Various Uses of Potential Map in a Soccer Game

Marián Lekavý, Jozef Wagner

Slovak University of Technology,
Faculty of Informatics and Information Technologies,
Ilkovičova 3, 842 16 Bratislava, Slovakia
lekavy@fiit.stuba.sk

Abstract. In many environments, it is possible to store the knowledge about the environment in a potential field, instead of expressing the knowledge explicitly. Potential field is also a powerful tool for agent and multi-agent control and coordination. Majority of approaches based on potential fields focus on the navigation and collision avoidance problem or on work or agent distribution. There also exist some approaches to situation evaluation and strategic decision making. This paper presents several approaches to coordination and teamwork in a simulated soccer game, based on the combination of a potential field map and analytical approach. The potential field is used to create a 'fitness map' of the field, while analytical approaches use this map to find the best target and chance of success for individual player skills. Combination of potential field with analytical approaches helps overcoming local extremes, which are a major problem for pure potential field approaches. The presented approaches based on the potential field include team passes (safe passes + predictive pass interception), shooting, dribbling, offside trap avoidance and adapting formation to opponent.

1 Introduction

Various approaches based on the idea of potential fields were used often in the domain of robotics. The information about the environment is converted into a (usually 2D) map of the environment, where high/low values represent good/bad places in the environment. This way, the surrounding world is represented in a more intuitive way, comparing to an array of object representations.

The potential map can be created and used in various ways, depending on the application area. Most applications of potential map (and potential field) focus on finding the best trajectory through the environment or on reactive movement control, letting the agents to move toward the highest/lowest potential. An overview of approaches and applications based on potential fields can be found in the "Related Work" subsection.

This paper presents several approaches based on the idea of a potential field, used in a simulated robotic soccer player in the RoboCup 2D simulation league. The approaches were verified in the simulated player developed at the Faculty of Informatics and Information Technologies in Bratislava [10, 4].

© Václav Snášel (Ed.): Znalosti 2008, pp. 112–123, ISBN 978-80-227-2827-0.
FIIT STU Bratislava, Ústav informatiky a softvérového inžinierstva, 2008.

Unlike most other approaches, we don't rely fully on the potential field, but use it as heuristic information for analytical decision algorithms. This helps us to overcome some problems of pure potential field approaches - the possibility to get stuck in a local extreme when used for reactive behaviour and sensitivity to parameters. It is not possible to remove these problems fully, but we eliminate the worst consequences by not allowing nonsense solutions like passing into a local extreme between two teammates instead of passing to one of them.

We chose the domain of simulated robotic soccer as a coordination and teamwork testing domain for several reasons. The main reason is, that the agents are fully autonomous and have limited communication possibilities. It is difficult to use arranged cooperation, so it is necessary to exploit the possibilities of implicit and predictive cooperation [8].

There are several other constraints that have to be taken into account, making the cooperation problem even more difficult and more interesting. The players act in a real environment, which means that the information provided by their sensors are both inaccurate and incomplete, so the players often need to make their decisions without all the necessary information. At the same time, actions carried out by the players are subject to random noise, so the result of an action is only partially predictable. The predictability of the environment gets even lower, as there are 11 opponent players trying to achieve an opposite goal. Other important constrain is the need to act in real-time, making it hard to use several techniques like deliberation methods or state-space search (especially when the state-space becomes as vast as to be practically infinite).

In section 2, we describe the used potential field, its components and parameters. Section 3 provides an overview of the applications of the potential field, together with their evaluation. In section 4, we describe some improvements and new usage possibilities planned for the future.

1.1 Related Work

Trajectory planning. Since the first approaches [5], much work has been done in this area, including various RoboCup (e.g. [11]) and FIRA ([6]) teams. Various search techniques and motion strategies were developed, including pre-planning and reactive control together with versions suitable for parallel execution (an overview can be found for example here [2]).

In [9], a composite potential field is used to find the best place to run to or kick to. The field is created as superposition of potential fields for players, ball, offside area, penalty area, formation and others. The best value in the composite field is then selected as the target of the next action (dash/dribble/pass). Basically, this is an enhancement of trajectory planning, using a more complex composition of potential fields. The unique idea on this approach is that a single (even if complex) potential field is used to describe the whole behaviour of an agent, even if only at the level of very primitive actions. Similar idea is also used by other teams (e.g. [12]) in a simpler settlement.

