Skills, Tactics and Plays for Distributed multi-robot control in Adversarial environments

Lotte de Koning[1], Juan Pablo Mendoza[2], Manuela Veloso[2], and René van de Molengraft[1]

[1] Departement of Mechanical Engineering, Eindhoven University of Technology, The Netherlands, c.k.m.d.koning@student.tue.nl, m.j.g.v.d.Molengraft@tue.nl
[2] School of Computer Science, Carnegie Mellon University, Pittsburgh, United States of America, jpmendoza@ri.cmu.edu, mmv@cs.cmu.edu

Abstract. This work presents a pioneering collaboration between two robot soccer teams from different RoboCup leagues, the Small Size League (SSL) and the Middle Size League (MSL). In the SSL, research is focused on fast-paced and advanced team play for a centrally-controlled multi-robot team. MSL, on the other hand, focuses on controlling a distributed multi-robot team. The goal of cooperation between these two leagues is to apply teamwork techniques from the SSL, which have been researched and improved for years, in the MSL. In particular, the Skills Tactics and Plays (STP) team coordination architecture, developed for centralized multi-robot team, is studied and integrated into the distributed team in order to improve the level of team play. The STP architecture enables more sophisticated team play in the MSL team by providing a framework for team strategy adaptation as a function of the state of the game. Voting-based approaches are proposed to overcome the challenge of adapting the STP architecture to a distributed system. Empirical evaluation of STP in the MSL team shows a significant improvement in offensive game play when distinguishing several offensive game states and applying appropriate offensive plays.

Keywords: Robot soccer, Multi-robot system, Distributed coordination, Team plan execution

1 Introduction

RoboCup is an international robot soccer organization, founded to promote research in robotics and artificial intelligence. The ultimate goal of RoboCup is to beat, by 2050, the winner of the most recent human soccer World Cup with a team of fully autonomous humanoid robot soccer players, complying with the official rules of FIFA. To accomplish this goal, teams compete annually in various soccer leagues, each of which focuses on the different challenges of robot soccer. This work is related to two of those soccer competitions: Small Size League[3]

[3] http://wiki.robocup.org/wiki/Small_Size_League

(SSL) and Middle Size League[4] (MSL) (Figure 1). SSL is specialized in fast-paced and advanced team play within a central coordinated multi-robot team. In MSL on the other hand, research focuses on coordination of a distributed multi-robot team. Both teams operate in adversarial environments. In order to reach the RoboCup goal, research and accomplishments of the different leagues should be brought together. This work presents a co-operation between the SSL team CMDragons from the Carnegie Mellon University in Pittsburgh (United States) and the MSL team Tech United from Eindhoven University of Technology (the Netherlands).

Much research [3, 8, 9, 12] has been done on team planning of centrally coordinated multi-robot teams. However, when the number of robots in the system increases, computations exponentially increase in complexity. Distributed multi-robot systems are a solution to this computation problem [11], as each individual robot computes its tasks. As high-level team planning is the same for both centrally coordinated and distributed multi-robot systems, the research on team planning in centrally coordinated systems can be beneficially used in distributed systems. This work presents an integration of a planning algorithm designed for a centrally coordinated multi-robot team into a distributed multi-robot team.

The case study used for this work is the RoboCup environment. Towards the RoboCup goal of beating a human soccer team, the number of robots in the team will increase to 11, therefore a distributed solution is desired [11]. The advanced planning strategies developed by the SSL over the past years [1, 2, 6], can be used in the distributed teams, such as the MSL. For this work, the Skills, Tactics and Plays (STP) planning algorithm [2], developed by the CMDragons team for small size soccer robots, is integrated into the Tech United middle size robots to increase the level of team play in Tech United attacks. Compared to planning algorithms designed for multi-robot teams [3, 8, 9, 12], the STP architecture is specifically designed to control an autonomous multi-robot team in a dynamic environment with the presence of adversary's. The architecture contains predefined team plans, the plays, which are chosen during game play based on the world state. The tactics and skills associated with a play define the single robot behaviour during the execution of a play.

This paper presents the approach and challenges of integrating an architecture developed for a centralized multi-robot team into a distributed multi-robot team. In particular, a distributed team of robots, with differences in their estimated state of the world, needs to agree on a team plan and role assignments in a fair way. To overcome this problem, voting systems are considered to determine which play to use, as well as their optimal role assignment [4, 5, 7]. The new algorithm in the Tech United MSL team is evaluated with simulations and show that using the SSL approach can significantly improve the level of team play.

