systems. Let's us briefly look at multi-agent research which address cooperative planning under dynamics environment, where various resource and communication restrictions exists.

In cooperative distributed planning for common global goals, important tasks include the generation of promising local plans at each agent and coordination of these local plans. When the dynamics of the problem space, e.g. the changing rate of goals compared with the performance of each planner, is relatively large, reactive planning that interleaves the plan generation and execution phases is known to be an effective methodology at least for a single agent [McDermott 78; Agre and Chapman 87; Maes 91; Ishida and Korf 91]. Whether this scheme extends naturally to the multi-agent environment is an interesting issue.

For cooperative plan schemes, there are frequent changes in the problem space or the observation of each agent is restricted locally. There is a trade-off between communication cost, which is necessary to coordinate the local plans of agents with a global plan, and the accuracy of the global plan (this is known as the predictability/responsiveness tradeoff). The study of the relationship between the communication cost and processing cost concerning the reliability of the hypotheses in FA/C [Lesser and Erman 80], and the relationship between the modification cost of local plans and the accuracy of a global plan in PGP [Durfee and Lesser 87] illustrate this fact. Also, Korf addressed it theoretically in [Korf 87]. Tambe specifically use RoboCup domain to test a scheme for joint intention generation[Tambe 96].

Schemes for reactive cooperative planning in dynamic problem spaces have been proposed and evaluated sometimes based on the pursuit game (predator-prey) [Benda *et al.* 85; Stephens and Merx 89; Gasser *et al.* 89; Levy and Rosenschein 92; Korf 92; Osawa 95]. However, the pursuit game is a relatively simple game. Tileworld[Pollack and Ringuette 90] was also proposed and studied[Kinny and Georgeff 91; Ishida and Korf 91]. However, the environment is basically for the study of a single agent arichtecture.

As it is clear from these research, RoboCup directly address a critical issue in multi-agent systems research — generation and execution of cooperative plan under the dynamic environment. RoboCup provides an interesting and critically important task for multi-agent cooperative planning.

3 RoboCup Simulator

3.1 Soccer Server

In the simulation section, we will use Soccer Server, a simulator of **RoboCup** developed by Dr. Itsuki Noda, ETL, Japan, which is a network-based graphical simulation environment for multiple autonomous mobile robots in a 2D space. Using the soccer server, each client program can control each player on a soccer field via UDP/IP. This allows us to compare different types of multi-agent systems through the server, and test how well techniques of cooperation of agents work in dynamical varied situations.

The soccer server provides a virtual field where players of two teams play a soccer (association football) game. Each player is controlled by a client program via local area networks. Control protocols are simple in that it is easy to write client programs using any kind of programming system that supports UDP/IP sockets.

Control via Networks: A client can control a player via local area networks. The protocol of the communication between clients and the server is UDP/IP. When a client opens a UDP socket, the server assigns a player to a soccer field for the client. The client can control the player via the socket.

Physical Simulation: The soccer server has a physical simulator, which simulates movement of objects (ball and players) and collisions between them. The simulation is simplified so that it is easy to calculate the changes in real-time, but the essence of soccer is not lost.

The simulator works independently of communications with clients. Therefore, clients should assume that situations on the field change dynamically.

Referee: The server has a referee module, which controls each game according to a number of rules. In the current implementation, the rules are: (1) Check goals; (2) Check whether the ball is out of play; (3) Control positions of players for kick-offs, throw-ins and corner-kicks, so that players on the defending team keep a minimum distance from the ball.

Judgments by the referee are announced to all clients as an auditory message.

