Proc. of the Int. RoboCup Symposium. Seattle, WA, 2001. To appear.

Using the Electric Field Approach in the RoboCup Domain

Stefan J. Johansson[1] and Alessandro Saffiotti[2]

[1] Dept. of Software Eng. and Computer Science
Blekinge Institute of Technology
S–372 25 Ronneby, Sweden
sja@ipd.hk-r.se
[2] Center for Applied Autonomous Sensor Systems
Dept. of Technology, Örebro University
S–701 82 Örebro, Sweden
asaffio@aass.oru.se

Abstract. In autonomous robotics, so-called artificial potential fields are often used to plan and control the motion of a physical robot. In this paper, we propose to use an artificial *electric field* to address the problem or real time action selection in embodied, autonomous agents. We attach positive and negative electric charges to the relevant objects in the agent's domain, and use the resulting electric field to estimate the heuristic value of a given configuration. This value is used to select the action that results in the best configuration. This allows us to consider in the same framework both navigation and manipulation actions. We apply the electric field approach in the RoboCup domain, and present results drawn from our experience in the Sony legged robots league.

1 Introduction

There is a concept from physics that has become extremely popular in autonomous robotics: the concept of a potential field. A major initiator of this popularity has been Khatib, who in 1986 proposed an obstacle avoidance method based on the idea that obstacles exert repulsive forces on the robot, while the target exerts an attractive one [4]. At each point of the space, the robot computes a resulting force by summing all the repulsive and attractive forces, and moves as if it were subject to this force. Equivalently, we may say that the robot selects, at each point, the action that brings it to a point of locally minimal potential.

Following Khatib's suggestion, many methods for robot navigation have been proposed based on artificial potential fields (e.g., [1, 2, 5, 6]). All of these methods perform action selection in a very specific case: when the objective is to place the robot at a certain location, and the possible actions consist in robot displacements. It is not obvious if, and how, these methods can be used to solve a more general action selection problem, where the objective may involve the achievement of certain configurations of the objects in environment, and actions may include the manipulation of these objects.

We propose the *Electric Field Approach* (EFA) as a generalization of traditional potential field approaches, which allows us to control both motion and object manipulation. The two main distinctive points of our approach are: the use of the potential

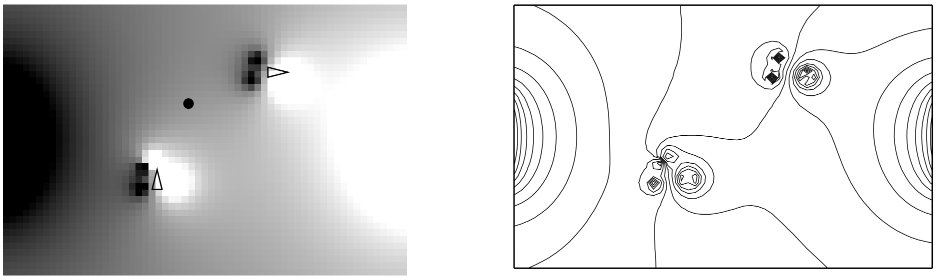

Fig. 1. Electric potential generated by a given configuration of charges (left), which give corresponding equipotential lines (right). Dark areas are negative, white areas are positive. The own net is on the left of the field. The players are indicated by triangles surrounded by charges, the ball (which is uncharged and the probe of this setup) by a circle.

field generated by a set of electric charges as a *heuristic* for action selection in complex domains; and the possibility of the robot to perform actions that *modify* the field. In the rest of this note, we outline the EFA, and show how we have used it to control a team of Sony AIBO legged robots in the RoboCup (RC) domain.

