An Architecture for a National RoboCup Team

Thomas Röfer

Center for Computing Technology (TZI), FB3, Universität Bremen
Postfach 330440, 28334 Bremen, Germany
`roefer@tzi.de`

Abstract. This paper describes the architecture used by the GermanTeam 2002 in the Sony Legged Robot League. It focuses on the special needs of a national team, i.e. a "team of teams" from different universities in one country that compete against each other in national contests, but that will jointly line up at the international RoboCup championship. In addition, the tools developed by the GermanTeam will be presented, e.g. the first 3-D simulation used in the Sony Legged Robot League.

1   Introduction

All RoboCup leagues share the problem that there are more teams that want to participate in the world championship than a normal contest schedule can integrate. Therefore, each league has its own approach to limit the number of participants to a certain amount. For instance, teams in the simulation league have to qualify by submitting a team description and two log files, one of a game against a strong opponent, the other of a game against an average one. Based on this material, a committee selects the teams for the championship. In contrast, participants in the Sony Legged Robot League only qualify by submitting a description of their scientific goals—before they have even worked with the robots. Each year, the Sony Legged League grows a little bit, e.g. from 16 to 19 teams from 2001 to 2002. So far, teams that were once accepted in the league were allowed to stay during the following years without a further qualification, amongst other reasons because of their investment in the robots. Thus, only a few new teams have the chance to join the league.

Currently, it is discussed how to provide to chance to participate in the league to more research groups. One solution would be to set up *national teams*, as the GermanTeam that was founded in 2001 [1]. The GermanTeam currently consists of students and researchers at five universities: the Humboldt-Universität zu Berlin, the Universität Bremen, the Technische Universität Darmstadt, the Universität Dortmund, and the Freie Universität Berlin. The members of the GermanTeam are allowed to participate as separate teams in national contests, but will jointly line up at the international RoboCup championship as a single team.

The other solution would be to install national leagues, of which the winning team will get the ticket to participate in the international contest. On the one hand, this approach would enforce the competition, but on the other hand, the goal of the RoboCup initiative is to promote research, and competing teams do not work together very well.

Therefore, it may be a good compromise to support collaboration on the lower, national level, while stressing the element of competition on the international level.

The GermanTeam is an example of a national team. The members will participate as separate teams in the German Open 2002, but will form a single team at Fukuoka. Obviously, the results of the team would not be very good if the members will develop separately until the middle of April, and then try to integrate their code to a single team in only two months. Therefore, an architecture for robot control programs was developed that allows to implement different solutions for the tasks involved in playing robot soccer. The solutions are exchangeable, compatible to each other, and they can even be distributed over a varying number of concurrent processes. The approach will be described in section 2. Finally, in section 3, the tools that were implemented to support the development of the robot control programs are presented.

2   Multi-Team Support

The major goal of the architecture presented in this paper is ability is to support the collaboration between the university-teams in the German national team. Some tasks may be solved only once for the whole team, so any team can use them. Others will be implemented differently by each team, e.g. the behavior control. A specific solution for a certain task is called a *module*. To be able to share modules, interfaces were defined for all tasks that could be identified for playing robot soccer in the Sony Legged League. These tasks will be summarized in the next section. To be able to easily compare the performance of different solutions for same task, it is possible to switch between them at runtime. The mechanisms that support this kind of development are described in section 2.2. However, a common software interface cannot hide the fact that some implementations will need more processing time than others. To compensate for these differences, each team can use its own *process layout*, i.e. they can group together modules to processes that are running concurrently (cf. section 2.2).

2.1   Tasks

Figure 1 depicts the tasks that were identified by the GermanTeam for playing soccer in the Sony Legged Robot League. They can be structured into five levels:

Sensor Data Processing. On this level, the data received from the sensors is preprocessed. For instance, the image delivered by the camera is segmented, and then it is converted into a set of blobs, i.e. image regions of the same color class. The current states of the joints are analyzed to determine the direction the camera is looking at. In addition, further sensors can be employed to determine whether the robot has been picked up, or whether it fell down.

