

Coordinated Team Play in the Four-Legged RoboCup League

Georgios Kontes and Michail G. Lagoudakis
Intelligent Systems Laboratory
Department of Electronic and Computer Engineering
Technical University of Crete
Chania, Crete, 73100, Hellas (Greece)
gkontes@isc.tuc.gr, lagoudakis@intelligence.tuc.gr

Abstract

For a decade now, the RoboCup competition promotes research in robotics through soccer games between autonomous robot teams. The ability to coordinate the players within such a team of robots is the key to the success of the team. Team coordination in a human soccer game is achieved through various team formations, tactics, and strategies. Unfortunately, research in the four-legged RoboCup league has focused mostly on single player skills, demonstrating only limited results in coordinated team play. In our work, we adapt and transfer formations, tactics, and strategies used by human soccer teams, such as the popular 4-4-2 scheme, to our four-legged RoboCup team Kouretes. We define roles for each player in all the cases we consider and we implement these roles using Petri Net Plans (PNP). The assignment of appropriate roles to players is performed dynamically during the game depending on the current game state using a simple communication scheme and a finite state machine. Our approach is implemented and tested on our four-legged RoboCup team. The proposed coordination scheme can be generalized and used in various robot team applications beyond robotic soccer, such as planetary exploration and search-and-rescue missions.

1 Introduction

In its short history, the RoboCup competition [4] has grown to a well-established annual event bringing together the best robotics researchers from all over the world. Its uniqueness stems from the real-world challenge it poses, whereby the core problems of robotics (perception, cognition, action, coordination) must be addressed simultaneously under real-time constraints. The proposed solutions are tested on a common benchmark environment through soccer games in various leagues, thus setting the stage for



Figure 1. One of our Sony Aibo ERS-7 robots.

demonstrating and promoting the best research approaches, and ultimately advancing the state-of-the-art in the area.

The Four-Legged League of the RoboCup competition is among the most popular leagues, featuring four robot players (SONY AIBO robots - Figure 1) on each team competing in a 4×6 meters field. The league is characterized by a unique combination of features: autonomous vision-based player operation, legged locomotion and action, uniform robotic platform. Given that the underlying robotic hardware is common for all competing teams, research efforts have focused on developing more efficient algorithms and techniques for visual perception, active localization, omnidirectional motion, skill learning, and coordination strategies. During the course of the years an independent observer could easily notice a clear progress in all research directions related to single robot skills and abilities, however, one can also notice that little progress has been made at the robot team level.

The ability to coordinate within a team is a crucial factor to the success of the team. In years 2006 and 2007 the

passing challenge in the four-legged league was specifically designed to promote research efforts in coordinated team play. According to the rules of the challenge, three robots are placed on the field in a triangular formation; they must pass the ball around without moving away from their initial position and they score points for accurate passing and successful grabbing. The ability of successful passing, even though it could be seen as a single robot skill, underlies any form of team coordination in robot soccer. Given an accurate passing mechanism, team coordination in a soccer game extends to various team formations and strategies depending on the actual state of the game. Players can choose between different ways of positioning themselves on the field in order to move the ball faster, trick the opponents, and eventually score goals.

Team formations, tactics, and strategies is largely an unexplored area in the four-legged RoboCup research. In our endeavors as a RoboCup team, we decided to take a radical step in behavior control and implement robot soccer strategies, which are inspired by human soccer strategies used in real soccer games. Taking under consideration that the ultimate goal of RoboCup is a game between robots and professional human soccer players, we believe that our current and ongoing work takes a critical step towards this goal.

The work presented in this paper is by no means limited to robotic soccer. In fact, the need to coordinate multiple robots arises in any application involving multiple robots. Such applications include planetary exploration, search-and-rescue missions, collaborative assembly, and area surveillance. The proposed scheme, which is based on Petri-Net Plans (PNPs) and Finite State Machines (FSM), can be adapted to various robot domains, optionally in conjunction with other coordination methods, such as market-based methods [2] and auction-based coordination [5].

2 From Human to Robot Play

Each team in a human soccer game consists of eleven players; one of them assumes the role of the goal keeper, whereas the others can freely move around the field. One of the oldest and most popular team formations is the so-called 4-4-2 formation (also described as 4-1-2-1-2 or 4-3-1-2) which is shown in Figure 2 (a). This is the favorite formation of the Argentina national soccer team; also, FC Porto under coach Jose Mourinho won the Champions League in 2004 playing mostly with this formation [1].