The remainder of the paper is organized as follows: Section 2 presents briefly the current Tech United strategy; Section 3 explains the Skill, Tactics and Play

[4] `http://wiki.robocup.org/wiki/Middle_Size_League`

(STP) architecture designed by the CMDragrons team; Section 4 describes the integration of STP in the Tech United software and finally; Section 5 gives the simulation results which show the improvement of the Tech United's offense; the paper is concluded in Section 6.

(a) Small Size League at RoboCup European Open 2016, Eindhoven, the Netherland

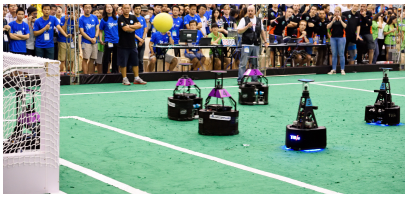

(b) Middle Size League at RoboCup 2015, Hefei, China

Fig. 1: Both competitions at RoboCup tournaments

2 Tech United MSL strategy

For this study STP is used to integrate offensive plays into the Tech United team. The current Tech United offensive strategy assigns roles to the active robots, which include specific tasks. The goalkeeper defends the goal, two defenders position between the ball and their own goal, and two attackers execute an attack on adversary's half. Two attacker roles are distinguished: the AttackerMain is the role that manipulates the ball and the AttackerAssist positions on the adversary's half to potentially receive the ball. Roles execute tasks which involve decision making and selecting skills to complete the task.
This current implementation does not allow the team to easily change strategies shortly before a game or add new strategies to the software. Using the STP architecture, this will be possible.

3 Skill Tactics Play algorithm

The Skills, Tactics and Plays (STP) architecture was designed by CMDragons to coordinate a centralized team of autonomous robots to achieve long-term goals [2]. This architecture enables simple specification of team plans as plays, which are then executed individually by each robot through tactics and skills. STP enables the team to adapt their team strategy as a function of the state of the game, resulting in highly versatile teamwork. This section describes each component of STP and how the CMDragons team uses STP for team planning.

3.1 STP Components

Skills At the most basic level, each robot is capable of performing a certain set of low-level skills. Skills are base behaviors that can be parameterized to achieve various goals. In the domain of robot soccer, these skills include navigating safely and quickly to a specific target location, shooting the ball with a specified velocity, and intercepting a moving ball, among others. Skills alone do not encode goal-oriented behavior, but are used in a set of skills to achieve goals.

Tactics Tactics are goal-oriented behaviors that each individual robot performs to carry out a team plan. These behaviors may be composed of skills, organized to work cohesively towards achieving a goal. In the CMDragons team, tactics are implemented as finite state machines that control the flow of the different skills that make up a tactic.

Each role has a specific set of tactics to perform. Some of the most prominent tactics in the CMDragons team include the Goalkeeper to block shots from the adversary, Defenders to prevent the adversaries from passing and shooting, the AttackerMain to manipulate the ball to score goals on the adversary, and AttackerAssists to position themselves to receive a pass from the AttackerMain. Tactics, like skills, may be parametrized. For example, the AttackerAssist may restrict their search for a good pass location to a specific region of the field.

Plays In STP, team strategy is encoded into a playbook, made up of a set of plays. A play is a team plan specified as a list of roles that the team of robots must fulfill. Each role, within the context of STP, is defined as a sequence of tactics to be completed sequentially. Additionally, plays specify applicability conditions: *preconditions* specify the set of game states under which a play should be considered for selection, while *invariants* specify the set of game states under which the team does not need to abort this play and select a new one.

Table 1 shows a slightly simplified example of a play in the CMDragons playbook. This play is applicable when all preconditions match the world state: the ball is in the adversary's half of the field (their_side), and there are no opponents on our half of the field (!opp_our_side), and our team has possession of the ball (offense). Furthermore, if this play is selected for execution, it will remain in execution until either the invariants become false, –i.e., the ball moves to our side of the field (their_side becomes false) or the opponent team gains possession of the ball (!defense becomes false), or all the robots have finished performing their role.

3.2 STP Procedure

During each step of execution, the team selects a joint plan, optimally assigns roles to each robot, and then each robot individually executes its role. This section describes this procedure, with an emphasis on the team components.