Object Recognition. On this level, the information provided by the previous level is searched to find objects that are known to exist on the field, i.e. landmarks (goals and
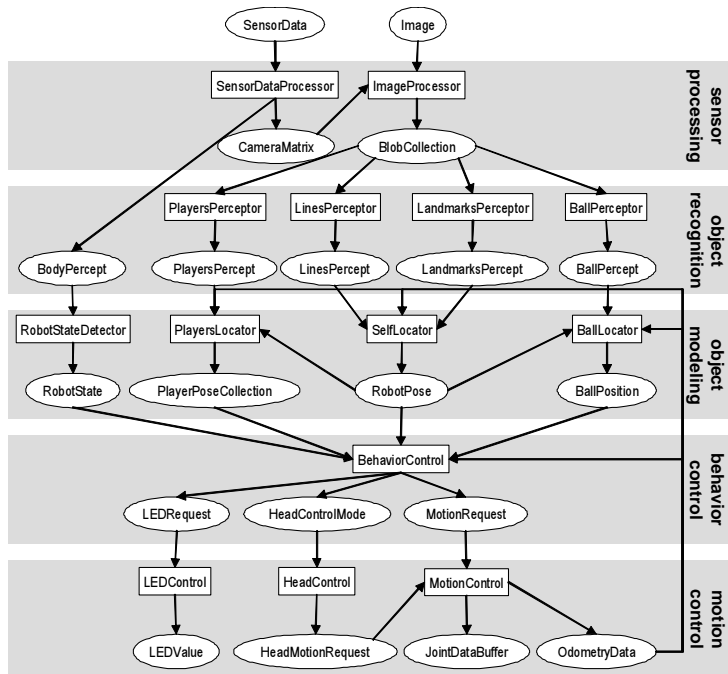
Fig. 1. The modules implemented by the GermanTeam 2002

flags), field lines, other players, and the ball. The sensor readings that were associated to objects are called *percepts*.

Object Modeling. Percepts immediately result from the current sensor readings. However, most objects are not continuously visible, and noise in the sensor readings may even result in a misrecognition of an object. Therefore, the positions of the dynamic objects on the field have to modeled, i.e. the location of the robot itself, the poses (i.e. the (x, y, θ) positions on the field) of the other robots, and the position of the ball. The result of this level is the estimated *world state*.

Behavior Control. Based on the world state and the role of the robot, the fourth level generates the behavior of the robot. This can either be performed very reactively, or deliberative components may be involved. The behavior level sends requests to the fifth level to perform the selected motions.

Motion Control. The final level performs the motions requested by the behavior level. It distinguishes between motions of the head and of the body (i.e. walking). When walking or standing, the head is controlled autonomously, e.g., to find the ball or to look for landmarks, but when a kick is performed, the movement of the head is part of the whole motion. The motion module also performs dead reckoning and provides this information to many other modules.

2.2 Debugging Support

One of the basic ideas of the architecture is that multiple solutions exist for a single task, and that the developer can switch between them at runtime. In addition, it is also possible to include additional switches into the code that can also be triggered at runtime. The realization is an extension of the debugging techniques already implemented in the code of the GermanTeam 2001 [2]: *debug requests* and *solution requests*. The system manages two sets of information, the current state of all *debug keys*, and the currently active solutions. Debug keys work similar to C++ preprocessor symbols, but they can be toggled at runtime. A special infrastructure called *debug queues* is employed to transmit requests to all processes on a robot to change this information at runtime, i.e. to activate and to deactivate debug keys and to switch between different solutions. The debug queues are also used to transmit other kinds of data between the robot(s) and the debugging tool on the PC (cf. section 3). For example, motion requests can directly be sent to the robot, images, text messages, and even drawings can be sent to the PC. This allows to effectively visualize the state of a certain module, textually and even graphically. These techniques work both on the real robots and on the simulated ones (cf. section 3.1).

2.3 Process-Layouts

As already mentioned, each team can group its modules together to processes of their own choice. Such an arrangement is called a *process layout*. The GermanTeam 2002 has developed its own model for processes and the communication between them:

Communication between Processes. In the robot control program developed by the GermanTeam 2001 for the championship in Seattle, the different processes exchanged their data through a shared memory [2], i.e., a blackboard architecture [3] was employed. This approach lacked of a simple concept how to exchange data in a safe and coordinated way. The locking mechanism employed wasted a lot of computing power and it only guaranteed consistency during a single access, but the entries in the shared memory could still change from one access to the other. Therefore, an additional scheme had to be implemented, as, e.g., making copies of all entries in the shared memory at the beginning of a certain calculation step to keep them consistent. In addition, the use of a shared memory is not compatible to the new ability of the Sony Aibo robots to exchange data between processes via a wireless network.

The communication scheme introduced in GT2002 addresses these issues. It uses standard operating system mechanisms to communicate between processes, and therefore it also works via the wireless network. In the approach, no difference exists between inter-process communication and exchanging data with the operating system. A single line of code is sufficient to establish a communication link. A predefined scheme separates the processing time into a communication phase and a calculation phase.

