

# Hierarchical Reactive Control for Soccer Playing Humanoid Robots

Sven Behnke, Jörg Stückler, Hauke Strasdat, and Michael Schreiber  
*University of Freiburg, Computer Science Institute  
 Germany*

## 1. Introduction

What drives thousands of researchers worldwide to devote their creativity and energy to make robots kick a ball into a goal? The answer lies not only in the fascination of the soccer game, but rather in the quest to advance the fields of artificial intelligence research and robotics. AI researchers started to investigate games early-on. Already in the 1950th, Simon predicted that computers would be able to win against the human chess world champion within ten years (Simon & Newell, 1958). Playing chess was viewed as the epitome of intelligence. The dominant view at that time was that human intelligence could be simulated by manipulating symbols. While the chess world champion was defeated by a machine in 1997 (Newborn, 1997), human intelligence is still far from being understood.

The basis for intelligent action is the perception of the world. Already this seemingly easy task frequently exceeds the capabilities of current computer systems. Perceptual processes, which interpret the flood of stimuli streaming into our senses and make it accessible for behavior control, are mostly unconscious. Hence, we are not aware of the difficulties involved. The performance of our perceptual system becomes clear only when trying to solve the same task with machines. This applies to behavior control as well. Human locomotion, for example, does not seem to be problematic. That walking and running on two legs is not an easy task becomes clear only when one tries to implement it on a real robot.

Based on these observations, a view on intelligence has established itself over the last two decades that does not rely on manipulating symbols, but emphasizes the interaction of an agent with its environment (Brooks 1990; Pfeifer & Scheier 1999). The term embodiment stresses the importance of having a body as the physical basis for intelligence. Situatedness of an agent in a rich environment enables feedback from its actions to its sensory signals. The complexity of the interaction is increased significantly when the environment does not only contain passive objects, but other agents as well.

### 1.1 RoboCup Competitions

Motivated by the success in the chess domain, the RoboCup Federation organizes since 1997 international robotic soccer competitions. Similar competitions are organized by FIRA. The long-term goal of RoboCup is to develop by the year 2050 a team of humanoid soccer robots that wins against the FIFA world champion (Kitano & Asada, 2000). The soccer game was

selected for the competitions, because, as opposed to chess, multiple players of one team must cooperate in a dynamic environment. The players must interpret sensory signals in real-time, select appropriate actions, and execute them. The soccer competitions do not test isolated components, but during a match two systems compete with each other. The number of goals scored is an objective performance measure for comparing systems that implement a large variety of approaches to robot construction, perception, and behavior control. The presence of opponent teams, which continuously improve their system, raises the bar every year. Such a challenge problem focuses the effort of many research groups worldwide and facilitates the exchange of ideas.

The RoboCup championships grew to the most important robotic competition worldwide. In the last RoboCup, which took place in June 2006 in Bremen, Germany, 440 teams from 36 countries competed, not only in RoboCupSoccer, but also in RoboCupRescue, RoboCupJunior, and RoboCup@home. The total number of participants was more than 2.600.

## 1.2 Humanoid League



Fig. 1. Some of the RoboCup 2006 Humanoid League participants.

The RoboCupSoccer competitions are held in five leagues for simulated, wheeled, four-legged, and biped robots. The Humanoid League was established in 2002. Here, robots with a human-like body plan compete with each other. Fig. 1 shows some participants of the 2006 competition. The players must have two legs, two arms, a head, and a trunk. Size restrictions make sure that the center of mass of the robots is not too low, that their feet are not too large, and so on. The participants are grouped in two size classes: KidSize (<60cm) and TeenSize (>80cm). The humanoid robots must be able to walk on two legs, and they must be fully autonomous. They may communicate with each other via a wireless network. Help from outside the field is not permitted, neither by humans nor by computers.

Because the construction and the control of humanoid robots is significantly more complex than that of wheeled robots, initially, there were only preliminary competitions held, but no

soccer games played, in the Humanoid League. The robots had to footrace around a pole and faced each other in penalty kicks. Since 2005, 2 vs. 2 soccer games take place in the Kid-Size class with rules derived from the FIFA laws. Some simplifications apply, however. For example, the offside rule is not observed and key objects are color-coded.

The complexity of playing soccer games is much higher than the complexity of kicking penalties. The ball might be at any position on the field and the robots need to search for it if they lost track of its position. The robots must also perceive at least the two goals and the other players. Higher-level behaviors require self-localization on the field. The distances to walk are much longer. Hence, the walking speed must be higher. As two robots play together, there is need for coordination. While some teams use one dedicated goalie and one field player, other teams use two field players. This makes dynamic role assignment necessary. Last, but not least, in soccer games robots of the two teams interact physically when going for the ball. This disturbs walking and leads to falls. The robots need to get up from the ground by themselves in order to continue play. As a result of these difficulties, in the RoboCup 2006 competition, only a fraction of the teams able to play penalty kick was able to play decent soccer games.

The other RoboCupSoccer leagues have been facing the complexity of soccer games for some years now. There, tools for structured behavior engineering have been developed. For example, Jaeger and Christaller proposed the Dual Dynamics architecture (Jaeger & Christaller, 1998), which has been used in the MiddleSize League. The architecture distinguishes elementary behaviors, which implement a target dynamics, and complex behaviors, which control the activation of elementary behaviors. Another tool used in the MiddleSize League is the BAP-framework (Utz et al., 2005), which allows for specifying hierarchical, event-driven, behavior-based control systems. In the Four-Legged League, the German Team developed XABSL (Lötzsch et al., 2004). It allows for XML-based specification of hierarchies of behavior modules that contain state machines for decision making. State transitions are modeled as decision trees. Parts of the German Team system are used now in the Humanoid League by Darmstadt Dribblers, Humanoid Team Humboldt, and BreDoBrothers. Another example for a behavior architecture used in more than one league is the architecture proposed by Laue and Röfer (Laue & Röfer, 2005), which combines action selection and potential field motion planning. It was used to control SmallSize and Aibo soccer robots.

