LEVELS OF MULTI-ROBOT COORDINATION FOR DYNAMIC ENVIRONMENTS

Colin P. McMillen, Paul E. Rybski, Manuela M. Veloso
School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, U.S.A.
mcmillen@cs.cmu.edu, prybski@cs.cmu.edu, veloso@cs.cmu.edu

Abstract   RoboCup, the international robot soccer competition, poses a set of extremely difficult challenges for multi-robot systems. To be competitive in RoboCup's rapidly-changing, dynamic, adversarial environment, teams need to make use of effective coordination strategies. We describe some of our experiences with effective coordination of robots teams and introduce several levels of strategies which encapsulate coordination from the level of individual robots to synchronized coordination of the entire team.

Keywords: Adversarial environments, robot soccer, multi-robot coordination

1.   Introduction

The RoboCup robot soccer competition is a domain in which teams must address the challenges of real-time perception, cognition, and action. Robots must be able to operate in a very dynamic environment in which they must reason not only about the actions of their teammates, but also about the actions of a team of adversarial agents. Teams of robots that operate without an effective teamwork strategy are likely to hamper each other's efforts and perform as an inferior team. Our primary research interests are in exploring the scientific challenges of developing effective teamwork strategies for autonomous robotic systems where all sensing and cognition is done on-board the robot.

We have had extensive experience with robot positioning ranging from early work in simulated robot soccer players (Veloso et al., 1999) to the Sony AIBO league (Uther et al., 2002). In this previous work, we made use of artificial potential field methods and have found them to be a

1

very powerful way of representing multiple constraints when positioning robots. However, there are a number of limitations in the kinds of behaviors that potential fields can express. We are actively exploring other coordination strategies which we will describe in more detail.

This paper focuses on developing team strategies for robots that compete in the RoboCup (Kitano et al., 1997) legged league. We discuss various approaches for teamwork and cooperation, and describe some empirical results from experiments.

1.1 Related Work

Potential field methods have been used very successfully for navigation tasks such as obstacle avoidance (Khatib, 1985). This idea has been extended such that a group of robots can maintain formations while using only local information in their potential calculations (Balch and Arkin, 1998; Balch and Hybinette, 2000).

Several groups have encoded domain-specific heuristics into potential fields. In the RoboCup domain, potential fields can be constructed that guide robots to an area near the opponent's goal or to an open position that is well-suited for pass reception (Castelpietra et al., 2001; Veloso et al., 1999; Weigel et al., 2001). A behavior architecture that relies on potential fields for motion planning and action selection is described in (Laue and Röfer, 2004). Their approach has been applied to real robots, including the German Team of the RoboCup legged league.

Potential field methods have several limitations that have been reported in the literature, including susceptibility to local minima and oscillations (Koren and Borenstein, 1991). An approach known as forward chaining dynamically reshapes the potential field using heuristics that guide the robot to the goal through a series of subgoals or waypoints that attempt to avoid local minima (Bell and Weir, 2004). The forward chaining approach is particularly interesting to us because it utilizes the idea of a positioning function that changes over time.

1.2 Coordination in Robot Soccer

We have been researching coordination strategies for the different RoboCup robot soccer leagues for several years. RoboCup robot soccer in general (Kitano et al., 1997) offers a very challenging research environment for multi-robot systems. One of the most interesting aspects of the RoboCup leagues is that we change the rules of the game and the playing environment every year in order to increase the difficulty of the problem towards matching real setups as much as possible.

Of particular interest to this paper is our work in the RoboCup legged league with the Sony AIBO four-legged robots, which need to be fully autonomous with on-board sensing, computation, and action. The legged league in particular has gone through several changes since 1998. Some of the most significant changes occurred in 2002 and have made the most impact in our multi-robot research efforts.