In a four-legged robot soccer game, each team consists of only four players (robots), one of them being the goal keeper. In order to adapt the 4-4-2 system, we assume that the RoboCup field corresponds to the center of an actual soccer field and we focus on the middle diamond formation. Since we lack a player in the middle we organize the players in a triangular formation. The actual positions are different

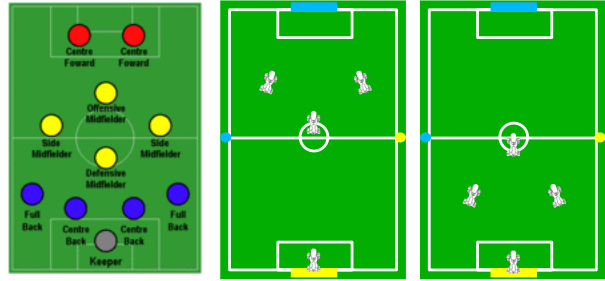


Figure 2. (a) The 4-4-2 soccer formation [1]. (b) Our robot team formation for offense. (c) Our robot team formation for defense.

for offense and defense as shown in Figure 2 (b), (c). In our work we consider four different tactics for applying the aforementioned formation: two for offense (Counter Attack and Passing Attack) and two for defense (Pressing Defense and Passive Defense). The decision to switch between tactics is made dynamically during the game.

In human soccer, when a defending team wins the ball, it strives to pass the ball from the Defenders to the Midfielders and from the Midfielders to the Attackers. Meanwhile, the Attackers and the Midfielders try to spread along the wings of the field, in order to break the opponent's defense and score.

In our RoboCup team, the switch from defense to offense is implemented with two different tactics. In the Counter Attack tactic, the robot which wins the ball attacks towards the opponent's goal supported by the robot further ahead, while the remaining robot stays behind and guards. Alternatively, according to the Passing Attack tactic, the robot with the ball tries to pass the ball to the robot further ahead in which case they switch to Counter Attack. In either case, a lost ball signals a switch from offence to defense.

In human soccer, when the attacking team loses the ball, all the teammates, except the attackers, try to get themselves behind the ball. Everyone, including the attackers, is pressing for a mistake in the opponent team. However, when the team is strongly defensive, even the attackers move behind the ball line.

In our RoboCup team, we implement the switch from offense to defense with two different tactics depending on the current state of the game. If a robot loses the ball, the tactic switches to Pressing Defense, except when the ball makes it into the opponent's half of the field and there is no teammate anywhere near the ball, in which case the tactic switches to Passive Defense. In both cases, a hold on the ball signals a tactic switch from defense to offense.

3 Roles in Tactics

In this section we provide a brief description of the role of each player in each tactic. Each player has a role in the field which can change dynamically during the game depending on the current player and ball locations. In our work, we consider three roles, the Attacker (ATT), the Defender (DEF), and the Midfielder (MID), since the Goal Keeper role does not change over time. Each role corresponds to a different behavior depending on the tactic followed at the time as described below.

Counter Attack

ATT The Attacker dribbles with the ball towards the opponent's goal straight from the place where he won the ball. When he reaches the opponent's goal area, he either shoots directly to the goal or passes the ball to the Midfielder who awaits on the opposite side of the field at the corner of the penalty area. If such a pass occurs, the Attacker notifies the Midfielder to catch the ball.

MID The Midfielder supports the Attacker, positions himself at the opposite side (at the corner of the penalty area), and waits for a pass from the Attacker. If he receives a successful pass, he shoots to the goal.

DEF The Defender stays at the center of the field, just past the center circle. In case the ball bounces back to him, he shoots towards the goal.

Passing Attack

MID The Midfielder dribbles towards the opponent's goal from the place he won the ball, looking for opportunities to pass the ball to the Attacker. Once a pass is made, he notifies the Attacker to catch the ball.

ATT The Attacker slowly moves towards the opponent's penalty area constantly facing the Midfielder. While moving or after reaching the opponent's penalty area, he looks for a pass from the Midfielder. As soon as a pass is successfully received, the tactic switches to Counter Attack.

DEF Same as in the Counter Attack above.