2 The Electric Field Approach

The two main concepts of the Electric Field Approach are *charges* and *probes*. The basic idea is as follows. The overall goals of the system are encoded in terms of positive and negative artificial charges placed on strategically important locations, usually attached to objects (including the robots themselves). We also chose certain probe locations, typically corresponding to the objects which we want to manipulate. *Action selection* is then formulated as the process of trying to increase the electric potential at the probed positions. This can be done by either performing actions that move the probed objects, or actions that move the charged objects. In the RC domain, for instance, the ball is probed, since our main goal it to get it at certain locations (the opponent net) and far from others (the own net). Actions include kicking (moving the probe) and navigating (moving the charges attached on the robot).

In general, charges are attached to areas that are strategically important for the probes. In RC, the opponent net is positively charged, while the own net is negatively charged. These charges are *static* (SC) in the sense that they do not move during the game. But the robots are also surrounded by charges. Positive *Environment-oriented Dynamic Charges* (EDCs) are typically placed on the side of the robot facing the opponent net, while negative EDCs are placed on the side facing the own net. Also, since the robot has better control over the ball when this is in front of it, positive *Agent-oriented Dynamic Charges* (ADCs) are placed in front of the robot. An illustrative situation is shown in Fig. 1.

The electric field encodes our knowledge about the heuristic value of the possible system configurations. Intuitively, we encode knowledge related to the structure of the system by SCs; and knowledge related to the current state by dynamic charges. The principle of superposition integrates this knowledge into an overall electric field. The heuristic value of a system configuration is then evaluated by computing the electric potential at the position of the probe.

This heuristic evaluation can be used for action selection: the goodness of each action a in a given situation s is given by the value of the situation that would result by performing a in s. Intuitively, the placement of the charges will lead the agent to prefer actions that cause the relevant objects to go "at the right places," and to avoid actions that result in a dangerous configurations. In Fig. 1, moving the left robot closer to the ball has a high heuristic value since this would increase the potential at the ball position. In general, action selection can be performed by the simple algorithm in Fig. 2.

```
procedure action_select (situation)
    highest_potential := −∞;
    possible_actions := find_applicable_actions(situation);
    foreach action in possible_actions
            simulate(action, situation);
            potential := calculate_potential(probe);
            if potential > highest_potential then
              highest_potential := potential;
              best_action := action;
    return(best_action);
```

Fig. 2. A skeleton EFA action selection algorithm for single probe environments.

3 Using the EFA in the RoboCup domain

We have applied the EFA in the "Team Sweden" entry at RoboCup-2000. The robots used were the Sony AIBO legged robots. The architecture implemented in each robot, described in [9], is a layered architecture for autonomy inspired by the Thinking Cap [8]. At the highest layer, a *reactive planner* (RP) selects an appropriate behavior to use based on the information from a *global map* GM. This behavior is executed by the middle layer *Hierarchical Behavior Module* (HBM), that sends back information to the RP when the behavior is either finished, or not possible to perform anymore.

The RP uses the EFA to dynamically decide which behavior, among those implemented in the HBM, should be used at every moment. For instance, the *GoToBall* behavior brings the robot in front of the ball at a close distance. *GoBehindBall* brings the robot close to the ball on the opposite side of the opponent net. *GoBetweenBallNet* brings the robot close to the ball on the side of its own net. *FaceBall* and *AlignWithBallNet* perform local adjustments in order to achieve an adequate posture for kicking. *Kick* and *Steal* respectively kick the ball in front of the robot using the legs, and laterally using the head.

In order to select a behavior, the RP has information about the preconditions, the expected effects, the perceptual needs and the post-conditions of each and one of the behaviors. It then uses an instance of the EFA algorithm given in Fig. 2 to decide which behavior to use. The core of the selection is the *prediction* of the effect of the behavior, and the *heuristic evaluation* of this effect by measuring the potential in the electric field at the selected probe point (the ball).