Communicating robots: The AIBO robots are now equipped with the ability to communicate wirelessly. In the initial years, when the robots could not communicate, we achieved teamwork through vision of the position of the ball – the robots' behaviors were vision-servoed. Three attacker robots searched for the ball; as soon as they saw the ball, they would move towards it and then move the ball towards the opponent goal. Because the ball was often not in the robots' field of view, due to the occlusion by other robots and the robots' own search for localization markers, not all the robots could see the ball. Teamwork was a property that emerged due to this discontinuity of ball perception: when one robot saw the ball, it would move toward the ball, tending to block the ball from the view of its teammates. Because of this occlusion, the teammates would remain spread out from the "attacker." Now that communication is available between robots, we have researched methods of sharing information (Roth et al., 2003) and dynamically changing robot roles and positioning (Vail and Veloso, 2003). In this paper, we discuss the different levels of dynamic coordination necessary in the presence of skilled opponent teams. We present the solutions that we have developed and plan to continue researching.

World space increase: The field's size has increased by approximately 50% from its initial size, and the number of robots in a team has increased from three to four. The increase in field size makes it infeasible for robots to see across the entire field. In the initial smaller field, individual robots could recognize objects across the complete field. Modeling the world state is now a task that needs to combine a robot's own perception and the communicated information from other robots. Multiple robots need to build an accurate world model and select joint actions that fit a team policy.

Rules of the game: The rules of the game set constraints on the legal positioning and actions of robots. For example, only one robot is allowed to defend the goal area. This type of rule creates hard constraints on the dynamic positioning of team members. In addition, robots encounter difficult motion situations when surrounded by

opponent robots. In these situations, a team member may need the help of robot teammates. This is a challenge that requires an effective dynamic coordination algorithm that monitors the progress of teammates. In addition, teamwork should change as a function of the opponent team, the specific state of the field, and the remaining time of the game.

The RoboCup legged league continues to motivate our research in multi-robot domains, inspiring incremental algorithmic successes and providing many issues to be addressed. Interestingly, the more we work on this adversarial multi-robot coordination problem, the more we understand how the problems we face go well beyond robot soccer and are of relevance to multi-robot systems in complex environments. In this paper, we present our findings aiming at such an abstract level.

2.   Dynamic Multi-Robot Coordination

Over the past few years, teams have experimented with different methods of team coordination. Many of these strategies involve keeping teammates from running into each other and placing teammates in good locations on the field so that they can be in good positions to receive passes or go after a free ball. While there have been some good approaches, no one strategy has emerged as being clearly superior to all others. One reason for this is that several different coordination strategies are likely to be applicable in a single situation. Since some strategies may work better than others, a team that selects the superior strategy will be at an advantage. Thus, one of the most important problems to address when designing a multi-robot soccer team is selecting the kind of coordination strategy that will be used during the game. Teams may choose to use a fixed coordination strategy defined *a priori*, but if chosen poorly, a fixed strategy may not fare well against the strategy of the other team. Thus, an important extension to the research problem of coordination strategies is the ability for a team to dynamically change their strategy at runtime to adapt to their opponents' strengths and weaknesses.

Dynamically selecting a different strategy depending on the situation can be very powerful technique, but can be very challenging to implement well. Robots that use a dynamic coordination system must be able to perceive and properly evaluate the state of the world as well as the state of their own progress. This information is vital when making the decision to switch from a poorly performing strategy to one that could potentially work better.

We have identified several different levels for dynamic coordination that can be applied to a robotic team. These include:

- A "first-order" approach, where the robots use a fixed coordination strategy and each robot modifies the parameters of its behavior according to the world state.

- A "second-order" approach, where the robots have multiple ways of handling different situations. In order to utilize a second-order strategy, the robots must be able to evaluate the world state so that they can choose between the different behaviors they have at their disposal.

- A "third-order" approach, where the robots have several different team strategies, or "plays," which describe the coordinated actions of all of the robots together. Depending on the world state, different plays may apply; the team collectively decides upon the right behavior to apply in a given situation.

We have implemented methods for first- and second-order coordination strategies, a description of which is provided below. Currently, the third level of coordination has been implemented in our small-sized league (Bowling et al., 2004) but not yet on the AIBOs.

2.1   Changing Single Robot Parameters

We define the first-order coordination strategy as the ability for the robots to set their own behavior based on the state of the world. In this kind of system, each robot is programmed with a single behavior set which is used to control the robot's behavior in its environment.