Work/agent distribution. If we want to distribute tasks and area among agents, we can set-up a repulsive field for agents and attractive field for tasks which have to be carried out. This way, we make sure, that the agents are placed near the tasks, but far enough from each other [1]. It is necessary to tune the parameters of such potential field and in some cases, it is necessary to implement some higher-level coordination to avoid local extremes.

Potential probing (Electric field [7, 3]). This is a powerful approach, allowing what-if modelling and strategic decision making. It sets-up a probe on an interesting place on the field (e.g. a candidate for a pass) and computes the sum potential on this place. Together with predicting of objects' movement, it can compute the overall fitness of some possible situation on the field, yielding the fitness of an action leading to this situation. This way, possible actions are evaluated and the best is selected. This approach needs a good generator of possible actions, which are then evaluated.

2 Potential Field

The potential map used in our player is a 2D grid with rectangular cells. The centre of the middle cell is in the centre of the play field to obtain symmetry. Here, we define the potential of a field part as high for a good position and low for an undesirable position. We use the analogy with physics, where we have to spend energy and effort to achieve a high peak (good position), while no effort is needed to reach a bad position, as the opponents try to achieve this.

The potential values are a superposition of a static and a dynamic component. The static component is an expression of the fact that we want to get from our defence to opponent's goal. Therefore, the basic potential increases linear with the x-axis towards the opponent goal. Similarly, we want the players' actions to be attracted to the opponent goal (e.g. to cause the player to dribble towards the goal from side of the field, instead of dribbling along the x-axis). Therefore, the values are modified in some distance (34m) from the opponent goal. The same way, the field around our goal is modified to repulse actions away.

The dynamic component of the value is the influence of players. Teammates have positive influence, opponents negative. If there is an unidentified player, we classify him as opponent to avoid unsafe actions.

If the velocity of a player is known, his future position is used in the map. This way, the potential map predicts the situation in the next cycle, improving the accuracy of all decisions.

The influence is configurable by two parameters: the influence on the cell the player stands on ($influencePower$) and the influence radius ($influenceRadius$). With distance, the influence of a player is decreasing linear, until it reaches zero in the distance $influenceRadius$. The basic player contribution to a cell is computed with respect to the centre of the cell, according to Equation 1.

$$basicContribution = influencePower * (1 - \frac{distanceFromPlayer}{influenceRadius}) \qquad (1)$$

The player influence is not identical in all directions. The influence is higher in the direction of player movement. In the ideal case, the influence would be

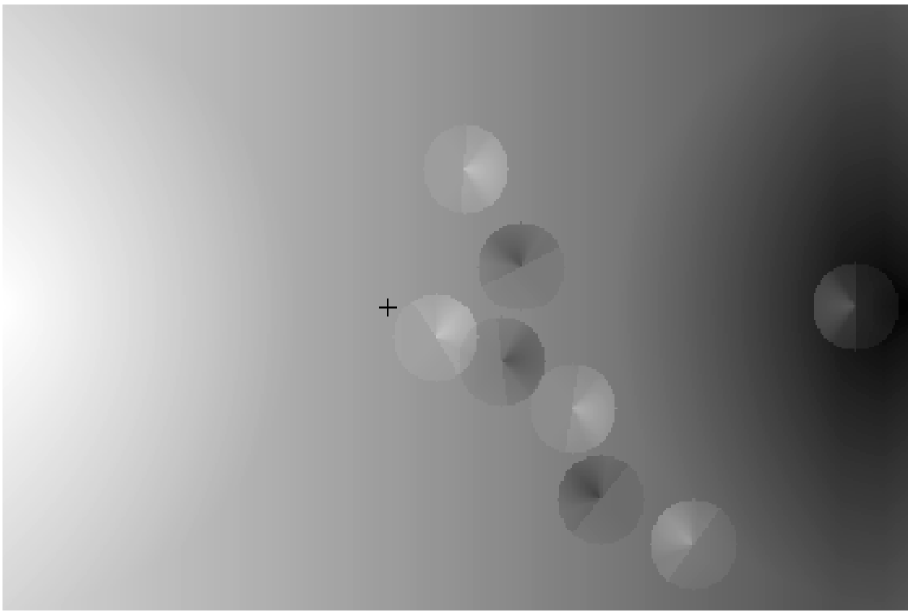

Fig. 1. An example of the potential map. Not all players are on the map, as the visual information is not complete. Cross marks the player creating the map.