To implement the behavior control software for the humanoid soccer robots of our team NimbRo, we used a framework that supports a hierarchy of reactive behaviors (Behnke & Rojas, 2001). This framework has been originally developed for the FU-Fighters SmallSize robots. It was later adapted to the FU-Fighters MiddleSize robots and also used by CMU in the Four-Legged League. We adapted it for the control of soccer playing humanoid robots by extending the agent-hierarchy to: joint – body part – player – team. The lowest levels of this hierarchy contain position control of individual joints and kinematic interfaces for body parts. At the next level, basic skills like omnidirectional walking, kicking, and getting-up behaviors are implemented. These are used at the player level by soccer behaviors like searching for the ball, approaching the ball, avoiding obstacles, and defending the goal. Finally, on the team level, the players communicate via a wireless network to share information about the world state and to negotiate roles like attacker and defender.

The remainder of this chapter is organized as follows. In the next section, we describe our robots. We cover mechanical design, electronics, and perception. Sec. 3 describes our behav-

ior control framework. The implementation of basic skills is covered in Sec. 4. Sec. 5 explains the design of our soccer behaviors. Finally, we present the results of using the soccer behaviors at RoboCup 2006.

## 2. NimbRo 2006 Robots

### 2.1 Mechanical Design

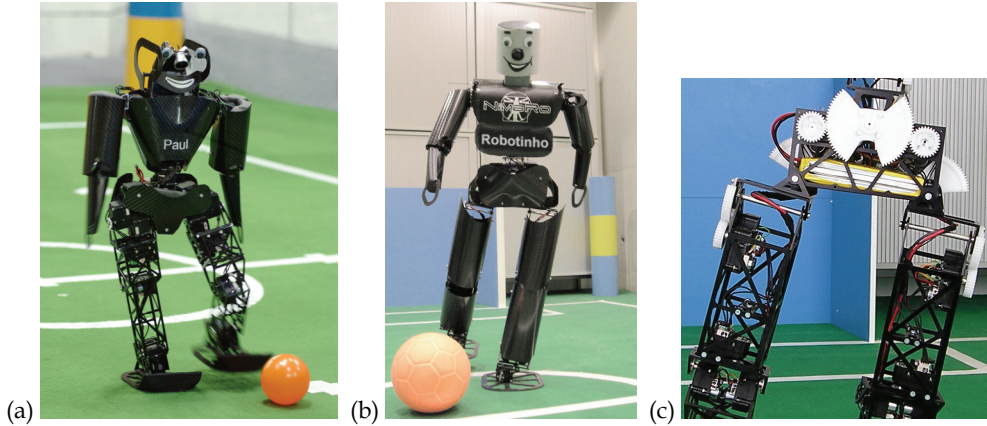


Fig. 2. NimbRo 2006 robots: (a) KidSize robot Paul; (b) TeenSize robot Robotinho; (c) close-up of Robotinho's mechanics.

Fig. 2 shows Paul, one of our 2006 KidSize robots, and Robotinho, our 2006 TeenSize robot. As can be seen, the robots have human-like proportions. Their mechanical design focused simplicity, robustness, and weight reduction. The KidSize robots have a height of 60cm and a weight of only 2.9kg, including batteries. They are driven by 24 Dynamixel actuators: 8 per leg, 3 in each arm, and two in the trunk. For the leg and the trunk joints, we use the DX-117 actuators (66g, 37kg-cm). Three orthogonal axes constitute the 3DOF hip joint. For the hip pitch and roll axes, we use two of these actuators in parallel. The actuators are coupled in a master-slave configuration. This doubles the torque and lowers operating temperatures. The master-slave pair of actuators has the same interface as the single actuators used for all other joints. Two orthogonal servos form the 2DOF ankle joint. One servo drives the knee joint. The trunk joints are in the pitch and yaw axes. The arms do not need to be as strong as the legs. They are powered by DX-113 actuators (58g, 10.2kg-cm). Two orthogonal servos constitute the shoulder joint and one servo drives the elbow joint.

The TeenSize robot Robotinho is 100cm tall and has a total weight of about 5kg. Its 21 DOF are driven by a total of 33 DX-117 actuators. The additional joint is the roll axis in the trunk. All joints in the legs and the trunk, except for the yaw axes, are driven by two parallel actuators. The hip and trunk yaw axes are reinforced by external 2:1 spur gears. The hip and trunk roll axes are reduced by 3:1, resulting in a holding torque of 222kg-cm at 16V.

The skeleton of the robots is constructed from aluminum extrusions with rectangular tube cross section. In order to reduce weight, we removed all material not necessary for stability. The feet and the forearms are made from sheets of carbon composite material. The elasticity

of the feet and the carpet, the robots walk on, helps to maintain non-degenerate foot-ground contact, even when the supporting foot is not parallel to the ground. Robotinho's head is made of lightweight foam. The upper part of the KidSize robots and the entire body of the TeenSize robot is protected by a layer of foam and an outer shell of thin carbon composite material.

## 2.2 Electronics

Our soccer robots are fully autonomous. They are powered by high-current Lithium-polymer rechargeable batteries, which are located in their hip. Four Kokam 910mAh cells are used for the KidSize robots. Robotinho has four Kokam 3200mAh cells. The batteries last for about 25 minutes of operation. The Dynamixel actuators have a RS-485 differential half-duplex interface. Each robot is equipped with a CardS12 microcontroller board, which manages the detailed communication with all Dynamixels. These boards feature the Motorola MC9S12D64 chip, a 16-bit controller belonging to the popular HCS12 family. We clock it with 32MHz. It has 4kB RAM, 64kB flash, two serial interfaces, CAN bus, 8 timers, 8 PWM channels, and 16 A/D converters.