Table 1: An example of a play

```
PRECOND their_side && !opp_our_side && offense
INVARIANT their_side && !defense
NUMROBOTS 1 6

ROLE 0
    Goalkeeper
ROLE 1
    AttackerMain
ROLE 2
    Defender
ROLE 3
    AttackerAssist front_center
ROLE 4
    AttackerAssist front_left
ROLE 5
    AttackerAssist front_right
```

Play Selection First, the team's central intelligence decides whether the play that was last used is still applicable. That is, it checks if the last play's invariants are still true, and if at least one of the robots has not yet finished its role. If these are both true, no new play is selected; if it is false, a new play is selected from the set of plays whose preconditions hold. If there are multiple plays whose preconditions hold, a play among the set is chosen randomly, with potentially different probabilities for each play.

Optimal Role Assignment Once a play has been selected, each of the roles in the play is assigned to a robot in the team. Given N active robots $\mathcal{P} = \{\rho_0, \rho_1, \ldots, \rho_N\}$, STP optimally assigns the first N roles $R = \{r_0, r_1, \ldots, r_N\}$ to them optimally, based on a cost function $C : \mathcal{P} \times R \to \mathbb{R}$. In the CMDragons, this cost function $C(\rho_i, r_j)$ is an estimate of the time that it would take robot ρ_i to complete role r_j. To do this, a cost matrix $\mathcal{C}$ of size $N \times N$ is created, such that the entry at row i and column j corresponds to the computed cost $C(\rho_i, r_j)$. From this cost matrix, algorithms such as the Hungarian algorithm [10] can be used to find the assignment from roles to robots that minimizes the total cost. It is important to note that the costs in matrix $\mathcal{C}$ are computed from the centrally-estimated state of the world, common to all the robots in the team.

The optimal role assignment computation used by the CMDragons is similar to the approach of Tech United, the only difference are the costs. Where the SSL team uses time, the MSL team uses distances. Assuming constant speed during repositioning we can consider time and distance to be equivalent.

Individual Role Execution After STP distributes the roles of the selected play to each of the robots, they proceed to individually execute their assigned tactic. In

SSL, this step also occurs in a centralized controller, with only the final motion command being sent to each individual physical robot. However, the computation of each tactic happens in parallel, with only very limited communication between robots at this stage.

3.3 STP in distributed team

The bulk of the problem of adapting the STP architecture to the MSL lies in the previous two steps: play selection and optimal role assignment. In SSL both steps are performed by a central controller, in MSL however, each robot computes strategy individually. Therefore, the play selection and role assignment is done in five separate computations. To execute a team plan towards a common goal, it is important for the team to agree on the selected play and role assignment. Section 4 discusses in detail the integration of this joint play selection and role assignment in MSL.

4 Integration of STP in MSL

The integration of STP involves four main components from the Tech United software: i) communication ii) the game state evaluation, iii) the role assignment and iv) the tactics execution. Figure 2 is an overview of the strategy components needed for the STP integration. Highlighed are the new designed components, the playbooks with plays, the play selection algorithm and the tactics, and the re-used and extended components, communication, game state evaluation and role assignment. This section explains each component in more detail.

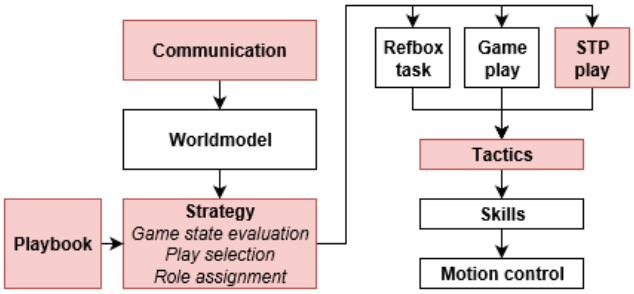

Fig. 2: Three main components for STP: playbook, strategy, play execution.

4.1 Overview STP algorithm in MSL

Algorithm 1 shows the complete STP integration algorithm. First the world state is evaluated on the preconditions and based on those preconditions a play

may be selected (lines 1-2). If a play is selected, the play roles are assigned to the active robots using the role assignment (lines 4-5). In the execution part, first the robot checks whether its team mates are still executing the same play, to ensure all robots execute the same team plan. In addition, the invariants are checked whether those are still true (lines 7-10). If one of those is false, the robot will abort the play and updates its play status to NOT_IN_PROGRESS. Finally, the robot uses the *play_selected* variable from the play selection, and the play status to determine whether a play or the default game play will be executed (lines 12-19). Default game play refers to the strategy as explained in Section 2.