We have tried two different methods for representing first-order coordination strategies. The first is a potential fields approach and the other is an approach that we call constraint-based positioning. In previous work (Vail and Veloso, 2003), we give a detailed description of our implementation of potential field-based coordination. In this approach, we use potential fields both to determine the role that each robot plays (attacker, supporting attacker, and defender) and also to determine where the supporting robots position themselves on the field of play. On efficiency issue with potential fields occurs when they are used to coordinate the actions of a team of robots in a very dynamic world. In this situation, the fields may need to be recomputed for each every new sensor reading. This does not tend to be true for implementations of potential fields that are used for navigation in more static environments. In general, however, it's possible for minor disturbances in the positions or strengths of individual attraction and repulsion fields to cause fairly significant changes in the local gradient surrounding the robot.

Constraint-based positioning is an approach to robot positioning that we have developed in the last year for the 2004 RoboCup competition. Under this regime, robots are still assigned roles using a potential function, but the field positions chosen by the supporting robots are subject to a set of constraints. This approach was developed because there are several hard constraints that we would like to enforce on the robots' positions which are difficult to specify clearly with potential fields. For instance, defender robots need to avoid their own goalie's defense box, because entering the defense box is a violation which will cause the robot to be removed from play for 30 seconds. Other constraints that we would like to enforce include not crossing in front of a robot that is about to take a shot on goal, not coming within a certain minimum distance of a teammate, and so on. Consider a situation in which a robot is near the defense zone and a teammate is directly approaching it. Should the robot move toward the goal, violating the defense-zone constraint, or stand still, violating the teammate-distance constraint? Our implementation of constraint-based positioning allows us to prioritize the constraints, so that the robot knows that entering the defense zone is a more serious violation than coming near a teammate. In theory, the priorities of these constraints could be represented as a potential field, but we have found that debugging the complex potential fields that result can be difficult. If no constraints are in danger of being violated, the robot can choose to move to a specific point that is chosen based on the current state of the world. In this case, the robot can still use potential fields to choose an open area on the field or to choose a path to navigate around local obstacles.

Our experience with RoboCup has been that a single positioning function defined for a particular role tends to be too limiting. Trying to capture all of the possible actions that a robot might accomplish can cause the complexity of the positioning function to grow beyond what is manageable. A soccer-playing robot might have multiple ways of approaching the goal, each of which has advantages depending on the relative position of the goalie and/or his other players. In some situations, the robot may want to try one approach and if it fails, try a different approach. Behaviors like these may be mutually exclusive and as such could be very difficult for a single function to capture.

2.2 Changing Single Robot Behaviors

An alternative is to factor the problem into subproblems and make multiple positioning functions available for the robot to use. In this case, a second-order decision process must exist whose purpose is to

evaluate the state of the world and/or the current performance of the robot. This decision process is responsible for deciding which positioning function should be used in a particular situation.

Designing multiple behaviors such as these with potential fields requires that an entirely new set of potential attractor/repulsor nodes be defined for each of the new behaviors. A single set of nodes cannot be used for independent behaviors because the individual nodes are not independent of each other. They all affect one another.

Another challenge with potential fields is that in the case of multiple specific and possibly exclusive behavior sets, a robot may be expected to approach a very specific location on the field and stay there. Specifying a specific (x, y, θ) location on the field would be fairly straightforward for a constraint-based system to handle, but designing the potentials such that they push the robot to a specific location on the field can be a very challenging task. An extreme solution for the potential fields approach is to have a single potential attractor that pulls the robot to the specified point. This suggests that having control over the attraction/repulsion nodes and being able to turn them on and off as necessary would make the potential field approach work in this situation.

In a constraint-based system, the decision process evaluates the points on the field and chooses a specific location for the robot to reach. In both positioning methodologies, a higher-level decision process is in charge of selecting the specifics of the behavior set by evaluating the state of the environment and selecting the one with the highest chance of success.

3.   Experimental Results

We have performed a set of experiments that show the need for second-order reasoning in the RoboCup domain. These experiments demonstrate that we can improve performance by having a higher-level decision process that changes the positioning strategy based on the environment. Specifically, we compare the performance of two positioning strategies under differing environmental conditions, and show that the strategy which is superior in one situation is inferior in the other situation.