The Dynamixel actuators can be configured in a flexible way. Not only target positions are sent to the actuators, but also parameters of the control loop, such as the compliance. In the opposite direction, the current positions, speeds, loads, temperatures, and voltages are read back. In addition to these joint sensors, each robot is equipped with an attitude sensor, located in the trunk. It consists of a dual-axis accelerometer (ADXL203,  $\pm 1.5g$ ) and two gyroscopes (ADXRS 300,  $\pm 300^\circ/s$ ). The four analog sensor signals are digitized with A/D converters of the HCS12 and are preprocessed by the microcontroller. The microcontroller communicates with the Dynamixels at 1MBaud and with a main computer via a RS-232 serial line at 115KBaud. Every 12ms, target positions and compliances for the actuators are sent from the main computer to the HCS12 board, which distributes them to the actuators. The microcontroller sends the preprocessed sensor readings back. This allows keeping track of the robot's state in the main computer.

We use a Pocket PC as main computer, which is located in the upper part of the robots. The FSC Pocket Loox 720 has a weight of only 170g, including the battery. It features a 520MHz XScale processor PXA-272, 128MB RAM, 64MB flash memory, a touch-sensitive display with VGA resolution, Bluetooth, wireless LAN, a RS-232 serial interface, and an integrated 1.3 MPixel camera. This computer runs behavior control, computer vision, and wireless communication. It is equipped with a Lifeview FlyCam CF 1.3M that has been fitted to an ultra-wide-angle lens. Robotinho's FlyCam lens also serves as nose. It looks in forward direction. For the KidSize robots, we took the integrated camera out of the Pocket PC and connected it via an extension cable. This camera uses the QuickCapture feature of the XScale chipset. Images of size  $640 \times 480$  can be captured at 15fps using DMA. The camera is fitted to a wide-angle converter. Located above the Pocket PC, it looks in forward direction. The FlyCam is looking in backward direction in the KidSize robots.

## 2.3 Perception

Our robots need information about themselves and the situation on the soccer field to act successfully.

- **Proprioception:** The readings of accelerometers and gyros are fused to estimate the robot's tilt in roll and pitch direction. The gyro bias is automatically calibrated and the low-frequency components of the tilt estimated from the accelerometers are combined with the integra-



ted turning rates to yield an estimate of the robot's attitude that is insensitive to short linear accelerations. As described above, joint angles, speeds, and loads are also available. Temperatures and voltages are monitored to notify the user in case of overheating or low batteries.

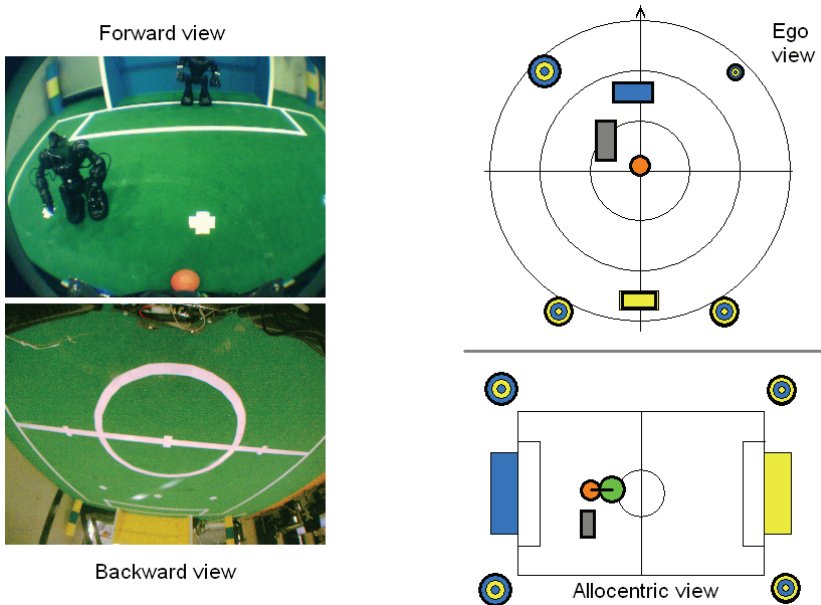


Fig. 3. Left: Images of the two cameras mounted on the robot. Upper right: Egocentric coordinates of key objects (ball, goals, corner poles, and obstacle) detected in the image. Lower right: Localization of the robot, the ball, and the obstacle on the soccer field.

- **Visual Object Detection:** The only source of information about the environment for our robots is their camera. The wide field of view of the cameras allows the robots to see their own feet and objects above the horizon at the same time (see left part of Fig. 3). Our computer vision software detects the ball, the goals, the corner poles, and other players based on their color in YUV space. Using a look-up table, the colors of individual pixels are classified into color-classes that are described by ellipses in the UV-plane. In a multistage process, we discard insignificant colored pixels and detect colored objects. We estimate their coordinates in an egocentric frame (distance to the robot and angle to its orientation), based on the inverted projective function of the camera. We correct first for the lens distortion and invert next the affine projection from the ground plane to the camera plane. The estimated egocentric coordinates of the key objects are illustrated in the upper right part of Fig. 3. Here, the objects detected by both cameras are fused, based on their confidence. The objects are also combined with previous observations, which are adjusted by a motion model, if the robot is moving. This yields a robust egocentric world representation.

- **Ball Tracking:** The limited computing power of the Pocket PC does not allow for processing all images at the frame rate of the camera (15fps). Because the ball is the most important object on the field, we implemented a tracking procedure for it. If the ball could be detected in the previous frame, a small window is placed at the predicted ball position. Only this

window is analyzed for every frame. Every third frame is processed entirely to detect the other objects. The full frame is also processed if the vision system loses track of the ball. Fig. 4 illustrates a typical problem when processing images captured from a walking robot. The walking induces camera motion that causes motion blur in the image. Because the orange of the ball blends with the green carpet to a brownish color, we use such a candidate color to detect blurred balls. In this case, however, it is important to make sure that a brownish color blob is surrounded by green carpet, in order to prevent false positive detections caused by brownish objects outside the field.