dependent on the number of cycles, in which the player is able to reach the place. This would increase the time complexity. Therefore, we used a simplified model. The final contribution is computed according to Equation 2.

$$playerContribution = \begin{cases} basicContribution* \\ *cos(bodyAngle), & |bodyAngle| \in \langle 0, 60) \quad (2a) \\ \dfrac{basicContribution}{2}, & |bodyAngle| \in \langle 60, 90) \quad (2b) \\ \dfrac{basicContribution}{4}, & |bodyAngle| \in \langle 90, 180 \rangle \quad (2c) \end{cases}$$

Tests showed, that the cell size should be within the interval $\langle 0.75m, 2m \rangle$. Higher resolution doesn't increase the game quality, but increases computational complexity significantly. With the resolution 0.75m, the whole team runs on a single PC (Athlon@2GHz) without problems. By using evolutionary algorithm to optimise the average of measured game parameters, we found the optimal values $influencePower = 15$ (with the static values from 0 to 100) and $influenceRadius = 10m$. An example of the potential map is in Figure 1.

All applications of the potential map (described in the section "Potential Field Applications") use the same potential map described above, except the formation adaptation. For the formation adaptation, the coach creates a map with lower resolution (2 meters, 52x34 cells). The cell value is increased by one for every cycle, in which an opponent occupies this cell.

2.1 Optimization

Optimizations of the potential map implementation are used to reduce memory and computational complexity. We use incremental computing, which only changes the map parts influenced by changes. As all objects in the field only have local influence (i.e. they only influence the area in a circle around them), this leads to a major performance improvement.

Other possibility would be to use on-demand computing, that computes only cells needed in the decision process. On-demand computation approach is equivalent to potential-probing algorithms (like the electric field approach [7, 3]). As most actions are localized near the player performing the computation, the time saving would be considerable. Only the formation analyzer in the coach has to compute the whole potential map to get an overview of the entire field. This is, however, not a problem, as the coach uses lower resolution for the statistical evaluation and coach response is not time critical.

3 Potential Field Applications

This section describes the uses of the potential field (described in the previous section). Each of the approaches described in the next six subsections was statistically evaluated (with the other approaches disabled) against multiple opponents and the most important results are described for each approach. All presented results were collected during at least 10 games (i.e. at least 60000 simulation cycles) started with the same parameters to collect significant data.

The composite effect of all changes can be found in Table 1 in the "Conclusion" section. The composite effect is not a simple sum of individual effects, but the approaches reinforce each other. The best example is the shoot improvement by 28%, but together with more intelligent dribbling and passing, the overall shooting success was increased by 139%.

3.1 Pass-ball

The simplest use of the map is evaluating a pass-candidate. For all pass candidates, we compute the average value from all cells on the pass trajectory. The pass with the best average value is then selected. If the highest value is too low, some other action is chosen by higher decision logic.

To avoid the evaluation of passes into empty areas and passes through opponents, the pass candidates have to be chosen carefully. Following algorithm is used to generate pass candidates:

1. Eliminate teammates who are too close or too far.
2. Eliminate teammates who are behind an opponent. If an opponent is inside an isosceles triangle with the teammate in the centre of the opposite side (see Figure 2), then the teammate is removed from the list of candidates.