Algorithm 1 Algorithm executing a play

```
 1: world.Precond ← evaluateWorldState()
 2: play_selected ← playSelection(world.Precond)
 3:
 4: if (play_selected) then
 5:     assignedRoles ← RoleAssignment(play_selected)
 6:
 7: if (PLAY_IN_PROGRESS) then
 8:     peerInProgress ← checkPlayStatePeers()
 9:     if (!peerInProgress OR !play.invariants) then
10:         PLAY_NOT_IN_PROGRESS
11:
12: if (PLAY_NOT_IN_PROGRESS) then
13:     if (play_selected) then
14:         execute_play(assignedRoles)
15:         PLAY_IN_PROGRESS
16:     else
17:         execute_default_play()
18: else if (PLAY_IN_PROGRESS) then
19:     execute_play(assignedRoles)
```

4.2 Plays and playbooks

A play P is a predefined team plan which executes a sequence of tactics T_i per robot. The playbook contains a selection of N plays: $\{P_0, P_1, ..., P_N\}$, this selection can specifically be chosen before a game, dependent on the opponent. During a game, one playbook is used. The active play is selected based on the world state. One of the two designed plays for this work, is given by Table 2. The play set-up is similar to the set-up designed by the CMDragons (Table 1). Preconditions and invariants are defined as world state conditions such as ball possession (ourBall), ball location (ballZone) and number of active opponents. Furthermore, the set of roles R is given with several parameters. First the type of role, which includes the sequence of tactics $\{T_1, T_2, ..., T_N\}$. Secondly, a target position, which is used for role assignment. The sequence of the N roles

given by the play $R = \{r_1, r_2, ..., r_N\}$ indicates the priority of each role when assigning roles to the set of robots $\mathcal{P}$. In this case, the AttackerMain is the most important role and the Defender the least important. Section 4.4 discusses this role assignment. Finally, in this play role 1 and 2, the two AttackerAssist roles, are assigned to a specific zone in the field. While positioning during the play execution, the role is bounded to this zone. Figure 3 shows the specified zones and target positions for the roles.

Table 2: An offensive play designed for MSL with three offensive robots and one defensive robot

AttackWithThree
PRECOND ourBall && ballZone1 && 4
INVARIANTS ourBall
role[0]
Goalkeeper
GOAL
role[1]
AttackerMain
BALL_POS
role[2]
AttackerAssist
PENALTYAREA_CORNER_LEFT
ZONE_A
role[3]
AttackerAssist
PENALTYAREA_CORNER_RIGHT
ZONE_B
role[4]
Defender
{0, -CIRCLERADIUS}

4.3 Play selection

For selection of a play from the playbook $\{p_1, p_2, ..., p_N\}$, first the game state is evaluated. The game state is evaluated for all preconditions of the plays, such as the number of opponents and the ball location in field. Each robot evaluates these preconditions individually based on the shared worldmodel.

Where in SSL play selection is done by one coordinator, the play selection in a distributed system is done by each robot individually. The robot compares the game state conditions with the preconditions in the plays. In this work, the playbook consists of two plays in addition to the default play. If no preconditions match the world state conditions, default game play is executed. When the world

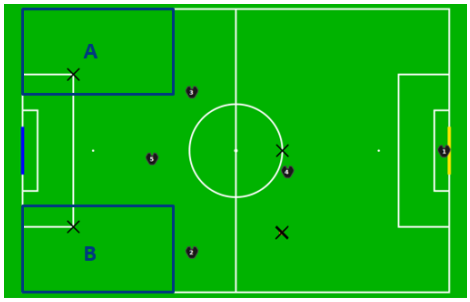

Fig. 3: Target positions for roles are indicated with a cross. The AttackerAssist roles are bounded for positioning to zone A and B.

state conditions do match the preconditions, the matching play is selected. To successfully execute a team plan, it is relevant that the team agrees on executing the same play. Therefore, each robot communicates the selected play with peers such that robots can individually determine whether its chosen play is correct. Each robot gathers the communicated selected plays in one list and sorts the selected plays. From this sorted list, the most common play is found (the mode). This play is feasible if the play is chosen by more than half of the number of active robots. If the mode is not feasible, no play will be selected and default game play is executed. If the mode is feasible, each robot will select the mode as the selected play.