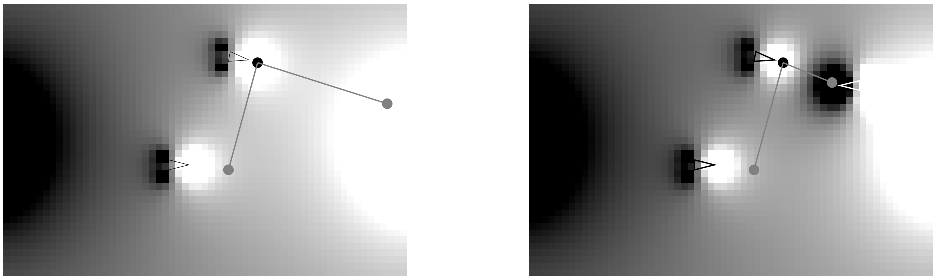

Fig. 3. Choosing between the 'Kick' and 'Steal' behaviors in two different situations. The black dots show the present position of the ball, and the lines and the grey dots, represent the calculated trajectories and predicted final positions (of the ball) respectively when performing the behaviors.

Fig. 3 (left) shows the result of evaluating two possible kicking actions in a given situation. The black circle shows the current position of the ball, the gray circles show the predicted positions of the balls after applying the *Kick* and the *Steal* behavior, respectively. The *Kick* behavior is selected since it would put the ball at a position with a higher potential. Fig. 3 (right) shows how the same example is modified when an opponent player is standing in front of the agent. The predicted position of the ball after executing a *Kick* lays now right in front of the opponent, in an area with a negative potential. Therefore, the *Steal* behavior is selected in this situation, thus passing the ball to the team mate. This example shows how the implemented action selection strategy can achieve some degree of (implicit) coordination.

An important issue is how often the RP should be called. Since the RoboCup domain is highly uncertain and dynamic, the RP cannot simply select a behavior and stick to it until that behavior has completed, since external factors may make this behavior inadequate. On the other hand, continuously re-evaluating the situation may lead to an unstable system, e.g., oscillating between two behaviors that have similar heuristic values. In our domain, we have adopted an extension of the approach proposed in [7]. We select a behavior, and commit to it until either:

– the behavior is close to completion;
– the behavior cannot be safely executed; or
– the situation has significantly changed since the time the behavior was selected.

When any of these conditions is true, the RP is called again and a new behavior is selected. Examples of the last condition are the fact that the ball has moved by a significant amount, or that the quality of the robot's self localization has significantly improved.

Fig. 4 shows a sequence of decisions produced by the EFA in a simplified (one agent) situation. (Similar sequences occurred in the actual competition, but the complexity of the situations makes them difficult to record and report.) At start, the applicable behaviors are evaluated: since the ball is far from the agent, only the navigation behaviors

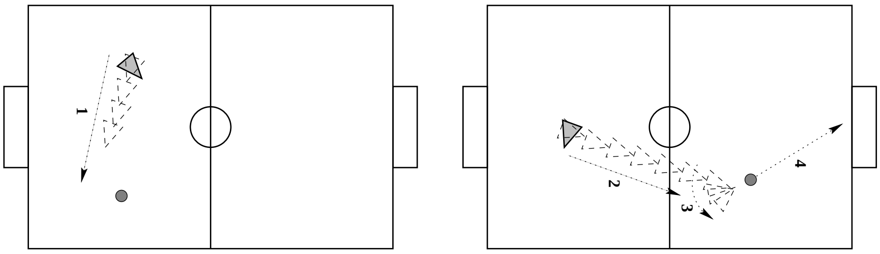

Fig. 4. A simple demonstration of the EFA: (1) *GoBetweenBallNet*, (2) *GoBehindBall*, (3) *AlignWithBallNet*, and (4) *Kick*.

are applicable. Of these, the *GoBetweenBallNet* results in the highest potential at the ball position (1 in the figure): intuitively, this behavior brings the agent to a defensive position where its positive charges best "shield" the negative potential produced by our own net. While executing the *GoBetweenBallNet*, the ball moves to the other side of the field. This change in the situation triggers a re-evaluation of the behaviors by part of the RP, and *GoBehindBall* is now the preferred one (2). When this behavior is close to completion, the RP is called again. The agent is now close to the ball, so the *FaceBall*, *AlignWithBallNet*, *Kick*, and *Steal* behaviors are applicable. Of these, *Kick* and *Steal* would move the ball to places with a lower potential than the current one, while *AlignWithBallNet* would slightly increase the potential on the ball. The robot so performs the *AlignWithBallNet* (3), and then re-evaluates the behaviors. The *Kick* behavior is now the preferred one (4), since it will bring the ball into the opponent net.