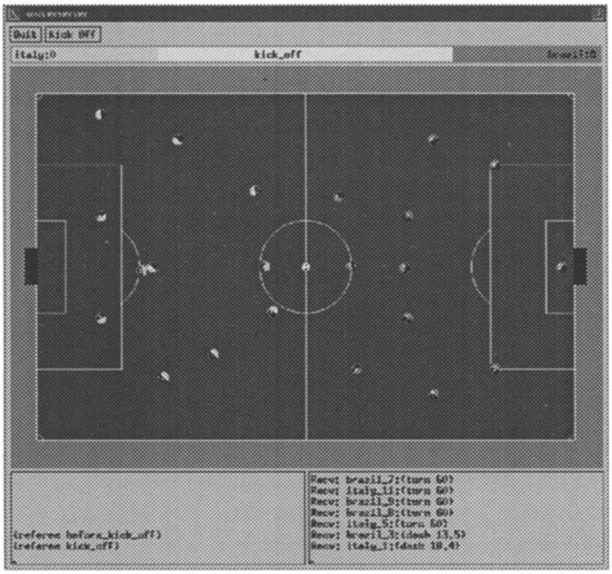

Fig. 4. Screen of Soccer Server

Although current version of the soccer server do not implements detailed physical and visual simulations, as seem in [Tu and Terzopoulos 94], we are planning to incorporate more realistic simulation in the future version.

3.2 Visualization

While the current version of the soccer server provides a top-down two-dimensional visualization, an independent server and browser systems provide a three dimensional visualization. A group of researchers at International Academy of Media Arts and Science (IAMAS) and Softopia, located in Gifu, Japan, is developing a three-dimensional visualization system. The first version of the system was used at RoboCup-97. A series of development project are planned for high-quality visualization, associated with automated camera-switching and commentary generation systems.

4 Conclusion

As it is clear by now, RoboCup provides fertile ground for AI and robotics research. The ultimate goal of RoboCup is so difficult that any near term accomplishment is not feasible. There is a clear paths to the stated goal, and each step toward the goal will generate a set of technologies which impacts industries for the next generation. Apart from these impact accessment, we believe that RoboCup contribute to AI and robotics community by providing exciting and publicly appealing goals, so that researchers in various field can collaborate for the common goal. We wish RoboCup offers an opportunities for AI and robotics community to revitalize their activities. "Let's get AI moving again!"

APPENDIX: Regulations for RoboCup

Regulations for RoboCup Real Robot Session (Summary)

General Policy

'Real worldness' in RoboCup mainly arises from the vast complexity of the overall situation due to interactions between behaviors and strategies of the ball and the players which cannot be fully predicted or controlled.

In the real robot session, we expect to have significantly greater complexity and hence much stronger reality than the simulation session. This is introduced by the uncertainty and uncontrollability in the structures and functions of the real robots along with real physical phenomena.

Therefore, we lean toward the least commitment policy in the game regulations, so that they do not obstruct surprises and creativity.

Due to the technical difficulty and unpredictability, the regulations can be adjusted to the overall situation of participating teams in each contest. However, the modifications must maintain the fairness to all the participants and must be announced in advance of the contest with an approval by the RoboCup technical committee.

The following sections summarize the regulations as of July 1997 very briefly. The most recent version can be obtained from the RoboCup web site. As rules being modified to cope with various technical and managment issues, please obtain the up-to-date rules from the web site. The rule have undergone two major changes since the first announcement in 1995. Also, prior to the announcement, several changes have been made since 1993, when we first drafted the RoboCup rule. The recent major changes include the size of the field, the size of robot, and the creation of defense zone. The field was defined based on a ping pong table so that most people can purchase it at low cost anywhere in the world. It is important to consider logistics of the material supply. The field, balls, and other materials should be so chosen that widest possible researchers can easily access and purchase with low cost. After a hearing period and discussions, the international committee for RoboCup have finalized the regulations for RoboCup-97. Further modifications will be made reflecting progress of research of participants. The RoboCup real robot league basically have two different classes based on the size of robots and the field — small size robot and midium size robot. Other classes, such as legged robots and humanoid robots may be created after the discussion by the committee.

The regulation for small robot league (excerpt)

Field Size: A ping pong table (a table tennis table) is used for the official match. The size and color of the table is officially determined as the international standard for ping pong. It is 152.5cm by 274cm, and color is green. Details shall be given in the figure 5.
Four small panels are attached to the corner to avoid ball to stuck. As shown in the figure below, it should be located 3 cm from the corner for each axis. Green strips of width 1cm shall be painted to identify the edge of the panel.