Pressing Defense

ATT While the ball is on the opponent's half of the field, the Attacker is chasing after the ball. When the ball enters his own half of the field, the Attacker moves around the middle line, continuously tracking the ball and placing himself between the ball and the opponent's goal.

MID While the ball is on the opponent's half of the field, the Midfielder moves in front of the middle line and

places himself between the ball and the own goal. Only when the ball enters his own half of the field, he is chasing after the ball. He acts similarly to the Attacker, but in complementary halves of the field.

DEF The Defender places himself between the ball and the own goal, with the only restriction that he doesn't move past the middle line, even when the ball is on the opponent's half of the field.

Passive Defense

ATT While the ball is inside the opponent's half of the field, the Attacker moves along the middle line and places himself between the ball and the own goal. As soon as the ball enters his own, the tactic switches to Pressing Defense.

MID The Midfielder supports the Attacker by staying behind him visually tracking the ball at all times.

DEF The Defender stays between the ball and the own goal in front of his own penalty line.

4 Petri Net Plans Implementation

We designed the roles of the players using Petri Nets Plans [9], expanding the infrastructure of the SPQR-Legged team. The formalism of Petri Net Plans allows for complex action interactions, such as synchronization, concurrency, non-instantaneous execution, and interrupts.

A Petri Net [7, 8] is a graphical language used in modeling dynamical systems. It allows the description of a system in terms of a weighted directed graph, where the nodes denote places (circles) or transitions (boxes) and the edges between them represent possible paths of execution as well as certain conditions (weights). The flow of execution is denoted by means of tokens, which move between places causing transitions to "fire", if all the specified conditions are met. The Extended Petri Net variation [8] allows for inhibitory edges, whose firing behavior is complementary to that of the regular edges.

A Petri Net Plan [9] is a collection of actions structured as an Extended Petri Net with unweighted edges. Each action in the plan is explicitly described in three phases (initiation, execution, termination) corresponding to an equal number of places with two transitions between them. This description allows for constructing sequences, loops, splits, joins, and parallel branches of actions. Petri Net Plans have been successfully used in various robot domains, such as four-legged RoboCup, search-and-rescue, and multi-agent communication.

Each Petri Net Plan is defined in terms of various base actions. The degree of complexity hidden inside the base actions is inversely proportional to the degree of complexity

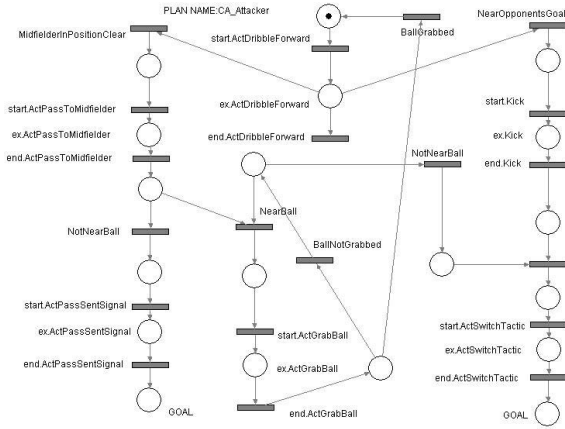


Figure 3. PN Plan for the Attacker role in the Counter Attack tactic.

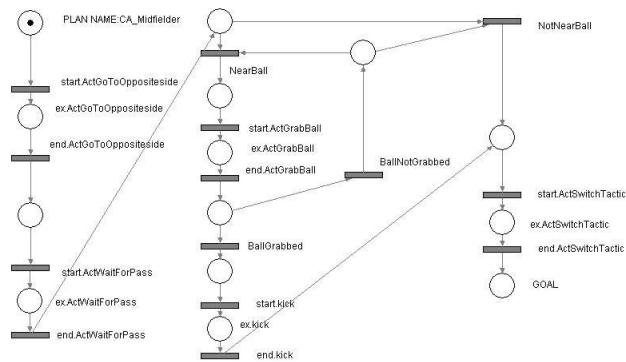


Figure 4. PN Plan for the Midfielder role in the Counter Attack tactic.

in the plans. We have chosen a set of base actions, which strikes a balance between the complexity of plans and actions, guided by the desire to design reusable actions that can be used in many different plans.