In each experimental trial, we placed the ball in one of the front corners of the field, and two robots (on the same team) attempted to score a goal within thirty seconds. We chose this initial position of the ball because it has traditionally been difficult to score a goal from the front corner of the field. In this situation, it is not usually possible to score a goal by a single direct kick; trying to do so will often send the ball rolling into the opposite corner. From the other corner, the attacker may very well choose to execute another strong kick toward the goal,
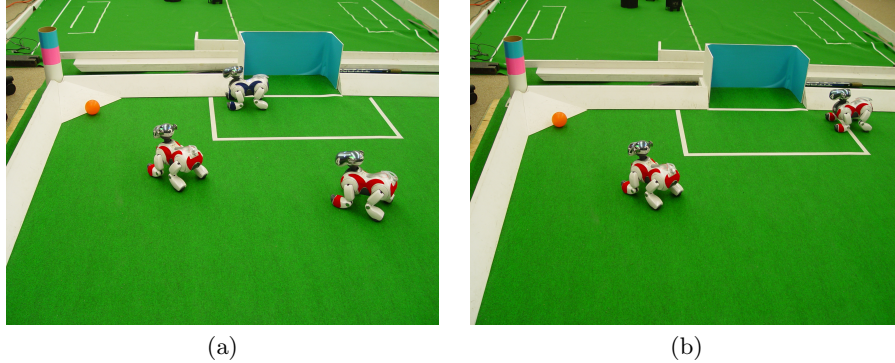
Figure 1. Two of the four initial configurations used in the experimental trials. Image (a) shows the supporter in the center position with a stationary goalie present on the field. Image (b) shows the supporter in the side position with no goalie.

which can lead to a series of "ping-pong" kicks across the goal until the goalkeeper clears the ball or until noise in the environment causes the ball to roll into a different area of the field. The 30-second time limit only gives the robots enough time to execute approximately three to five kicks, so we feel that a goal scored within that time limit indicates that the robots were performing reasonably well during that trial.

In half of the trials, we placed a goalie robot in the defense zone, facing the corner where the ball was initially placed. The position chosen was the one that our own goalie would adopt if the ball were placed in that position. However, the goalie was paused, and therefore did not attempt to clear the ball or attempt to move from this initial position unless it was pushed by the other robots. In the other half of the trials, no goalie was placed on the field.

One of the two robots on the team (the attacker) was placed 75 cm away from the ball, facing the corner of the field. The supporting robot was positioned according to one of two different potential fields. Both fields simply contained a single linear attractor that pulled the supporter to a desired point. In the *side* potential field, the supporter was drawn toward a point on the opposite corner of the goal; in the *center* potential field, the supporter was drawn toward a center point about 100 cm from the front of the goal. See Figure 1 for pictures of the initial configurations of the field, including the supporter positioning induced by the two different potential fields.

We ran 40 trials for all four different possible setups (with or without goalie, combined with center or side positioning), for a total of 160 trials. For each trial, the success or failure of the run was recorded. If the run

was a success (i.e., it terminated in a goal), we also recorded the amount of time it took for the robots to score the goal.

Each run started by unpausing the attacker robot; the 30-second timer was started as soon as the attacker touched the ball. If any robot crashed or ran out of batteries during a trial, the robot was rebooted and the trial was restarted from the beginning. Normal RoboCup penalties, such as player pushing, goalie pushing, and holding, were not enforced. If the ball was knocked out of the field, it was immediately placed back in-bounds at the place where it went out, as per the RoboCup 2004 rules.

The results of these experimental runs are summarized in Figure 2. Figure 3 shows the individual completion times for every trial. Note that the results are only shown for the runs that were counted as successes; therefore, each graph has a different number of points plotted.

	Successes	Failures	Mean Time per Success
Side positioning, no goalie	31	9	9.97s
Center positioning, no goalie	27	13	16.91s
Side positioning, with goalie	12	28	23.63s
Center positioning, with goalie	17	23	18.55s

Figure 2.   Summary of the results obtained in the experimental trials.