Robot: The maximum diameter of a circular robot shall be 15cm, while the maximum length of a rectangular robot shall be 18cm with a width of 10cm. These provide for the same size of robot in terms of surface area. This is approximately 1/10 of the length of the shorter end of the field.
(At the end of 1997, this rule was rewritten as:
For the small-size, each robot shall be within a floor area of 180 square centimeter, and the maximum length of the robot body shall be within 18cm. This is approximately equivalent to the previous definition that maximum diameter of a circular robot shall be 15cm. These provide for the same size of robot in terms of surface area. Height of the robot shall be within 15cm when the team uses the global vision system. The allowance was made for those who only uses the on-board vision systems (cameras are on-board, but not necessary all processors.) that up to 22.5cm height is permitted for 1998. This will be reviwed during RoboCup-98 Paris and the decision will be made whether this allowanced shall be abolished in future, and if so when it shall be abolished.)

Team: A team should consists of no more than 5 robots.

Goals: The width of the goal is 50 cm, which is approximately 1/3 of the length of the shorter end of the field.

Ball: Orange golf ball shall be used.

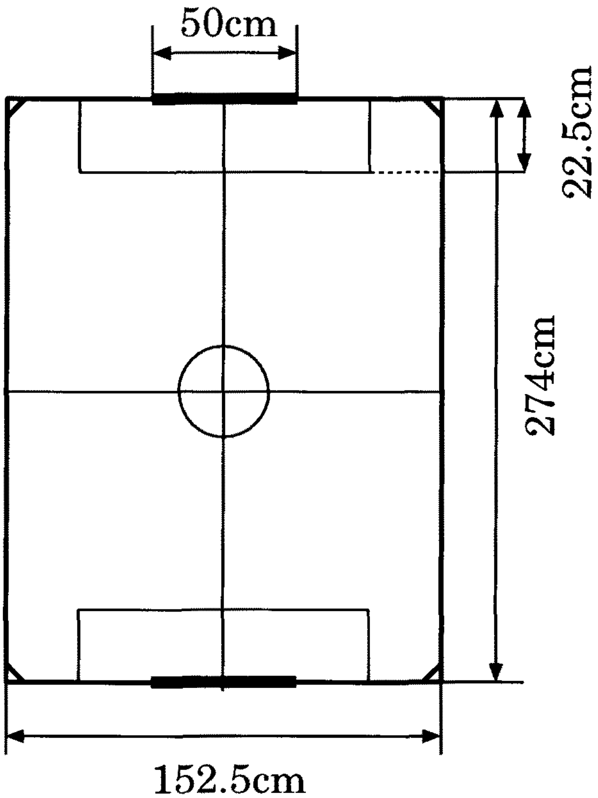

Fig. 5. Top view of the field for small robots

Colorings: Colors of each part of the field are as follows:

- Field shall be green.
- Wall shall be white.
- Ball shall be orange.
- Lines are drawn in white.
- Some markers on corners and goals are in green.

Length of the game: The games consits of the first half, break, and the second half. Each of them is 10 minutes.

Wall: A wall which is the same height as the golf ball shall be placed all around the field, except in goals. The wall shall be painted in white.

Defense Zone: Defense zone will be created in surrounding the goal of each side. It is 22.5 cm from the goal line, and width of 100 cm. The boarder of the defense zone will be painted in white, with the width of 1cm. Only one defense robot can enter this area. A brief passing and accidental entry of other robots are permitted, but intensional entry and stay is prohibited.

Global Vision System / External Distributed Vision System: The use of a global vision system and an external distributed vision system is permitted, not required, to identify and track the position of robots and balls. The rule explicitly allow the use of multiple cameras for tracking. The use of the global vision system shall be notified at the time of register, and detailed arrangements shall be discussed with the RoboCup organizing committee. (This rule was added at the end of December 1997. The rule make it explicit that for small-size, multiple camera and distributed vision is permitted to foster distributed vision research.)

Robot marking: Each robot should put at least one colored ping pong ball on top of their body, approximately between 15 cm to 20 cm in height. Teams may, at their discretion, use two pingpong balls (of differing colours) per robot to determine the orientation of the robot as well as its postion. The color(s) of the ping pong ball(s) will be used to identify friend and enemy, as well as positions using the global vision system.

Goal keepers: Goal keeper can hold and manipulate a ball for up to 10 seconds within its penalty area. After releasing the ball, the keeper must not hold the ball until it touches any opponent, or an alley outside the penalty area. If the ball released by the keeper reaches the other half end of the court without touching any other player, the opponent is given an indirect free kick positioned anywhere along the half way line (borrowed from Futsal rule).