The Petri Net Plan implementing the Attacker role for the Counter Attack tactic is shown in Figure 3. Execution of this plan begins at the place with the single token (black dot). The first action (`ActDribbleForward`) is executed continuously until it is interrupted by one of two events. Either the Attacker reaches the opponent's goal area (condition `NearOpponentsGoal`) or figures out that the Midfielder is in position and cleared for a pass (condition `MidfielderInPositionClear`). In the first case, the Attacker executes the action `kick` and signals the team to choose the next tactic. In the second case, the Attacker executes the `ActPassToMidfielder` action to pass the ball. A sensing action follows. If the ball cannot be seen

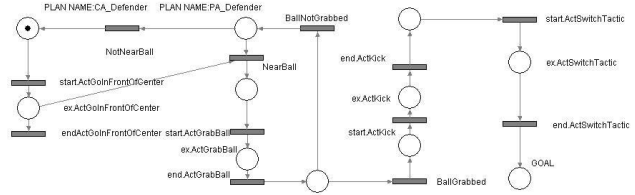


Figure 5. PN Plan for the Defender role in the Counter and Passing Attack tactics.

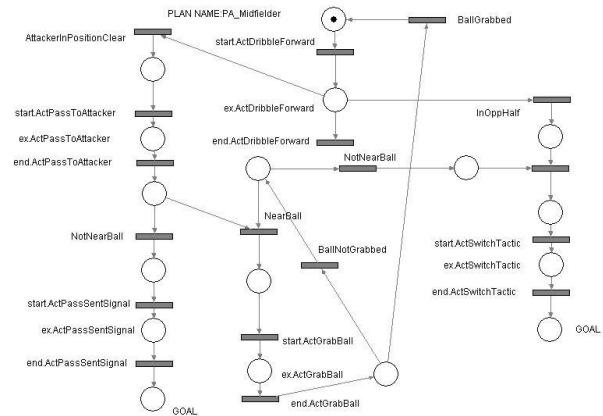


Figure 6. PN Plan for the Midfielder role in the Passing Attack tactic.

anymore near the Attacker, he notifies the Midfielder that a (hopefully successful) pass was made. However, if and as long as the ball is still seen near the Attacker, he repeatedly attempts to grab it (action `ActGrabBall`) and repeat the plan. Otherwise, if the ball cannot be seen anymore near the Attacker, he signals the team to choose the next tactic.

The Petri Net Plan for the Midfielder role in the Counter Attack tactic is shown in Figure 4. The Midfielder moves to the opposite side of the field and waits for a pass (along with the signal). If the pass is successful (`NearBall`), he grabs and kicks the ball to the opponent's goal, otherwise he repeatedly tries to grab the ball as long as it can be seen near the robot. In any case, if the ball moves far away from the Midfielder, he calls for a switch of tactic.

The Petri Net Plan for the Defender role in the Counter Attack tactic is shown in Figure 5. The Defender takes position in front of the center circle and waits. In the unexpected event that the ball bounces back to him, he attempts to grab it and kick it towards the opponent's goal. If successful, he signals for a switch of tactic, otherwise he returns to the top of the plan.

The Petri Net Plan for the Midfielder role in the Passing

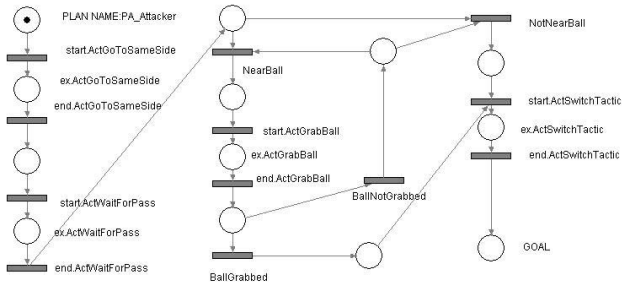


Figure 7. PN Plan for the Attacker role in the Passing Attack tactic.

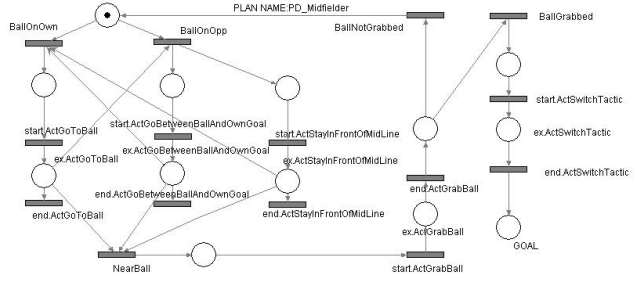


Figure 9. PN Plan for the Midfielder role in the Pressing Defense tactic.