Fig. 4. The orange ball behind a white line on a green field. (a) Clean image captured from a standing robot. (b) Same situation with motion blur due to humanoid walking movements. The green blends with the orange to a brownish color.

• **Self-Localization:** The relative coordinates suffice for many relative behaviors like positioning behind the ball while facing the goal. To keep track of non-visible goals or to communicate about moving objects with other team members, we need the robot coordinates in an allocentric frame ( $(x, y)$  -position on the field and orientation  $\theta$ ). We solve self-localization by triangulation over pairs of landmark observations, i.e. detected goals and corner poles. When observing more than two landmarks, the triangulation results are fused based on their confidence. Again, the results of self-localization are integrated over time and a motion model is applied. The lower-right of Fig. 3 illustrates the resulting allocentric representation.

### 3. Behavior Architecture

We control the robots using a framework that supports a hierarchy of reactive behaviors (Behnke & Rojas, 2001). This framework allows for structured behavior engineering. Multiple layers that run on different time scales contain behaviors of different complexity. When moving up the hierarchy, the speed of sensors, behaviors, and actuators decreases. At the same time, they become more abstract. This is illustrated in Fig. 5.

The framework forces the behavior engineers to define abstract sensors that are aggregated from faster, more basic sensors. One example for such an abstract sensor is the robot's attitude that is computed from the readings of accelerometers and gyros. Abstract actuators give higher-level behaviors the possibility to configure lower layers in order to eventually influence the state of the world. One such abstract actuator is the desired walking direction, which configures the gait engine, described below, implemented in the lower control levels. The behaviors within one layer of the behavior framework are activated according to the current state of its sensors. Activation is indicated by an activation factor in the interval  $[0, 1]$ . Each active behavior can manipulate the actuators in its layer. If multiple behaviors try to




Fig. 5. Sketch of the hierarchical framework for reactive control. Complex behaviors are evaluated less often than elementary behaviors. They make decisions based on aggregated fast sensors or rely on slow physical sensors. Complex behaviors use slow actuators to configure lower levels or to directly influence the environment.

manipulate the same actuator, the actuator is set to the weighted sum of desired values, where the activation factors are used as weights. To prevent conflicting behaviors from being active at the same time, behaviors can inhibit other behaviors. If an inhibiting behavior is not completely active, the inhibited behaviors share the remaining activation, such that the activation factors sum to one.

The control hierarchy of our soccer robots is arranged in an agent hierarchy:

- multiple joints (e.g. left knee) constitute a body part (e.g. left leg),
- multiple body parts constitute a player (e.g. field player), and
- multiple players constitute a team.

In our system, two teams can be controlled simultaneously. The behavior framework manages all but the motor control loop within the Dynamixel actuators, which has been implemented by Robotis. The behaviors on the lower level in the framework implement basic skills which generate target positions for individual joints at a rate of 83.3Hz.

To abstract from the individual joints, we implemented here a kinematic interface for the body parts. The leg interface, for example, allows to independently change leg extension  $\eta$ , leg angle  $\theta_{\text{Leg}}$ , and foot angle  $\theta_{\text{Foot}}$  as illustrated in Fig. 6. A detailed description of the kinematic leg interface is given in (Behnke, 2006).




Fig. 6. Kinematic interface to a leg.



## 4. Basic Skills

Several basic skills use this kinematic interface. Fundamental for playing soccer are the abilities to walk and to kick. As body contact between the physical agents is unavoidable, the capability of getting up after a fall is also essential. To act as a goalkeeper, the robot must be able to perform special motions. The basic skills are implemented on the body part layer. Fig. 12 illustrates the inhibitory structure of the basic skills and the interface that they provide for the next higher level of our behavior control system.

### 4.1 Omnidirectional Walking

Omnidirectional locomotion is a concept that has proven to be advantageous in dynamic environments and in restricted spaces. The ability to move into any direction, irrespective of the orientation, and to control the rotational speed at the same time has advantages in many domains, including RoboCupSoccer. Omnidirectional drives are used by most teams in wheeled leagues, and omnidirectional walking is heavily used in the Four-legged League. It is much easier to position robots for kicking and to outmaneuver opponents when using omnidirectional locomotion.

We use the leg interface to implement omnidirectional walking for our humanoid soccer robots. Shifting the weight from one leg to the other, shortening of the leg not needed for support, and leg motion in walking direction are the key ingredients of this gait. In contrast to the low-frequency gait of our 2005 robots (Behnke, 2006), we were able to increase the step frequency significantly to 3.45Hz for the KidSize robots and to 2.44Hz for Robotinho.

Fig. 7 shows in its left part the trajectories generated for forward walking. Note that the leg is not only shortening during swing, but also in the middle of the stance phase. Walking forward, to the side, and rotating on the spot are generated in a similar way. The three basic walking directions can be smoothly combined. The robots are able to walk in every direction and to change their heading direction at the same time. The gait target vector ( $v_x, v_y, v_\theta$ ) can be changed continuously while the robot is walking. This makes it possible to correct for