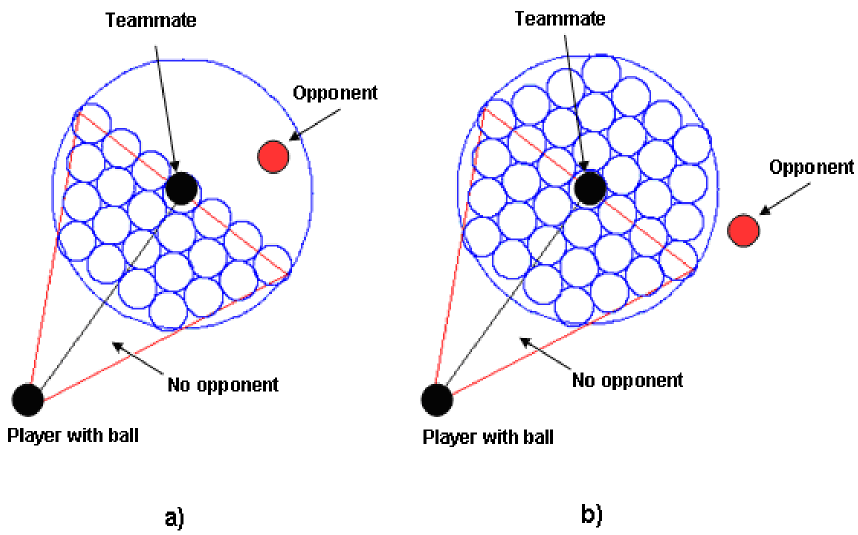

Fig. 2. Pass ball algorithm. A hexagonal grid is used to cover the area around the player. If an opponent is behind the player, only positions in front of the player are evaluated (a), otherwise the whole circle area is evaluated (b). If an opponent is placed in the triangle between the two teammates, the pass is not used.

3. Create sampling points around teammates. We have to choose a point to pass to. The points are selected from a circle area around the teammate. The area grows linearly with distance from the passing player, so for longer passes a larger pass area is taken into account. We use hexagonal grid to cover the whole area (see Figure 2), because hexagonal grid provides the best coverage from all types of grids.
4. Evaluate the sampling points. As already mentioned, the pass value is the average of all values on the pass trajectory.
5. Select the pass with the highest value.

Pass success was increased by 24% (together with local predictive positioning, described below). For direct passes (passes directly to the teammate position), the success increased by 31%. The analytical approach prevents passes through an opponent, to places far from a teammate and into local maxima between teammates. This way, the potential field becomes reliable, because the worst possibilities are simply filtered-out (increasing computational efficiency at the same time). The success was not achieved by being afraid to pass in some situations. In fact, the frequency of passes was increased by 15%.

3.2 Local Predictive Positioning

Local positioning is used to maximize the chance of successfully receiving a pass by changing the player position to "show" himself to the player owning the ball.

A predictive algorithm is used to achieve this. Players without ball predict possible passes by applying the algorithm described in the previous section. The player finds the point, his teammate should pass to, and runs to this point. This way, the player achieves two advantages. First, the player goes to a position, which is suitable for a pass. Second, if his teammate passes to him, he is likely to be already heading for the pass position, saving several simulation cycles and decreasing the probability of pass interception by an opponent.

To disallow the players to drift away from their position in formation, they are only allowed to move in some distance around their home position. Also, movement beyond the offside line and borders of the playfield are not allowed. It would be possible to incorporate the home position into the potential field directly [9]. This way, a player would be attracted to its home position. On the other hand, doing this would tell the player, that his home position is "better" than other positions in the field, which is not true in general. For example, while dribbling, it is necessary to avoid opponents and the home position should be ignored by the dribbling player. Also, highlighting the home position on the map may cause the player to prefer it, even if there is a better position somewhere near. Another reason for restricting the search area is the improvement of computational complexity.

Local positioning was evaluated together with passes, so the results in the previous subsection apply to this subsection too.

3.3 Shoot

Shooting behaviour is, in fact, similar to passing. The difference is, that the shoot target is the opponent goal. All sampling points are placed on the opponent goal. The remaining algorithm steps are identical to pass behaviour - eliminate the points behind an opponent, evaluate the remaining points and select the candidate with the highest value.

When testing the shoot alone (with potential map disabled for other skills), shoot success was increased by 28%. The composition of all approaches increased the shoot success by another 111%. This is mainly caused by dribbling and passing to "better" places of the field, creating better chances to score.

3.4 Offside Trap Overcoming

Similar to pass and shoot, we can also use the pass evaluation algorithm to overcome the opponents' offside trap. The best way to come through the offside trap is to kick forward through the offside line and run for the ball. The pass evaluation algorithm is used to choose the best target for the kick through the offside line.

By using this algorithm, number of offsides per game decreased by 67%. This is mainly an advantage against teams, who intentionally create offside traps. Against teams without offside trap implementation, this has a minimal effect, as the overall number of offsides is very low.