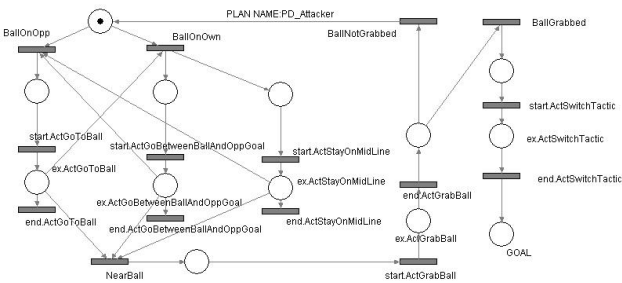


Figure 8. PN Plan for the Attacker role in the Pressing Defense tactic.

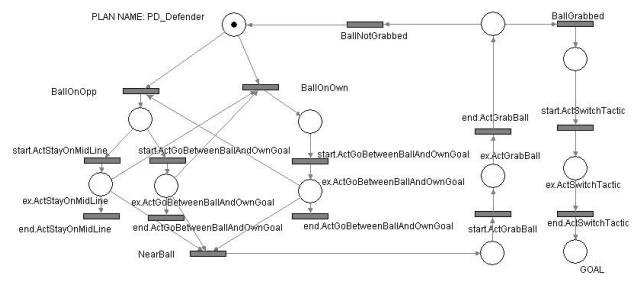


Figure 10. PN Plan for the Defender role in the Pressing Defense tactic.

Attack tactic is shown in Figure 6. The Midfielder starts by moving forward until one of two events occurs. Either he enters the opponent's half and signals a switch of tactic to Counter Attack or identifies the Attacker further ahead and passes the ball. Similarly to the Attacker in Counter Attack, he either notifies for a possibly successful pass or repeatedly tries to recover the ball and repeat the plan.

The Petri Net Plan for the Attacker role in the Passing Attack tactic is shown in Figure 7. The Attacker moves to the same side of the field as the Midfielder, but further ahead into the opponent's half, constantly facing backwards, and waits for a pass. If the pass is successful, he grabs the ball and calls for a switch of tactic, otherwise he repeatedly tries to grab the ball as long as it is seen near the robot. If the ball is lost, he calls for a switch of tactic.

The Petri Net Plan for the Defender role in the Passing Attack tactic is identical to the one in the Counter Attack shown in Figure 5.

The Petri Net Plan for the Attacker role in the Pressing Defense tactic is shown in Figure 8. Initially, a choice of course is made depending on the current location of the ball. If the ball is on the opponent's half of the field, the Attacker continuously chases after the ball. If the ball is in his own half, then a fork occurs and the Attacker continuously exe-

cuts two actions in parallel; one which keeps him around the middle line and one which moves him between the ball and the opponent's goal. If the ball switches from one half to another, the choice is adjusted appropriately. Finally, the Attacker breaks out of these loops, if he finds himself next to the ball. In this case, he repeatedly attempts to grab it; if successful, he signals a switch of tactic to offence, otherwise he repeats the plan.

The Petri Net Plan for the Midfielder role in the Pressing Defense tactic is shown in Figure 9. The careful reader will notice that this plan is almost identical to that of the Attacker. The only difference is that the conditions `BallOnOpp` and `BallOnOwn` have been switched. This is expected as these two roles are complementary with respect to the two halves of the field.

The Petri Net Plan for the Defender role in the Pressing Defense tactic is shown in Figure 10. The Defender never chases after the ball, but positions himself between the ball and his own goal, without ever leaving his own half.

The Petri Net Plan for the Attacker role in the Passive Defense tactic is shown in Figure 11. Again, a choice is made depending on the current location of the ball. If the ball is found in the own half of the field, the Attacker signals a switch of tactic to Pressing Defense. Otherwise, he stays

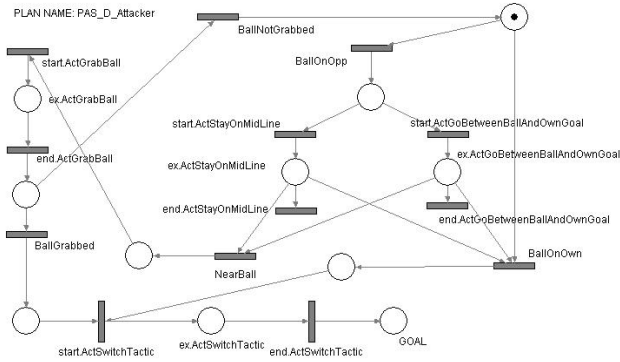


Figure 11. PN Plan for the Attacker role in the Passive Defense tactic.