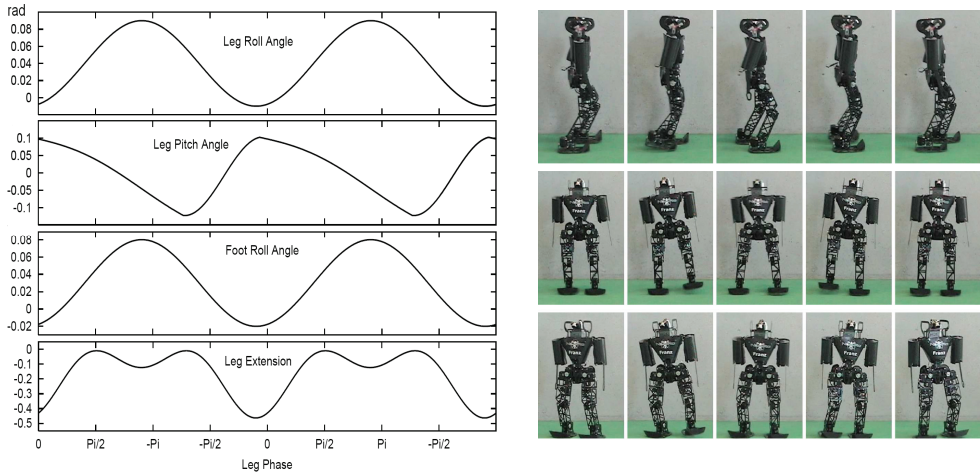


Fig. 7. Trajectories for forward walking of KidSize robots (left) and resulting robot motion during forward, lateral, and rotational walking (right).

deviations in the actual walking direction and to account for changes in the environment by using visual feedback. When using this flexible gait, the maximal forward walking speed of the robots is approx. 25cm/s. The right part of Fig. 7 shows image sequences of the robot Franz walking forward, laterally, and turning. Behaviors of the upper level can control the gait target vector with an actuator that enforces maximal speeds and accelerations.

## 4.2 Kicking

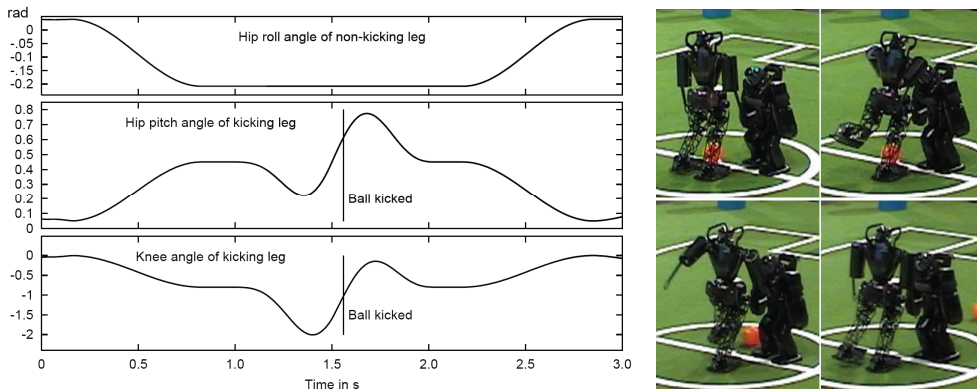


Fig. 8. Trajectories for kicking (left) and resulting robot motion (right).

In addition to walking, we implemented kicking. An actuator allows behaviors in the upper level to trigger kicks with both, the left and the right leg. Fig. 8 shows some of the trajectories generating the kicking motion. After inhibiting the walking behavior and coming to a stop, the robot moves its weight to the non-kicking leg (see hip roll angle). Then, it shortens the kicking leg, swings it back, and accelerates forward. The kicking leg reaches its maximal speed when it comes to the front of the robot. At this point, the hip pitch joint and the knee both move the foot forward and the ball is kicked. The kicking movement continues with deceleration of the foot and slow motion back to the bipedal stand. The resulting kick can be seen in the right part of Fig. 8.

## 4.3 Getting up from the Floor

Since in soccer games physical contact between the robots is unavoidable, the walking patterns are disturbed and the robots might fall. Hence, they must be able to detect the fall, to assess their posture on the ground, and to get back into an upright posture. After falling, the robot's center of mass (COM) projection to the ground is outside the convex hull spanned by the foot-contact points. Additional support points like knees, elbows, and hands must be used in order to move the COM back inside the foot polygon. Using their attitude sensors, the robots detect a fall, classify the prone or supine posture and trigger the corresponding getting-up sequence. We designed the getting-up sequences in the simulator using sinusoidal trajectories (Stückler et al., 2006). Fig. 9 illustrates the four phases of getting up from the prone and the supine posture. The getting-up sequences work very reliably. Under normal circumstances, i.e. appropriate battery voltage, the routines worked with 100 successes in 100 tests.




Fig. 9. Standing up from the supine posture (left) and the prone posture (right).

#### 4.4 Goalkeeper Motions

The goalkeeper is capable of diving into both directions or to bend forward with spread arms. Fig. 10 shows Franz diving to the left. First, it moves its COM and turns its upper body towards the left while shortening the legs. As soon as it tips over its left foot, it starts straightening its body again. While doing so it is sliding on its hands and elbows. The fully extended robot covers the entire goal half. After the dive Franz gets up again, as described above.



Fig. 10. Diving motion of the goalkeeper.

### 5. Playing Soccer

The next higher level of our behavior control framework contains soccer behaviors which are executed at a rate of 41.7Hz. They build on the basic skills and have been designed for 2 vs. 2 soccer games.

#### 5.1 Representation of the Game Situation

The soccer behaviors require knowledge of the current game situation. The visual perception supplies relative distance, angle, and perceptual confidence for the ball, the own goal, the opponent goal, and the nearest obstacle. In the attacking role, the relative position and confidence of the opponent goal is used as the target to kick at (ball-target). The decision for the kicking leg is made at every time step, depending on the relative position of the ball and the line from ball to ball-target, which we denote as ball-to-target-line. If the robot has to approach the ball-to-target-line from the right, it kicks with the left leg, and vice versa. To avoid oscillations it is important that the decision may only be changed if the distance of the robot to the ball-to-target-line exceeds a threshold.

To kick the ball with the chosen leg, the robot has to position itself behind the ball with lateral and sagittal offsets,  $\delta_l$  and  $\delta_s$  that depend on the distance between the legs and the length of the feet. To generate smoothly approaching trajectories, the sagittal offset is in-

creased by an amount  $\delta_a$  that is proportional to the angle between the robot's heading direction and the ball-target. The ball approach is illustrated in Fig. 11.