Fouls: Following fouls are defined:

Multiple Defense: When more than one defense robots enter the defense zone to substantially affects the game. The foul will be called, and the penalty kick will be declared.

Ball Holding: A player cannot 'hold' a ball unless it is a goal keeper in its penalty area. Holding a ball means taking a full control of the ball by removing its entire degrees of freedom; typically, fixing a ball to the body or surrounding a ball using the body to prevent accesses by others. A free kick will be decleared. If this happens in the defense zone by the denfense team, a penalty kick will be declared.

Court Modification: Modification or damage to the court and the ball is forbidden. Should this occur, the game is suspended and the appropriate restoration is done immediately before the game resumes.

Robot Halting: All the players must be halted prior to kick-off or restarting of the game. The judges check or adjust the placements of the players and declares the completion of adjustment by 5 seconds before cueing a kick-off or a restart action. During this 5 seconds, the players can move.

Offside Offside rule is not adopted.

Charging: Unless during striving for a ball, a player must not attack another. In case the umpire clearly observes such an act, it is regarded as a violent action. Then the umpire presents a red card to the responsible player ordering it to leave the game. The judgment is done based on an external appearance.

Throughout the game, if a player utilizes a device or an action which continuously exerts, or whose primal purpose appears to be, serious damages to other robot's functions, the umpire can present a yellow card as a warning to the responsible player, and order it to go outside the court and correct the problem. Once the correction is made, the robot can resume to the game under an approval by the umpire. In case the problem is repeated, the umpire presents a red card to the responsible player telling it to leave the game.

Aside from the above items, no regulations are placed against possible body contacts, charging, dangerous plays, obstructions etc.

The Regulations for Medium Size Robots (excerpt)

The regulations for medium size robots basically applies the same rule as the rule for small robot. All sizes are multiled by 3. This means that:

Field: 457.5cm by 822cm

Goals: The width of the goal is 1500mm and the height is 500mm. Each goal is coloured in one colour (defined in Colouring) and has different colour from the other.

Corner: Four small panels are attached to the corner to avoid stucking the ball. As shown in the figure below, it should be located 200mm from the corner for each axis. Green strips of width 30mm shall be painted to identify the edge of the panel.

Penalty Area: Penalty area will be created in surrounding the goal of each side. It is 675mm from the goal line, and width of 3000mm. The area is shown by white lines with the width of 35mm.

Robot: The size of the robot should be within a circle of 500mm diameter or 450mm square.

Team: A team should consists of no more than 5 robots.

Ball: FIFA Size 4 Futsal ball painted red.

Colourings: Following parts are painted in different colour:

- A field is green
- Walls are white
- Lines drawn on the field is white
- Markers on corners is green (see the photo above)
- A ball is red
- Goals are blue and yellow

Regulations of Simulation Track (excerpt)

1. Field of Play

   The field of play is provided by Soccer Server, a simulator of a soccer field. A match is carried out in a server-client style: The server, Soccer Server, provides a virtual field and simulates all movements of a ball and players.

Clients become brains of players and control their movements. Communication between a server and each client is done using UDP/IP sockets via local area networks.

2. Players and Teams
   The simulation track of PreRoboCup consists of 'small track' and 'standard track'. In the small track, each team has 1 ~ 5 players. In the standard track, each team has 6 ~ 11 players. There is no goalkeeper because players have no hands. Even a team consists of fewer players than another team, a match is carried out without any penalties.
   Client programs can be written by any programming systems. with the following restrictions.
   (a) A client controls only a player. Or if a client controls multiple players, the different control modules of players are separated logically from each other.
   (b) Clients may not communicate directly with each other. Communication between clients must be done by facilities provided by Soccer Server.
3. Rules
   The referee module in Soccer Server controls a match according to 3 rules: goal, out-of-field, clearance. Moreover, a human referee also controls a match. When he/she judges player' action is too un-gentle, for example, surrounding the ball, he/she suspends the match and restarts by a free kick of the opposite team.
4. Format of the Competition
   The competition shall be played in two rounds. In the first round, teams shall be divided into several groups of 4 teams. The system of play shall be the league system, each team playing one match against each of the other teams in the same group. The two teams coming first and second in each group shall qualify for the second round.
   The second round shall be played by a system of elimination (cup system).