In the no-goalie case, the side positioning succeeded slightly more often than the center positioning, and the mean time per success was significantly lower for the side positioning. (Statistical significance of the mean time determined by Student's two-tailed t-test, with $p = 0.001$.) However, in the runs with the goalie, the center positioning significantly outperformed the side positioning, with a higher number of successes and a faster mean time per success. (Statistical significance of the mean time determined by Student's two-tailed t-test, with $p = 0.047$.)

The advantages and disadvantages of each position are easily explained through a qualitative analysis. The side position does much better in the no-goalie case because the position of the supporter puts it in a very good location to intercept the attacker's kick. After a successful interception, a single head kick is usually sufficient to score a quick goal. The center positioning does not enable the easy interception of a missed shot, so it is more likely that the ball will end up in the opposite corner and require more time before a goal can be scored.

However, when a goalie is added to the field, the weaknesses of the side positioning become apparent. The initial kick often bounces off the goalie and stops close to the center of the field, instead of traveling across the field to the other side. In this situation, the supporter positioned in the center is much more likely to be able to assist the at-
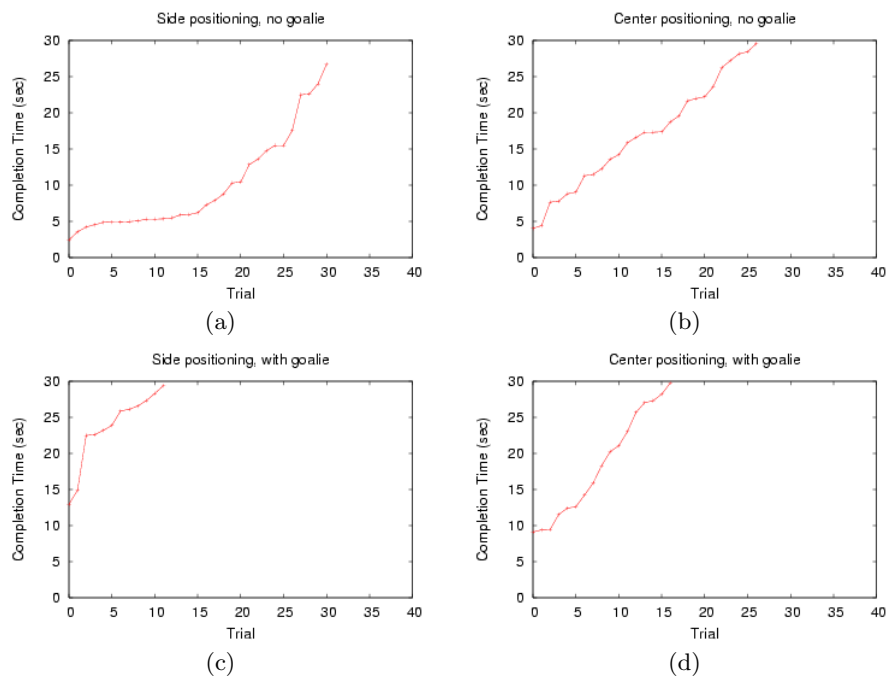
Figure 3. Graphs showing the amount of time it took to successfully score a goal. Each trial was stopped after 30 seconds if a goal had not yet been scored. Graphs (a) and (b) show the results for the no-goalie case; graphs (c) and (d) show the results for the with-goalie case. Trials are sorted from fastest to slowest completion time.

tacker. Furthermore, it is difficult for the side-positioned supporter to react quickly to changes in the ball's location, since the supporter's view of the ball is often occluded by the goalie. The center positioning is a more general approach that allows the supporter to chase down the ball relatively quickly wherever it goes on the field, while the side positioning is superior in the special case where the opposing goalie is temporarily outside the defense box.

Though the center positioning is the approach that we would prefer the majority of the time, there is a definite benefit to being able to use side positioning to exploit the situation when the goalie is not guarding the goal. For example, one of the only two goals scored in the (very defensive) final game of the 2004 US Open occurred when the opposing goalie temporarily left the defense zone and was inadvertently blocked from returning to the goal by another robot that had gotten behind it. The results presented in this section suggest that there is definitely a

benefit to be gained from using second-order reasoning in multi-robot systems, especially in an adversarial, dynamic environment.