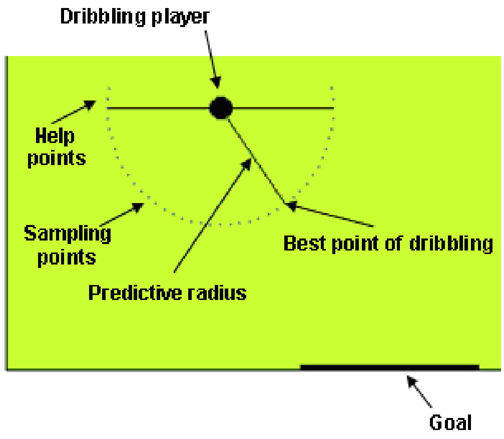

Fig. 3. Dribble target selection. Sampling points are on a semicircle around the x-axis.

3.5 Dribble

The goal of dribbling is to get the ball close to the opponent goal, while avoiding opponents. At the same time, an important effect of dribbling is to lure the opponents from other teammates. The teammates can then free themselves to receive a pass.

Potential map is used to select the best direction of dribbling. This is the application of the "classical" path planning algorithm. The player is attracted to the opponent goal if near to it. At the same time, he tries to avoid the opponents. To avoid dribbling towards a teammate, teammates are removed from the potential map for the needs of dribbling.

The whole path through the field is not constructed, as the situation on the field changes quickly and paths pre-planned for a longer time (e.g. traversing more than 1/4 of the field) have minimal chance of being accurate. Moreover, while dribbling, the player is subject to interception after a short time (typically one or few seconds). Pre-planning the path (perhaps using some graph-based or flood algorithm) is therefore waste of valuable computation time. The basic algorithm follows:

1. Select points on a semi-circle around the dribbling player in the angle $\langle -90, 90 \rangle$ (see Figure 3)
2. Evaluate the points. Value of a point is equal to the average value of cells on the connection of the point with the dribbling player.
3. Average the values. The final value of a point is equal to the average of N points ($(N-1)/2$ points to the left and to the right).
4. Choose the point with the highest value as the target of dribbling.

The averaging in step 3 is important, because this way, a triangle is searched, instead of just the path.

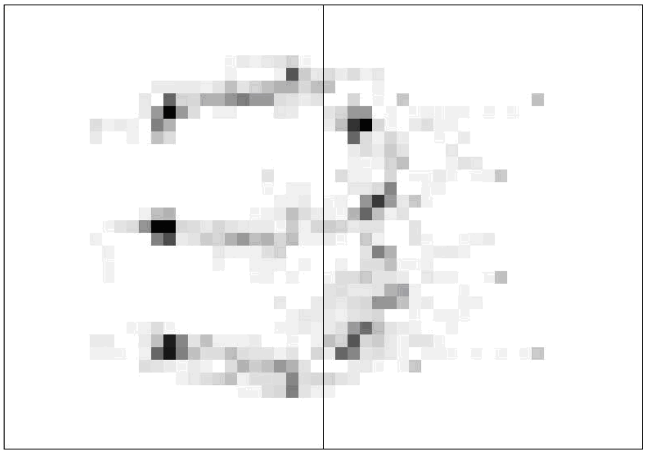

Fig. 4. Example opponent formation analysis after 300 simulation cycles. Dark fields show places, where enemy players are often positioned.

This dribbling (without other improvements and with the same higher decision logic) increased ball possession by 60%. This is mainly caused by intelligent avoiding opponents and by dribbling to "good" places with high potential. Together with the other approaches, it also increased shoot success.

3.6 Formation Adaptation

The positional game is an important factor on the path of achieving good results in soccer. The ability to predict the behaviour of the opponent can have a major role in making the correct decisions in a partially observable environment, like robotic soccer.

The ideal tool for analyzing the opponent is the on-line coach, which has a complete overview of the situation on the field. The coach analyzes the opponent by creating a map of the field with cells of size 2x2 meters. Each cycle, the coach increments cells, where an opponent is placed. After a period of time, the map contains relative amounts of time the enemy spent in each sector (Figure 4). This map is then used to compute the optimal offence and defence sub-formation. These sub-formations are then communicated to the players.