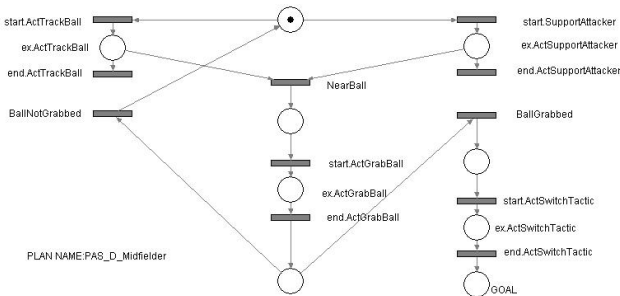


Figure 12. PN Plan for the Midfielder role in the Passive Defense tactic.

around the middle line between the ball and his own goal, looking for opportunities to grab the ball.

The Petri Net Plan for the Midfielder role in the Passive Defense tactic is shown in Figure 12. The Midfielder supports the Attacker by staying behind him and at the same time visually tracks the ball. He breaks out of this loop only if he manages to steal the ball or a signal is received to switch tactic.

Finally, the Petri Net Plan for the Defender role in the Passive Defense tactic is shown in Figure 13. This plan is almost identical to that of the Midfielder except that the Defender moves in front of the own penalty line and tries to position himself between the ball and the own goal. Again, he breaks out of this loop only if he manages to steal the ball or a signal is received to switch tactic.

5 Dynamic Tactic Selection

The tactic to be played at any moment is selected dynamically depending on the current situation of the game. Coordination between the robots takes place through a Fi-

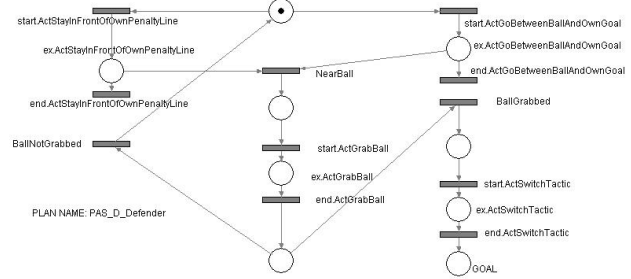


Figure 13. PN Plan for the Defender role in the Passive Defense tactic.

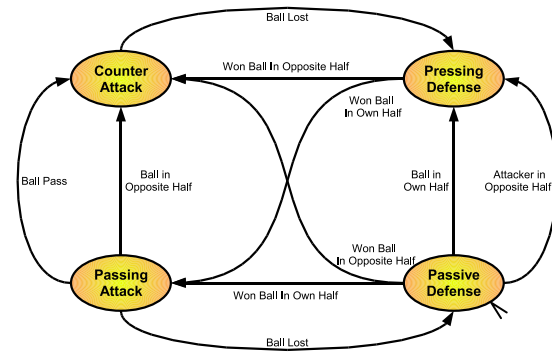


Figure 14. FSM for Switching Tactics.

nite State Machine (FSM), which implements a simple selection protocol. Each robot executes its own local copy of the FSM and the uniqueness of role allocation is guaranteed by appropriate conditions over the information shared between the robots. In particular, the FSM is triggered by the signals broadcast by the robots, whenever a switch of tactic is deemed necessary. Any information needed by the FSM to make a transition, such as a pass attempt and ball/player position, is provided directly by the robots. The complete state diagram of the FSM along with the necessary conditions is shown in Figure 14.

Apparently, any robot winning/losing the ball signals the team to choose an offending/defending tactic. While on offense/defence, a pass, the ball position, or the attacker location may also change the tactic. Transitions to the same state are possible in case of false alarms for tactic switch. Finally, if a player is penalized, the team automatically switches to Passive Defence (with a missing Attacker).

6 Robot Communication

Communication between robots in the four-legged RoboCup league faces several challenges. Besides signals transmitted through the physical world (e.g. sound or light),

robots are allowed to exchange information through a low-bandwidth, unreliable, wireless network. This fact implies that only limited amounts of information can be exchanged, messages are not guaranteed to be delivered to all recipients, several packets may be delivered with delay or out of order, and simultaneous transmission of messages will result in loss with high probability. In addition, experience has shown that the entire network can go down unexpectedly or individuals robots can go off-line due to hardware failure. It is therefore understood that, under these conditions, careful network usage is required for reliable communication.