For more detail, please refer the following WWW homepage and FTP cite.

http://ci.etl.go.jp/~noda/soccer/regulations/regulations.html
http://ci.etl.go.jp/~noda/soccer/manual.newest/main.html
ftp://ci.etl.go.jp/pub/soccer/server/

References

[Agre and Chapman 87] P. Agre and D. Chapman. Pengi: An implementation of a theory of activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pp. 268–272, 1987.

[Aloimonos 94] Y. Aloimonos. "Rply: What i have learned". *CVGIP: Image Understanding*, 60:1:74–85, 1994.

[Asada et. al, 96] M. Asada, S. Noda, and K. Hosoda. Action-based sensor space categorization for robot learning. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS96)*, pages ???–???, 1996.

[Asada et. al, 95] M. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda. Vision-based reinforcement learning for purposive behavior acquisition. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 146–153, 1995.

[Asada et al 94a] M. Asada, S. Noda, S. Tawaratsumida, and K. Hosoda. "Purposive behavior acquisition on a real robot by vision-based reinforcement learning". In *Proc. of MLC-COLT (Machine Learning Conference and Computer Learning Theory) Workshop on Robot Learning*, pages 1–9, 1994.

[Asada et al 94b] M. Asada, E. Uchibe, S. Noda, S. Tawaratsumida, and K. Hosoda. "coordination of multiple behaviors acquired by vision-based reinforcement learning ". In *Proc. of IEEE/RSJ/GI International Conference on Intelligent Robots and Systems 1994 (IROS '94)*, pages 917–924, 1994.

[Benda *et al.* 85] M. Benda, V. Jagannathan, and R. Dodhiawalla. On Optimal Cooperation of Knowledge Sources. Technical Report BCS-G2010-28, Boeing AI Center, 1985.

[Brooks 86] R. Brooks, "A robust layered control system for a mobile robot," *IEEE J. Robotics and Automation*, RA-2, 14-23, April, 1986.

[Connel and Mahadevan 93a] J. H. Connel and S. Mahadevan. "Rapid task learning for real robot". In J. H. Connel and S. Mahadevan, editors, *Robot Learning*, chapter 5. Kluwer Academic Publishers, 1993.

[Connel and Mahadevan 93b] J. H. Connel and S. Mahadevan, editors. *Robot Learning*. Kluwer Academic Publishers, 1993.

[Tu and Terzopoulos 94] X. Tu, and D. Terzopoulos, "Artificial Fishes: Physics, Locomotion, Perception, Behavior," *Proc. of SIG GRAPH-94*, Orland, Florida, July 24-29, pp43-49, 1994.

[Durfee and Lesser 87] E. Durfee and V. Lesser. Using Partial Global Plans to Coordinate Distributed Problem Solvers. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, 1987.

[Fujita and Kageyama 97] M. Fujita and K. Kageyama, An Open Architecture for Entertainment Robot. In *Proceedings of the First International Conference on Autonomous Agents*, 1997.

[Gasser *et al.* 89] L. Gasser, N. Rouquette, R. Hill, and J. Lieb. Representing and Using Organizational Knowledge in Distributed AI Systems. In Les Gasser and Michael N. Huhns, editors, *Distributed Artificial Intelligence, Volume II*, pp. 55–78. Morgan Kaufmann Publishers, Inc., 1989.

[Ishida and Korf 91] T. Ishida and R. Korf. Moving Target Search. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 204–210, 1991.

[Kennedy 61] Kennedy, J. F., "Urgent National Needs," *Speech to a Joint Session of Congress*, 25 May 1961, Congressional Record – House (25 May 1961) p.8276.

[Kinny and Georgeff 91] D. Kinny and M. Georgeff. Commitment and Effectiveness of Situated Agents. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 82–88, 1991.

[Kitano et. al, 93] H. Kitano, B. Wah, L. Hunter, R. Oka, T. Yokoi, and W. Hahn, "Grand Challenge AI Applications," In *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, 1993.

[Korf 87] R. Korf. Planning as Search: A Quantitative Approach. *Artificial Intelligence*, Vol. 33, No. 1, pp.65–88, 1987.