Each pre-defined sub-formation (with size of 10x34 segments) comes with a field coverage, which is computed from the positions of players in the given subformation. The sector the player stands on has coverage of 1.0, the sector, which is 10 sectors away from the home sector (the one the player stands on according to the subformation), has coverage of 0.0. The sectors in-between have a linearly decreasing coverage.

The algorithm looks for a maximum value (for defensive sub-formation) or minimal value (for offensive sub-formation) of the sum 3, where, $f(x, y)$ is the value of the sector (x, y) from the opponent map, $c(x, y)$ is the coverage of the sector (x, y) for the given sub-formation and offset is a value between 0 and 25 representing the offset in sectors of the sub-formation from the home goal line (1 sector = 2 meters, thus the sub-formation can cover any part of the defensive/offensive half of the playfield).

$$fitness = \sum_{i=0}^{10} \sum_{j=0}^{34} f(i + offset, j) * c(i, j) \tag{3}$$

The sum is computed for each sub-formation and for each offset. The result of the sum is then averaged to represent the average coverage per player in a sub-formation, so the algorithm doesn't overvalue large numbers of defensemen. The best sub-formation and its offset are then recommended by the coach.

The best observable results produced by this method are excellent coverage of enemy attackers by the defensive sub-formation, basically producing an implementation of personal defence. In the offensive part of the formation the method successfully chooses sub-formations exploiting weaknesses in the enemy defence. The best results are in situations, where the enemy defence has some parts of the field uncovered. From measurable results, ball possession was increased by 12%. More information about the formation adaptation can be found here: [13].

4 Future Work

The potential field can be further improved and used for other skills and tactics. This way, the player behaviour can be improved without the need of significantly increasing the computational complexity, as the already existing map would be used. This section describes some of the possible improvements:

Including "what-if" simulation to predict actions outcome (similar to the electric field approach [3]).

Defensive behaviours - personal defence and blocking of potential passes and shoots. For personal defence, players should move to places with low potential (with high concentration of opponents and low concentration of teammates). Blocking possible passes is similar to the pass-ball skill, covering possible safe opponent passes. Pure-potential approaches have problems with more complex situations (like the defence against several attackers is) as there may be a small local extreme between two objects influencing the potential field, causing the defender to protect the wrong place. However connecting it with a simple analytical approach, considering only connections between passing points (similar to the pass-ball behaviour) solves this easily.

Communication of key fragments of the map, allowing the player to notify others about interesting places (with very high or very low value or significant for some type of behaviour).

Improved modelling of player influence on the map, better corresponding to the simulation physics. This would slow-down the computation, but would increase the accuracy of the approach.

Worse potential on places which were not seen for a long time. If a place was not seen for a long time, its value should decrease to avoid blind actions.

5 Conclusion

We successfully used the idea of potential field in several areas of a simulated soccer game, including the "classical" path planning (dribble); action success evaluation (pass, shoot); predictive cooperation (pass + predictive pass interception); arranged cooperation and learning by run-time adaptation of formation to opponent. The applications of the potential field significantly improved the players. We have evaluated the described approaches by collecting statistical data from several hundreds simulated games against the original player (without potential maps) and against other teams. It is difficult to evaluate an approach exactly, as different strategies are successful against different players. Anyway, the presented approaches significantly improved all monitored parameters against all tested opponents. Statistical evaluation against the original player is in table 1.

Table 1. Statistical evaluation against the original player.

Parameter	Old value	New value	Change
Number of shoots to goal per game	2.5	5.9	+136%
Shoot success	16.7%	39.8%	+139%
Number of opponent shoots to goal per game	2.5	0.3	-88%
Opponent Shoot success	16.7%	2.5%	-85%
Number of own offsides per game	3.9	0.8	-79%
Number of passes per game	227.6	271.1	+19%
Pass success rate	47.3%	61.9%	+30%
Ball possession	50%	80.1%	+60%

As we can see, the sum of all approaches increased the overall game quality significantly. The synergic effect of all approaches caused, that the overall improvement is higher, than the improvement of individual approaches. The most significant improvement is in the ball possession, caused by improved dribbling and passing, together with better player global and local positioning. This also resulted in more scoring chances. This, combined with the improved shooting, resulted in an increased goal rate. There are still some problems in the player (as we can see on the relatively low pass success rate), but this is caused by improper control of the visual sensor (which will be the next to improve), rather than by the potential map approach.