When playing as defensive field player, the own goal is used as ball-target, such that the position behind the ball is set to a position between ball and own goal. The distance kept to the ball depends on the distance to the own goal. A threshold for the minimal distance to the goal lets the robot stay out of its goal, as long as the ball is still far away. If the ball and the robot are near the goal, the robot keeps behind the ball at a minimum distance.




Fig. 11. Two examples showing sequences of robot poses, target positions behind the ball (blue crosses), and ball positions while approaching the ball with the left leg as kicking leg.

The robot maintains additional hypotheses about the relative ball location that are used for searching the ball. If a kick is triggered, one hypothesis is set in front of the robot at a distance depending on kick strength. The confidence of the hypothesis is discounted by the time since the kick started. Its relative position is altered according to the motion model.

Additionally, hypotheses are maintained for the perceptions of the ball by other players on the field. The confidences of these hypotheses depend on the self-localization and ball perception confidences of the other players and the self-localization confidence of the robot itself.

## 5.2 Soccer Behaviors

According to the current game situation, behaviors like searching the ball, positioning behind the ball, or avoiding obstacles are activated. These behaviors are implemented on the player level and use the actuator interface that the basic skills of the lower layer provide. For example, they set the gait target vector or trigger a kick. Fig. 12 illustrates the inhibitory structure of the soccer behaviors and the actuator interface used for configuring the basic skills.

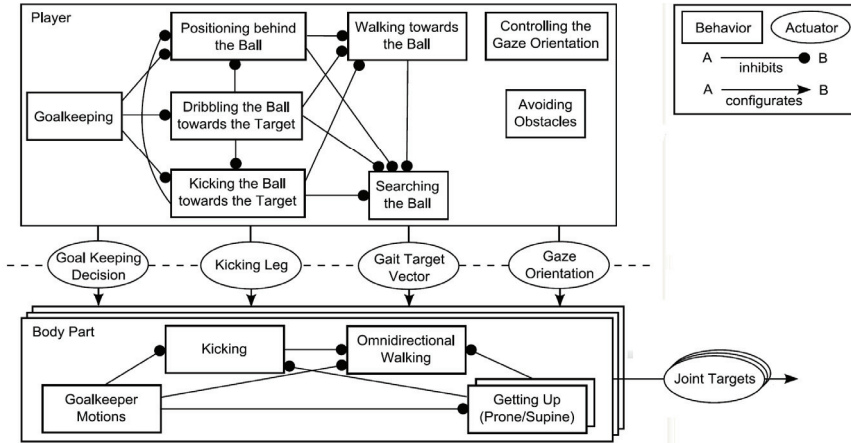


Fig. 12. Behaviors on the body part and player levels of the behavior hierarchy.

• **Searching the Ball:** Exploring the environment for the ball is always active, but inhibited by behaviors that activate when the ball has been perceived with a certain confidence. If a ball hypothesis with confidence over a certain threshold exists, the robot walks towards the most confident hypothesis. Otherwise, it turns towards the most confident hypothesis for a short time. If the ball still is not visible, it starts to walk around the center circle in a constant distance in order to inspect all parts of the field.

• **Walking towards the Ball:** The robot walks straight towards the ball, if it perceives the ball. The own goal must be either not visible or far away to avoid scoring an own goal. This behavior controls the gait target velocity to keep the robot near the ball, e.g. if visual perception fails to detect the opponent goal. The behavior inhibits searching the ball.

• **Positioning behind the Ball:** If the ball and the ball-target are perceived, the robot positions itself behind the ball, facing towards the ball-target. The robot is positioning on the behind-ball-position by controlling the gait target velocity. If the distance to the target position is large, the robot rotates towards the target position, such that it can approach it by mainly combining forward walking with turning. If it is near the target position, the robot aligns itself towards the ball-target. For intermediate distances, the gait rotation is interpolated linearly between both alignment targets. The behavior also handles the case when the ball is located between the robot and the behind-ball-position. Here, the robot walks around the ball by walking towards the target position but avoiding the ball-to-target-line. When playing as defensive field player, the robot rotates towards the ball at any distance. It does not avoid the ball-to-target-line, because the ball-target is the own goal. This behavior inhibits walking towards the ball, such that the inhibited behavior may only be active, if the ball-target has not been perceived. It also inhibits searching the ball.

• **Kicking the Ball towards the Target:** This behavior is activated as soon as the behind-ball position has been reached with a certain precision in angle to the ball-target and in distance to the target position. If the precision conditions hold, a kick is triggered. Obviously, ball and ball-target must be perceived and the own goal must not be in front of the robot. If the ball comes into a kicking position by chance, the behavior initiates a kick with the corresponding leg. As the robot has to come to a complete stop before the kicking motion can be



executed, the robot can cancel the kick, if the ball moves away in the meantime. This behavior inhibits searching the ball, walking towards the ball, and positioning behind the ball.

• **Dribbling the Ball towards the Target:** If positioning behind the ball was not successful for a longer time, or the game started with a kick-off for the player, the robot activates dribbling the ball towards the ball-target for some time. Additional preconditions for activation are that the ball and ball-target are perceived and the angle towards the ball-target is small. Dribbling is performed by steering towards the ball. The forward walking speed is inversely related to the angle to the ball. In combination with positioning behind the ball, the robot is kept behind the ball, facing the ball-target when dribbling. Dribbling inhibits searching the ball, walking towards the ball, and positioning behind the ball. As we want the decision for dribbling to be strict, it also inhibits kicking the ball towards the target.