4.   Conclusion / Future Work

In this paper, we have proposed a classification scheme that identifies various levels of dynamic multi-robot coordination. We have provided examples showing the limitations of first-order coordination strategies in the robot soccer domain, and presented experimental results that show that there is a substantial benefit to our use of second-order reasoning about team coordination.

In the future, we intend to improve upon our existing coordination strategies by adding third-order functionality to our team. We plan to take inspiration from the idea of using a playbook for team coordination, which has been a successful strategy in the RoboCup small-size league (Bowling et al., 2004). The effectiveness of playbooks in the small-size league is largely due to the fact that this league makes use of an overhead camera and so the state of the entire team can be very easily determined. The legged league has no such overhead camera system and so a team state estimate must be computed in a distributed fashion by merging the local sensory information from each of the robots. We are actively researching methods for accomplishing this task so that we can pursue the development of third-order coordination strategies, such as a playbook, for our RoboCup legged league team.

References

Balch, T. and Hybinette, M. (2000). Social potentials for scalable multirobot formations. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 1, pages 73–80.

Balch, T. R. and Arkin, R. C. (1998). Behavior-based formation control for multiagent robot teams. *IEEE Transactions on Robotics and Automation*, 14(6):926–939.

Bell, G. and Weir, M. (2004). Forward chaining for robot and agent navigation using potential fields. In *Proceedings of the 27th conference on Australian computer science*, volume 26, pages 265–274, Dunedin, New Zealand.

Bowling, M., Browning, B., Chang, A., and Veloso, M. (2004). Plays as team plans for coordination and adaptation. In Polani, D., Browning, B., Bonarini, A., and Yoshida, K., editors, *RoboCup 2003: Robot Soccer World Cup VII*, volume 3020 of *Lecture Notes in Computer Science*, pages 686–693. Springer Verlag, Berlin, Germany.

Castelpietra, C., Iocchi, L., Nardi, D., Piaggio, M., Scalzo, A., and Sgorbissa, A. (2001). Communication and coordination among heterogeneous mid-size players: ART99. *Lecture Notes in Computer Science*, 2019:86–95.

Khatib, O. (1985). Real-time obstacle avoidance for manipulators and mobile robots. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 500–505.

Kitano, H., Asada, M., Kuniyoshi, Y., Noda, I., and Osawa, E. (1997). RoboCup: The robot world cup initiative. In Johnson, W. L. and Hayes-Roth, B., editors, *Proceedings of the First International Conference on Autonomous Agents (Agents'97)*, pages 340–347, New York. ACM Press.

Koren, Y. and Borenstein, J. (1991). Potential field methods and their inherent limitations for mobile robot navigation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1398–1404, Sacramento, CA.

Laue, T. and Röfer, T. (2004). A behavior architecture for autonomous mobile robots based on potential fields. In *8th International. Workshop on RoboCup 2004 (Robot World Cup Soccer Games and Conferences), Lecture Notes in Artificial Intelligence*, Lecture Notes in Computer Science, Berlin, Germany. Springer Verlag.

Roth, M., Vail, D., and Veloso, M. (2003). A real-time world model for multi-robot teams with high-latency communication. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 3, pages 2494–2499.

Uther, W., Lenser, S., Bruce, J., Hock, M., and Veloso, M. (2002). CM-Pack'01: Fast legged robot walking, robust localization, and team behaviors. In Birk, A., Coradeschi, S., and Tadokoro, S., editors, *RoboCup 2001: Robot Soccer World Cup V*, Lecture Notes in Computer Science, pages 693–696. Springer Verlag, Berlin, Germany.

Vail, D. and Veloso, M. (2003). Dynamic multi-robot coordination. In *Multi-Robot Systems: From Swarms to Intelligent Automata, Volume II*, pages 87–100. Kluwer Academic Publishers.

Veloso, M., Stone, P., and Bowling, M. (1999). Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer. In *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II*, volume 3839, Boston.

Weigel, T., Auerbach, W., Dietl, M., Dümler, B., Gutmann, J.-S., Marko, K., Müller, K., Nebel, B., Szerbakowski, B., and Thiel, M. (2001). CS Freiburg: Doing the right thing in a group. *Lecture Notes in Computer Science*, 2019:52–63.