The presented potential field approach has also many further possibilities of use, which would improve the quality of the game, without the need of significantly increasing the computational complexity. The use is not restricted to the simulated league, for which it was designed.

Potential fields and maps are very successful in domains of autonomous agents and multi-agent systems. This and other works in this area show, that even after many years since the first potential field approaches, their future potential is far from exhausted.

Acknowledgments. This work was partially supported by the Slovak State Programme of Research and Development "Establishing of Information Society" under contract No. 1025/04; by the Science and Technology Assistance Agency under the contract No. APVT 51-024604; by the Grant Agency of Slovak Republic grant No. VG 1/3102/06.

References

1. de Jong, S., Tuyls, K., Sprinkhuizen-Kuyper, I.: Robust and Scalable Coordination of Potential-Field Driven Agents. The International Conference on Intelligent Agents, Web Technologies and Internet Commerce (IAWTIC 2006), 29 November - 1 December, Sydney - Australia (2006)
2. Henrich D.: Fast motion planning by parallel processing – A review. Journal of Intelligent and Robotic Systems, vol 20, no 1 (1997) 45–69
3. Johannson, S.J., Saffiotti, A.: Using the Electric Field Approach in the RoboCup Domain. Birk, A., Coradeschi, S., Tadokoro, S., eds.: RoboCup 2001: Robot Soccer World Cup V. Volume 2377 of Lecture Notes in Artificial Intelligence., Springer (2002)
4. Kapustík, I.: RoboCup ako multiagentový systém. Kognice a umely zivot VII., Slezska univerzita, Opava (2007) 167–172
5. Khatib, O.: Real-time Obstacle Avoidance for Manipulators and Mobile Robots. The International Journal of Robotics Research 5 (1986) 90–98
6. Kim, J., Kim, D., Kim, Y., Seow, K.: Soccer Robotics, Springer Tracts in Advanced Robotics 11, Springer-Verlag (2004)
7. Laue, T., Schindler, T., Penquitt, F., Burchardt, A., Birbach, O., Elfers, C., Stoye, K.: B-Smart (Bremen Small Multi-Agent Robot Team) Team Description for RoboCup 2006 http://b-smart.informatik.uni-bremen.de/B-Smart_TDP2006.pdf (2006)
8. Lekavý, M., Návrat, P., Kapusta, R., Ulický, M., Košťál, M., Košík, M.: Multiagent coordination in the domain of robotic soccer. ZNALOSTI 2005, Poster proceedings. (2005) 65–68
9. Meyer, J., Adolph, Robert and Stephan, Daniel and Daniel, Andreas and Seekamp, Matthias and Weinert, Volker and Visser, Ubbo Decision-making and Tactical Behavior with Potential Fields RoboCup 2002, Lecture Notes in Artificial Intelligence, Springer (2002) 304–311.
10. Návrat, P., Kapustík, I.: RoboCup Exhibition at IIT-SRC 2006. Student Research Conference 2006: Proceedings in Informatics and Information Technologies. Vydavatelstvo STU v Bratislave (2006) 327–328
11. Nguyen, B. Q., Chuang, Y-L., Tung, D., Hsieh, C., Jin, Z.: Virtual Attractive-Repulsive Potentials for Cooperative Control of Second Order Dynamic Vehicles on the Caltech MVWT. Shi, L., Marthaler, D., Bertozzi, A. L., Murray, R. M., American Control Conference 2005, Jun. 8-10, Portland, Oregon (2005)
12. Tews, A., Wyeth, G. F.: Multi-Robot Coordination in the Robot Soccer Environment. Proceeding of the Australian Conference on Robotics and Automation (ACRA-99), Brisbane (1999) 90–95
13. Zajac, J. (Lekavý, M., supervisor): Adaptation of Formation according to Opponent Analysis. Bieliková, M., ed., Proceedings of IIT.SRC 2005: Student Research Conference in Informatics and Information Technologies, Bratislava. Faculty of Informatics and Information Technologies, Slovak University of Technology in Bratislava (2005) 22–26