• **Avoiding Obstacles:** After a fall, the robot needs valuable time to get back on its feet. The main reason for our robots to fall is physical contact with other robots. Hence, obstacle avoidance is an important feature. The visual perception supplies the behavior with the nearest obstacle. If it is detected closely in front of the robot, obstacle avoidance is activated by a factor that interpolates linearly between a minimum and a maximum distance for the obstacle. The avoidance sets the gait target actuator to a constant and a variable part of the direction from obstacle to robot. The strength of the variable part depends on the distance to the obstacle, similar to the activation factor. If the ball is between obstacle and robot, the variable avoidance is weakened, such that the robot moves more aggressively behind the ball. A stuck situation is indicated by a resulting gait target vector that is small in length for a longer time. In this case, the robot may sidestep the obstacle, if the ball is not between the obstacle in the front and the robot and is perceived on one side of the obstacle. The action is cancelled, if either the preconditions for sidestepping do not hold anymore or a certain amount of time has elapsed since sidestepping has been activated. The deactivation of sidestepping after some time is important, because the decision for the sidestep direction is made only once on activation.

• **Controlling the Gaze Direction:** Although the robot has wide-angled views to the front and the rear, it cannot perceive objects on the sides. Thus, a gaze control behavior is always active and primarily keeps the ball within an angular range of  $\pm\pi/4$  by twisting the upper trunk with the trunk yaw joint. If the ball is not visible or within range and the robot is localized, it aligns the upper body with the line between the goals to keep the localization landmarks visible. This is achieved by keeping the angle to the line within the angular range of  $\pm\pi/4$ . The twisting of the trunk is limited to  $\pm\pi/4$ .

• **Goalkeeping:** The goalkeeper's objective apparently is to keep the ball out of the own goal. While the ball is visible and is not in kicking distance to the robot, the goalkeeping behavior is active. Otherwise, the robot behaves like a field player and tries to shoot the ball towards the opponent goal. Hence, goalkeeping inhibits positioning behind the ball, kicking the ball, and dribbling the ball towards the target. Walking towards the ball and searching for the ball is not activated when playing as a goalkeeper. The goalkeeper stands still until it reacts on the ball. Balls close to the robot let it react immediately. It uses the ball angle to determine the appropriate type of motion (diving left/right or bending forward). To achieve fast reaction on an approaching ball, the visual perception supplies the difference of the ball position between the last two images. The magnitude of this vector is interpreted as approaching speed. The goalkeeper does not react on small speeds. The type of the goalkeeper motion is determined by the intersection point of the moving ball direction and the goal line. At kick-

off, the goalkeeper is placed in the goal. After a diving motion, it gets up and repositions itself in the goal while facing the opponent goal.

### 5.3 Team Behaviors

The importance of team behaviors is still low in the Humanoid League, as only two players per team have competed so far. In Bremen 2006, most teams assigned one player to keep the goal clear and used the other player as field player. In our team, the players share perceptions via wireless communication. The ball perceptions communicated by other players are used for search. For the soccer play with two field players, we implemented simple but effective role negotiation between the players. As soon as one of our players has control of the ball, the other player goes to a defensive position between the ball and the own goal. A player takes control of the ball, if it is close to the ball and perceived it with high confidence. It loses control, if the ball gets too far away or has low confidence. The thresholds for taking and losing control implement hysteresis to prevent oscillations of the control state.

## 6. RoboCup 2006 Results

Our robots performed well at RoboCup 2006, where 21 teams from eleven countries competed in the Humanoid League. In the 2 vs. 2 soccer round robin, the KidSize robots played 2 games and scored 12:0 goals. In the quarter final, they won 6:1 against team RO-PE from

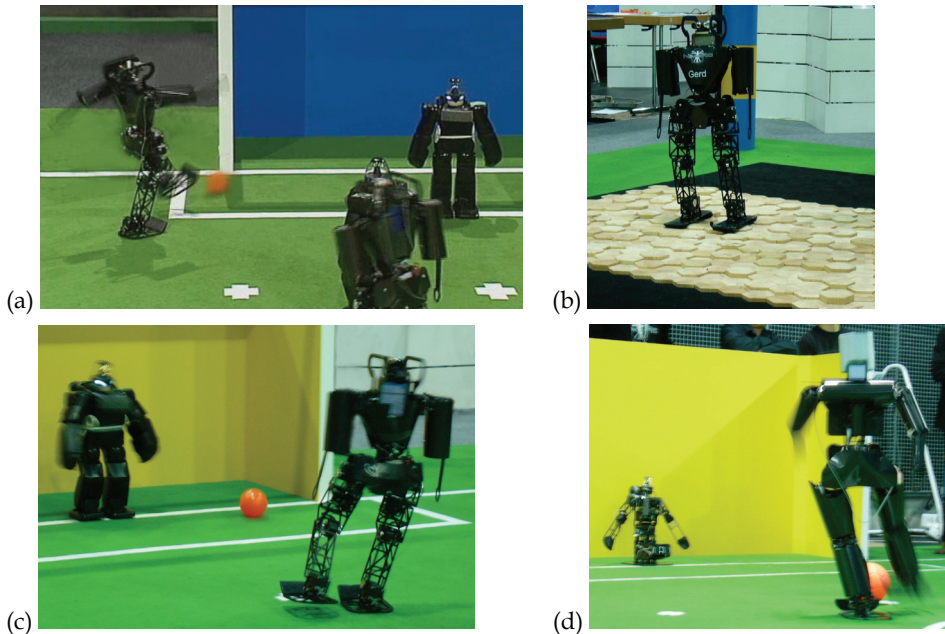


Fig. 13. RoboCup 2006: (a) 2 vs. 2 Soccer final NimbRo vs. Team Osaka. (b) NimbRo robot Gerd walking over rough terrain. (c) KidSize Penalty Kick final NimbRo vs. Team Osaka. (d) TeenSize Penalty Kick final NimbRo vs. Team Osaka.