The EFA leads itself to efficient implementations, since the field only needs to be evaluated at a few positions and the contribution of all the static charges can be precomputed. The Team Sweden implementation of the EFA, detailed in [3], runs in real time using only a small fraction of the onboard computational resources in the AIBO.

4 Discussion

Our experience has shown that the electric field approach is an efficient tool for higher-level robot control which is general, robust, and maintainable: *General*, in the sense that classical potential field approaches are (at least on a conceptual level) special cases of the EFA. These can be encoded in the EFA by attaching negative charges to the obstacles, a positive charge to the goal, and using the robot itself as a probe. The only actions available in this case are one-step motions of the robot. *Robust*, with respect to errors in the calibration of the charges or in the estimated state of the environment: if these errors are small, the selected action will be close to the optimal one with respect to its effects. *Maintainable*, since we have found that the EFA can be easily adapted to perform new tasks in the same domain by modifying a few elements, and without changing the basic algorithm. Moreover, it is easy to include more objects or probes in the environment. This is due to the additive nature of the field, and the fact that all objects are modeled in similar ways. For example, one part of the RoboCup-200 competition was a series of "technical challenges." The program thaat we used for these challenges was the same one used for normal games, with only a few modifications to the charges and preconditions used in the RP.

There are open questions in the current state of development of the EFA. For example, we have only explored the use of EFA for one-step look-ahead decision making. This may obviously lead to myopic behavior, and it would be worth to study the use of EFA to generate short term plans. Another issue is that we currently account for information gathering actions in a *ad hoc* way. We believe that there are ways to include a uniform treatment of perceptual actions in the EFA framework. A third problem is the current lack of coordination between several robots running an EFA-based action selection mechanism as described here: the robots may decide to head toward the ball simultaneously. We are studying the inclusion of explicit coordination in the EFA through communication. We also intend to test the EFA approach on domains besides the RoboCup.

Acknowledgements

The work reported in this note was partially supported by the Swedish KK Foundation. John Johansson implemented the EFA on the AIBO. The authors are grateful to the members of Team Sweden for worthy insights and collaboration and to the Center for Applied Autonomous Sensor Systems and the Societies of Computations research groups for their support.

References

1. R. C. Arkin. Motor schema based navigation for a mobile robot. In *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 264–271, 1987.
2. J. Borenstein and Y. Koren. The vector field histogram: fast obstacle avoidance for mobile robots. *IEEE J. of Robotics and Automation*, 7(3):278–288, 1991.
3. J. Johansson. An electric field approach — a strategy for sony four-legged robot. M.Sc. Thesis MSE-2001-03, Blekinge Institute of Technology, Sweden, April 2001.
4. O. Khatib. Real-time obstacle avoidance for manipulators and mobile robots. *The International Journal of Robotics Research*, 5(1):90–98, 1986.
5. B.H. Krogh and C.E. Thorpe. Integrated path planning and dynamic steering control for autonomous vehicles. In *Proc. of the Intl. Conf. on Robotics and Automation*, 1986.
6. J. C. Latombe. *Robot Motion Planning*. Kluver Academic Publishers, Boston, MA, 1991.
7. S. Parsons, O. Pettersson, A. Saffiotti, and M. Wooldridge. Robots with the best of intentions. In M. Wooldridge and M. Veloso, editors, *Artificial Intelligence Today*, pages 329–338. Springer-Verlag, 1999.
8. A. Saffiotti. The Thinking Cap homepage. http://www.aass.oru.se/˜asaffio/Software/TC/.
9. Team Sweden. Official website. http://www.aass.oru.se/Agora/RoboCup/.