In our implementation, all communication takes the form of messages sent by a single robot to all other team members. This is required because each player executes its own local copy of the FSM and therefore trigger events must become common knowledge to ensure that all robots switch to the same tactic. Note that any tactic switching in the FSM is triggered by a single message. Thus, there is no need for simultaneous and/or synchronized messages and as a consequence the network is not congested with multiple different messages, which need to go through at the same time.

We use the UDP protocol for broadcasting each message to all robots including the original sender. Any player originating a message adds a special tag to the message indicating its own identity and a time stamp with the original submission time. Given that only a single message is broadcast over the network at any point in time, all players in the team (including original senders) are required to retransmit any message they receive stamped with the most recent submission time to ensure that all messages eventually reach all players. Messages stamped with an older submission time are not retransmitted and eventually cease from the network. In other words, the message that was broadcast last over the network is continuously circulated until the next message appears. We adopted this scheme to cope with the unreliability of the network. Given that the size of all these messages is rather small, even the resulting continuous traffic on the network does not lead to network congestion or violation of the bandwidth limit. The above communication scheme achieves delivery times to all recipients comparable to the network lag (1.5 seconds) for 72% of the messages and delivery of virtually all messages (99.7%) within 10 seconds. This result was obtained using all four robots and a standard wireless router.

If a robot is inactive (does not control the ball) and receives no message for a predefined period of time (currently, 3 seconds), it assumes that the network is down and adopts the most conservative tactic (Passive Defense) until a new message arrives. That also covers the possibility of a robot being off-line, which is not rare in RoboCup games. This convention may result in a robot temporarily adopting a tactic that is possibly different than the team tactic, however it prevents the robot from entering a deadlock, where it waits

indefinitely for an incoming message before taking action.

The passing mechanism needed for the offense tactics is also realized using communication; teammates exchange over the network the necessary information for making a successful pass. In particular, as soon as a robot is ready to pass the ball, it broadcasts an *intent-to-pass* signal along with its own coordinates in the field (available through the localization module). Provided a robot is available to receive the pass, it orients itself facing the passing robot, broadcasts its own position, and waits. The passing robot kicks the ball towards the receiving robot and broadcasts a *sent-pass* signal. Assuming that the area between them is not obstructed and given the fairly good accuracy of localization, there is high chance that the ball will end up near the receiving robot, thereby realizing a successful pass.

7 Results

For successful operation, our approach must be supported by an accurate localization module (for moving to and broadcasting correct positions), a passing module (for passing the ball in the correct direction), and a communication module (for transmitting signals). To provide these support services, we embedded PNP coordination within the software architecture of our RoboCup team Kouretes.

The performance of the proposed coordination approach cannot be demonstrated on paper, therefore we have produced video clips showing the roles and tactics in action¹. The proposed coordinated team play and tactics were employed during the games of the RoboCup German Open 2007 competition. However, they were not fully utilized due to severe vision and localization problems and thus their performance in practice during actual games could not be fully assessed. Home tests revealed that indeed teammates make better use of the field with better positioning, even in situations where only a single robot is aware of the ball position. As indicated by our experiments captured on video, we have reasons to believe that the proposed coordination approach enforces a good degree of teamwork which was largely missing from the four-legged league.

8 Related and Future Work

Most teams in the four-legged league follow a role switching model for coordination. Team ARAIBO emphasizes individual skills and implements a rather simple role switching module. All robots calculate the time they need to get to the ball and transmit it through the network; the robot with the smallest time becomes the attacker, the robot with the second smallest time becomes the supporter, and the remaining robot assumes a predefined role. Despite its

¹Available from <http://www.intelligence.tuc.gr/kouretes>

simplicity, this approach is not robust against network failures and may lead to a poor offensive strategy. A similar strategy is followed by the German Team. Again, robots assume roles according to the time they need to get to the ball, however there is a backup role assignment in case of network failures. Robots are required to share their individual cost functions over the network at all times. The rather complicated backup strategy depends solely on the position of the ball and may lead to bad formations in certain cases (e.g. when the ball is on the middle line).