4.4 Distributed role assignment

Once the optimal play has been selected given the state of the world, the roles for that play must be filled by the different robots. This is done following the current role assignment algorithm used by Tech United. The role assignment is based on the distance of the current robot position $x(\rho_i)$ to the fixed role position of the selected refbox task $x(r_j)$. Depending on the number of active robots N, N roles are assigned based on priority $R = \{r_1, r_2, ..., r_N\}$. Each robot computes a cost matrix C with the distances from each robot ρ_i to role r_j: $C : \mathcal{P} \times R \rightarrow \mathbb{R}$. Then, computations are done to determine the optimal role assignment where the total to be travelled distance by the robots from current position to role position is minimal: ${}_C(\sum \rho_i \in \mathcal{P}[C^i])$. To ensure all robots have the same role assignment, such that all roles are begin executed, all robots communicate their optimal role assignment. Currently, from these role assignment calculations, the role assignment of the active robot with lowest ID is used to divide the roles optimally.

4.5 Tactics execution

To execute the selected play, a sequence of tactics is performed by each role. Algorithm 2 shows an example of the tactics to be performed by the Attacker-Main role. A tactic may involve decision making, –e.g. choosing a pass receiver as

shown in the example. To complete the tactic, a set of skills is selected using the function doAction. Tactic transitions take place either after finishing a tactic, –e.g., when the AttackerMain possesses the ball the role transits into the give_pass state, or when a peer robot completes a certain tactic, –e.g., when the AttackerMain possesses the ball, the AttackerAssist roles transits to the receive_pass tactic. Therefore, communication is also required while executing tactics.

Algorithm 2 Sequence of tactics to be executed by the AttackMain.

```
1: switch AttackerMain.tactic do
2:     case intercept_ball
3:         doAction(intercept_ball)
4:     case give_pass
5:         pass_receiver ← choosePassReceiver()
6:         doAction(shoot_at_target, pass_receiver)
7:     case wait_to_end_play
8:         target ← findOptimalPosition()
9:         doAction(go_to_target, target)
```

5 Simulations and results

The integration of STP in the Tech United software is empirically evaluated with simulations. Two offensive plays were integrated, each for different world state, where the ball location precondition differs. Figure 4a show these ball locations. Play 1 will be executed when the ball is located in zone 1. Play 2 when the ball is located in zone 2. For all other ball locations, regular offensive game play, as explained in Section 2, is executed. The simulations are attacks from one of these zones with 4 robots: one goalkeeper and the first three roles defined by the play. These attackers are the black robots in Figure 4a. The attacks are performed against two defenders, the red robots in the Figure. The results of 60 attacks using the STP integration, are compared to the results of 60 attacks playing without STP.

5.1 Results

The difference in the attack using STP and not using STP is shown in Figure 4. The black robots are the attacking robots, the two red robots the defenders and the ball is orange and located in zone 1. The two defenders defend in between the ball and home goal. They follow the ball and try to block and intercept passes from the attackers. Figure 4b shows the positioning before the attackers are in ball possession. The AttackerMain (AM) intercepts the ball while the AttackerAssist (AA) roles are positioning within their assigned zone. The target

positions for the AA roles are indicated with the blue cross, the red cross indicates the target position of the AM. The difference with the attack without using STP, as shown in Figure 4c, is that two robots are positioning on the adversary's half. In Figure 4c can be seen that the AM intercepts the ball, while only one AA is positioning on adversary's half and the DefenderMain (DM) on their own half.