The inter-object communication is performed by *senders* and *receivers* exchanging *packages*. A sender contains a single instance of a package. After it was directed to send the package, it will automatically transfer it to all receivers as soon as they

Fig. 2. Process layout of the Bremen Byters. The broken lines indicate the debugging part.

have requested the package. Each receiver also contains an instance of a package. The communication scheme is performed by continuously repeating three phases for each process:

1. All receivers of a process receive all packages that are currently available.
2. The process performs its normal calculations, e.g. image processing, planning, etc. During this, packages can already be sent.
3. All senders that were directed to transmit their package and have not done it yet will send it to the corresponding receivers, if they are ready to accept it.

    Note that the communication does not involve any queuing. A process can miss to receive a certain package if it is too slow, i.e., its computation in phase 2 takes too much time. In this aspect, the communication scheme resembles the shared memory approach. Whenever a process enters phase 2, it is equipped with the most current data available.

    The whole communication is performed automatically; only the connections between senders and receivers have to be specified. In fact, the command to send a package is the only one that has to be called explicitly. This significantly eases the implementation of new processes.

Different Layouts. The figures 2 and 3 show two different process layouts. Both contain a debug process that is connected to all other processes via debug queues. Note that debug queues are transmitted as normal packages, i.e. a package contains a whole queue. Comparing the two process layouts, it can be recognized that on the one hand, the Bremen Byters try to parallelize as much as possible; on the other hand, Humboldt 2002 focuses on using only a few processes, i.e. the first four levels (cf. Fig. 1) are all integrated into the process *Cognition*. In the layout of the Bremen Byters, one process is used for each of the levels one, four, and five, and three processes implement parts of the levels two and three, i.e. the recognition and the modeling of individual aspects of the world state are grouped together. Odometry is used to decompose information that is dependent: although both the *players* process and the *ball* process require the current pose of the robot, they can run in parallel to the self-localization process, because the odometry can be used to estimate the spatial offset since the last absolute localization. This allows running the ball modeling with a high priority, resulting in a fast update rate, while the self-localization can run as a background process to perform a computationally expensive probabilistic method as, e.g., the one described in [4] or the method used by the GermanTeam 2001 [2].

Fig. 3. Process layout of Humboldt 2002. The broken lines indicate the debugging part.

Currently, it is not known which process layout will be the more successful one. The Darmstadt Dribbling Dackels are using a third approach that is a compromise between the two discussed here, and all three will compete against each other at the German Open. So the best can be used for the world championship.

3 Development Tools on the PC

Two tools were implemented on the PC to ease the development of the robot control programs. The first is a 3-D simulator, and the second is a general tool that provides nearly any support imaginable, even the simulator is integrated.

3.1 SimRobot

SimRobot is a kinematic robotics simulator that was developed in the author's group [1]. It is written in C++ and is distributed in public domain [6]. It consists of a portable simulation kernel and platform specific graphical user interfaces. Implementations exist for the *X Window System*, *Microsoft Windows 3.1/95/98/ME/NT/2000/XP*, and *IBM OS/2*. Currently, only the development for the 32 bit versions of Microsoft Windows is continued.

A simulation in SimRobot consists of three parts: the simulation kernel, the graphical user interface, and a controller that is provided by the user. The German-Team 2002 has implemented the whole simulation of up to eight robots including the inter-process communication described in section 2.3 as such a controller, providing the same environment to robot control programs as they will find on the real robots. In addition, an object called *the oracle* provides information to the robot control programs that is not available on the real robots, i.e. the robots own location on the field, the poses of the teammates and the opponents, and the position of the ball. On the one hand, this allows implementing functionality that relies on such information before the corresponding modules that determine it are completely implemented. On the other hand, it can be used by the implementators of such modules to compare their results with the correct ones.