Other teams follow a higher-level team strategy approach. In team Cerberus, robots exchange bids and run a virtual auction over the network to decide which robot will do what. Despite the low network utilization, crucial tasks may end up to inappropriate robots, if bid information is lost over the network. Team NUbots employs a scheme, which takes into account the current score and the remaining time of the game. The field is divided into localized positions (left defender, center defender, etc.) and the robots can take upon one of two roles: chasing or positioning. The robot with the smallest distance to the ball takes the chasing role, while the other two robots go to the closest localized positions. Network traffic is kept at minimum levels, however the team does not always act in a coordinated manner, due to poor communication between the robots.

The main difference of our approach is that coordination takes place at team and not player level. The team decides on a common tactic (through the current message on the network and the FSM) and each robot automatically knows its role within that tactic and the conditions under which the tactic and roles may change. Therefore, there is no need for low-level, intense communication to determine role assignments to robots. This abstraction adds flexibility to switching between a variety of predefined simple or complex tactics, as well as robustness against network problems.

In the future, we plan to extend our coordination method along two dimensions. Auction-based coordination [2, 5] borrows methods and techniques from economics (auctions, trading) and applies them to agent problems. Tasks are assigned to agents through an auction procedure, whereby agents place bids on tasks and winners are determined through a clearance procedure that optimizes a team objective. Auctions could be used for efficient dynamic role assignment with a chosen team tactic. In addition, multi-agent reinforcement learning offers opportunities for collaborative multi-agent learning (where many agents learn to collaborate as a team) [3] and competitive multi-agent learning (where two teams learn to compete against each other, but collaborate within the team) [6]. The scaling properties of these algorithms through exploitation of domain knowledge make them attractive for the RoboCup domain as factorization of the representation can be done on the basis of proximity between players during a game. Us-

ing such learning methods, appropriate tactic switching decisions could be learned directly through trial-and-error during actual RoboCup games, thereby enhancing team coordination with adaptive capabilities and replace the current static FSM-based scheme.

9 Conclusion

We presented a method for achieving coordinated team play in a RoboCup team, a feature that has not been emphasized much in the Four-Legged league. The formalism of Petri Net Plans allows the specification of tactics and roles for coordination in an intuitive and systematic way. The proposed method can be adapted and used beyond the four-legged league to any multi-robot domain with well-defined roles and objectives. Our long term goal is to develop a tool to let users set-up automatically various formations, tactics, and roles for any number of players in a team.

Acknowledgments

We would like to thank Fr. Demetrios Alexandrakis and the parish of Panagia in Kounoupidiana for providing valuable laboratory space to the team when it was most needed. We also like to thank the administration of the Technical University of Crete for the financial support of the team.

References

- [1] Anonymous. Formation (football). *Wikipedia*, 2007.
- [2] M. Dias, R. Zlot, N. Kalra, and A. Stentz. Market-based multi-robot coordination: a survey and analysis. *Proceedings of the IEEE*, 94(7):1257–1270, 2006.
- [3] C. Guestrin, M. G. Lagoudakis, and R. Parr. Coordinated reinforcement learning. In *Proceedings of the 19th International Conference on Machine Learning*, pages 227–234, 2002.
- [4] H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, , and H. Matsubara. Robocup: A challenge problem for AI. *AI Magazine*, 18(1):73–85, 1997.
- [5] M. Lagoudakis, V. Markakis, D. Kempee, P. Keskinocak, S. Koenig, C. Tovey, A. Kleywegt, A. Meyerson, and S. Jain. Auction-based multi-robot routing. In *Proceedings of Robotics: Science and Systems*, pages 343–350, 2005.
- [6] M. G. Lagoudakis and R. Parr. Learning in team markov games using factored value functions. In *Proceedings of NIPS*2002: Neural Information Processing Systems: Natural and Synthetic*, pages 1659–1666, 2002.
- [7] T. Murata. Petri nets: Properties, analysis and applications. *Proceedings of the IEEE*, 77(4):541–580, 1989.
- [8] J. L. Peterson. *Petri Net Theory and the Modeling of Systems*. Prentice Hall, Upper Saddle River, NJ, 1981.
- [9] V. A. Ziparo and L. Iocchi. Petri net plans. In *ATPN/ACSD Fourth International Workshop on Modelling of Objects, Components, and Agents*, pages 267–289, Finland, 2006.