[Korf 92] R. Korf. A Simple Solution to Pursuit Games. In *Proceedings of the Eleventh International Workshop on Distributed Artificial Intelligence*, 1992.

[Lesser and Erman 80] V. Lesser and L. Erman. Distributed Interpretation: A Model and Experiment. *IEEE Transactions on Computers*, Vol. 29, No. 12, pp.1144–1163, 1980.

[Levy and Rosenschein 92] R. Levy and J. Rosenschein. A Game Theoretic Approach to Distributed Artificial Intelligence. In Eric Werner and Yves Demazeau, editors, *Decentralized A.I. 3*. Elsevier/North Holland, 1992.

[Maes 91] P. Maes. Situated agents can have goals. In Pattie Maes, editor, *Designing Autonomous Agents: Theory and Practice from Biology to Engineering and Back*, pp. 49–70. The MIT Press, 1991.

[Mataric 94] M. J. Mataric. Learning to behave socially. In *Proc. of the 3rd Int. Conf. on Simulation and Adaptive Behaviors – From animals to animats 3 –*, pages 453–462, 1994.

[McDermott 78] D. McDermott. Planning and Action. *Cognitive Science*, Vol. 2, pp.71–110, 1978.

[Nourbaksh et al 93] I. Nourbaaksh, et al., The Winning Robots from the 1993 Robot Competition. *The AI Magazine*, Vol. 14, No. 4, 51-62, The AAAI Press, 1993.

[Osawa 95] E. Osawa. A Metalevel Coordination Strategy for Reactive Cooperative Planning. In *Proceedings of the First International Conference on Multi-Agent Systems*, 1995.

[Pollack and Ringuette 90] M. Pollack and M. Ringuette. Introducing the Tileworld: Experimentally Evaluating Agent Architectures. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pp. 183–189, 1990.

[Rizzi and Koditschek 93] A. A. Rizzi and D. E. Koditschek. Further progress in robot juggling: The spatial two-juggle. In *Proc. of IEEE Int. Conf. on Robotics and Automation*, pages 919–924, 1993.

[Sahota 94] M. K. Sahota. Reactive deliberation: An architecture for real-time intelligent control in dynamic environments. In *Proc. of AAAI-94*, pages 1303–1308, 1994.

[Schaal and Atkeson 94] S. Schaal and C. G. Atkeson. Robot learning by nonparametric regression. In *Proc. of IEEE/RSJ/GI International Conference on Intelligent Robots and Systems 1994 (IROS '94)*, pages 478–485, 1994.

[Stephens and Merx 89] L. Stephens and M. Merx. Agent Organization as an Effector of DAI System Performance. In *Proceedings of the Ninth Workshop on Distributed Artificial Intelligence*, pp. 263–292, 1989.

[Stone and Veloso 96] P. Stone, and M. Veloso, "Beating a defender in robotic soccer: Memory-based learning of a continuous function," In David S. Tourtzky, Michael C. Mozer, and Michael E. Hasselmo, editors, *Advances in Neural Information Processing Systems 8*, Cambridge, MA, MIT Press, 1996.

[Tambe 96] M. Tambe, "Tracking Dynamic Team Activity " In *Proceedings of AAAI-96*, Portland, 1996.

[Uchibe et. al, 96a] E. Uchibe, M. Asada, and K. Hosoda. Reasonable performance in less learning time by real robot based on incremental state space segmentation. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS96)*, pages ???–???, 1996.

[Uchibe et. al, 96b] E. Uchibe, M. Asada, and K. Hosoda. Behavior coordination for a mobile robot using modular reinforcement learning. In *Proc. of IEEE/RSJ International Conference on Intelligent Robots and Systems 1996 (IROS96)*, pages ???–???, 1996.

[Watanabe et al. 94] H. Watanabe, Y. Nihna, y. Masutani, and F. Miyazaki. Vision-based motion control for a hitting task - hanetsuki -. In *Proc. of IEEE/RSJ/GI*

International Conference on Intelligent Robots and Systems 1994 (IROS '94), pages 910–916, 1994.

[Whitehead and Ballard 90] S. D. Whitehead and D. H. Ballard. "Active perception and reinforcement learning". In *Proc. of Workshop on Machine Learning-1990*, pages 179–188, 1990.