3.2 RobotControl

RobotControl is the successor of *DogControl*, the debugging tool used by the GermanTeam in 2001 [2]. Its purpose is to integrate all functionality that is required during the development of the control programs for the Sony Aibo robots.
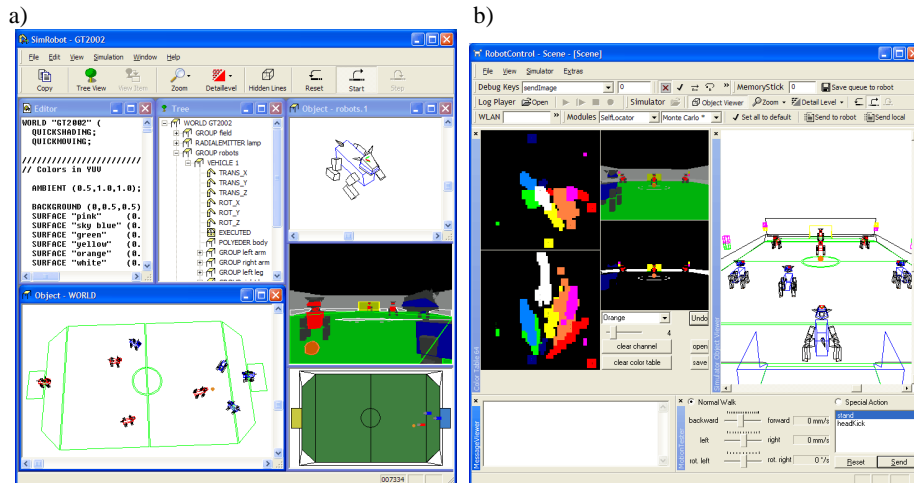
Fig. 4. a) SimRobot simulating the GermanTeam 2002. b) RobotControl: color tool, simulation, message viewer, and motion tester.

Running Robot Controllers. First of all, RobotControl has the ability to run the process-layouts that make up the robot control programs. The simulation kernel of SimRobot is integrated into RobotControl, but in contrast to SimRobot, the robot controller cannot only be provided with simulated inputs. It is also possible to connect it to real robots via the wireless network, and, as a third possibility, the inputs can be generated by replaying log files.

Log Files can either be recorded with a robot, storing them on a memory stick, or the data can be transferred from the robot via the wireless network to RobotControl, and then it will be recorded to a file on the harddisk of the PC. As the same robot controller code runs in all environments, even a simulated robot can produce log files. Log files can contain sensor data and intermediate data as, e.g., blob collections. RobotControl is able to replay log files in real-time or step by step. As RobotControl can run under a debugger, all normal debugging features are available (setting breakpoints, inspecting variables, etc.).

Extensibility. The main purpose of RobotControl is to function as the user interface of the Aibo robots. Therefore, it provides the infrastructure to easily add new toolbars and dialogs. The window layout of RobotControl is always stored and restored on restart. Figure 4b shows a screenshot of RobotControl, in which several toolbars and dialogs can be seen. Toolbars control the replay of log files, they control running the simulation, they allow sending debug keys to real or simulated robots, provide the ability of switching solutions and configuring the wireless network, etc. Dialogs allow generating color tables (for image segmentation, shown in Fig. 4b left), display the simulator scene (Fig. 4b right), control the motions of the robot (Fig. 4b lower right pane), and display debug messages (Fig. 4b lower left pane) and debug drawings (Fig. 4b center). As a result, RobotControl is a very powerful and flexible tool.

4 Conclusion and Future Work

The paper has presented the architecture used by the GermanTeam 2002 in the Sony Legged Robot League. The architecture has been designed for a national team, i.e. a team from different universities that compete against each other in national contests, but that will form a single team at the international RoboCup world championship. The architecture is currently implemented on two different systems, i.e. the Sony Aibo robots and on Microsoft Windows—integrated into the simulator SimRobot and the control software RobotControl. SimRobot is the first 3-D simulator used in the Sony Legged League, and has also been integrated into RobotControl, a universal tool to support the development of the robot soccer team.

Acknowledgements

References

1. Burkhard, H.-D., Düffert, U., Jüngel, M., Lötzsch, M., Koschmieder, N., Laue, T., Röfer, T., Spiess, K., Sztybryc, A., Brunn, R., Risler, M., v. Stryk, O.: GermanTeam 2001. Technical report. Only available online: http://www.tzi.de/kogrob/papers/GermanTeam2001report.pdf (2001).
2. Brunn, R., Düffert, U., Jüngel, M., Laue, T., Lötzsch, M., Petters, S., Risler, M., Röfer, T., Spiess, K., Sztybryc, A.: GermanTeam 2001. In *RoboCup 2001*. Lecture Notes in Artificial Intelligence. Springer (2001), to appear.
3. Jagannathan, V., Dodhiawala, R., Baum, L.: *Blackboard Architectures and Applications*. Academic Press, Inc. (1989).
4. Lenser, S., Veloso, M.: Sensor resetting localization for poorly modeled mobile robots. In *Proc. of the IEEE International Conference on Robotics and Automation* (2000).
5. Röfer, T.: Strategies for Using a Simulation in the Development of the Bremen Autonomous Wheelchair. In Zobel, R., Moeller, D. (Eds.): *Simulation-Past, Present and Future*. Society for Computer Simulation International (1998) 460-464.
6. SimRobot homepage. http://www.tzi.de/simrobot.