Singapore. They met the German-Japanese team Darmstadt Dribblers and Hajime in the semi-final. Our robots won 6:2. The final game was between our robots and Team Osaka, as in 2005. Our robots played well in the first half and scored a lead of 4:0. Fig. 13(a) shows one of the shots. After a goal directly from kick-off, the score at halftime was 4:1. Due to hardware problems of our robots, Team Osaka was able to reach a draw of 4:4 after regular playing time. As we already had taken the available two substitutions, we needed to continue playing with impaired robots in the extra time. The final score was 9:5 for Team Osaka.

Our KidSize robots also kicked penalties very reliably. In the Penalty Kick competition they scored in 31 of 34 attempts. In the KidSize Penalty Kick final (Fig. 13(c)) our robots won 8:7 against Team Osaka.

In the technical challenge, our KidSize robot Gerd was one of the two robots able to walk across the rough terrain (Fig. 13(b)). Our KidSize robots also scored in the passing challenge. Our TeenSize robot Robotinho used a simplified version of the KidSize behaviors. It also reached the final of its Penalty Kick competition (Fig. 13(d)).

In the overall Best Humanoid ranking, our KidSize robots came in second, next only to the titleholder, Team Osaka. Videos showing the performance of our robots at RoboCup 2006 can be found at <http://www.NimbRo.net>.

## 7. Conclusions

This chapter described the design of the behavior control software for our humanoid soccer robots, which successfully took part as team NimbRo at the RoboCup 2006 competitions. We implemented the control software in a framework that supports a hierarchy of reactive behaviors. This structure restricts interactions between the system variables and thus reduces the complexity of behavior engineering.

A kinematic interface for body parts made it possible to abstract from individual joints when implementing basic skills like omnidirectional walking. These basic skills made it possible to abstract from body parts when implementing more complex soccer behaviors. At this player level, our humanoid robots are very similar to wheeled or four-legged soccer robots. Finally, at the team level, the players are coordinated through role negotiation.

Playing soccer with humanoid robots is a complex task, and the development has only started. So far, there has been significant progress in the Humanoid League, which moved in its few years from remotely controlled robots to soccer games with fully autonomous humanoids. Indeed, the Humanoid League is currently the most dynamic RoboCupSoccer league. We expect to see the rapid progress continue as more teams join the league. Many research issues, however, must be resolved before the humanoid robots reach the level of play shown in other RoboCupSoccer leagues. For example, the humanoid robots must maintain their balance, even when disturbed. Postural reflexes that are triggered by deviations from the normal walking patterns are one way to minimize the number of falls (Renner & Behnke, 2006).

In the next years, the speed of walking must be increased significantly. We work on automatic gait optimization to increase both speed and stability. At higher speeds, running will become necessary. We recently started to explore this direction. The visual perception of the soccer world must become more robust against changes in lighting and other interferences. We continuously improve our computer vision software to make it more reliable.

The 2006 competition has shown that most teams were able to kick penalties, but that soccer games are much richer and more interesting. In the team leader meeting after the competition, the majority voted for abandoning penalty kick as a separate competition. Instead, the KidSize teams will focus on soccer games. Unfortunately, most teams do not feel ready to increase the number of players to more than two per team. This limits the possibilities for team play.

As the basic skills of the humanoid soccer robots improve every year, teams will be able to focus on the more complex soccer behaviors and on team play. This will make structured behavior engineering a key factor for success.

## Acknowledgment

Funding for this project is provided by the DFG (Deutsche Forschungsgemeinschaft), grant BE 2556/2-1,2,4.

## References

- Behnke, S. & Rojas, R. (2001). A Hierarchy of Reactive Behaviors Handles Complexity, In: *Balancing Reactivity and Social Deliberation in Multi-Agent Systems*, pp. 125-136, Springer, LNCS 2103
- Behnke, S. (2006). Online Trajectory Generation for Omnidirectional Biped Walking, *Proceedings of IEEE International Conference on Robotics and Automation (ICRA'06)*, pp. 1597-1603, Orlando, Florida
- Brooks, R. A. (1990). Elephants Don't Play Chess. *Robotics and Autonomous Systems*, 6:3-15
- Jaeger, H. & Christaller, T. (1998). Dual dynamics: Designing behavior systems for autonomous robots. *Artificial Life and Robotics*, 2(3):108-112
- Kitano, H. & Asada, M. (2000). The RoboCup Humanoid Challenge as the millennium challenge for advanced robotics. *Advanced Robotics*, 13(8):723-737
- Laue, T. & Röfer, T. (2005). A Behavior Architecture for Autonomous Mobile Robots Based on Potential Fields, In: *RoboCup 2004: Robot Soccer World Cup VIII*, pp. 122-133, Springer, LNCS 3276
- Löttsch, M.; Bach, J; Burkhard, H.-D. & Jüngel, M. (2004). Designing Agent Behavior with the Extensible Agent Behavior Specification Language XABSL, In: *RoboCup 2003: Robot Soccer World Cup VII*, pp. 114-124, Springer, LNCS 3020
- Newborn, M. (1997). *Kasparov versus Deep Blue: Computer chess comes of age*, Springer
- Pfeifer, R. & Scheier, C. (1999). *Understanding Intelligence*, The MIT Press, Cambridge MA
- Renner, R. & Behnke, S. (2006). Instability Detection and Fall Avoidance for a Humanoid using Attitude Sensors and Reflexes, *Proceedings of IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2967-2973, Beijing, China
- Simon, H. A. & Newell, A. (1958). Heuristic Problem Solving: The Next Advance in Operations Research. *Operations Research*, 6(1):1-10
- Stückler, J.; Schwenk, J. & Behnke, S. (2006). Getting Back on Two Feet: Reliable Standing-up Routines for a Humanoid Robot, *Proceedings of the 9th International Conference on Intelligent Autonomous Systems (IAS-9)*, pp. 676-685, Tokyo, Japan
- Utz, H.; Kraetzschmar, G.; Mayer, G. & Palm, G. (2005). Hierarchical Behavior Organization, *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2598-2605, Edmonton, Canada