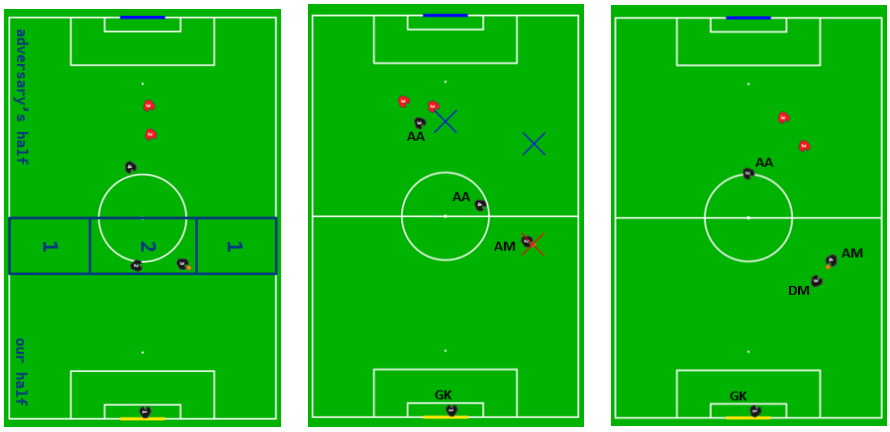

(a) For ball in zone 1, play 1 is executed, for ball in zone 2, play 2 is executed. For all other situations, regular game play is executed. Black robots: attackers, red robots: defenders.

(b) Attack with STP: AM intercepts the ball (target red cross), AA's position in given zone (target blue cross).

(c) Attack without STP: AM intercepts the ball, AA positioning on opponent half, DM positioning on own half.

Fig. 4: Simulations of attacks using STP and not using STP. Black robots are attackers, red robots defenders.

The results of the attacks from the two zones are given in Figure 5. Attacks were performed from each of the given zones (Figure 4a). When the attacking team made four passes, the attack was rated as ball possession for the team. If the attack finished before these four passes, the attack results either in a goal, a goal attempt or the ball was lost. For both plays the similar conclusions can be drawn.

The total number of lost balls, either a failed pass or an interception by the opponent, is equal for both the STP integration and no STP integration. Therefore, it can be concluded that the STP integration does not have an influence on this parameter. Furthermore, attacking without the plays, resulted mostly in ball possession for the team, the attackers did not find the chance to shoot at

goal. Using the STP integration on the other hand results in significantly more goals and goal attempts.

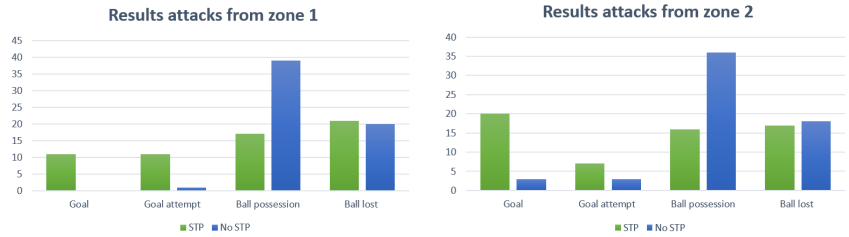

(a) Attacks from zone 1.

(b) Attacks from zone 2.

Fig. 5: Comparison of 60 attacks between a team using STP and a team not using STP. Graph shows results after max. 4 passes starting in the zone: goal, goal attempt, ball possession or the ball was lost.

6 Conclusions and future work

This paper presented a novel effort in adapting team-level strategy from the centrally-controlled Small Size League of RoboCup to the distributed Middle Size League. In particular, the Skills, Tactics and Plays (STP) team-planning architecture was successfully integrated in the Tech United team. The main challenges for this integration were the agreement on a common team plan to execute and optimally assign the team plan roles to the active robots. Voting-based approaches were used to overcome these challenges. Each robot individually selects a play and computes the costs for optimal role assignment. Both are communicated among all robots. The most common selected play is chosen by each robot and role assignment is computed based on an averaged cost matrix. Both methods are proven to be robust by the performed simulations.

Simulations with and without the STP integration were performed. Three robots performed attacks against to adversary's for game state situations for which two plays were designed. The results showw a significant improvement while applying appropriately-offensive play for specific game situations chosen for this work: ball possession in zone 1 and 2.

This work shows a successful cooperation between two different RoboCup leagues. While the SSL has been developing strategy algorithms for fast-paced team play over the years, in MSL teams have been focusing on controlling a distributed team. Integrating a well-developed strategy algorithm from the SSL, such as STP, into the MSL, helps the league many steps forward. Such collaborations, where knowledge and algorithms are shared among leagues, are desirable in order to accomplish the ultimate RoboCup goal: beating the human World Cup winner of 2050 with a fully autonomous humanoid robot soccer team.

References

1. Biswas, J., Mendoza, J.P., Zhu, D., Choi, B., Klee, S.D., Veloso, M.M.: Opponent-driven planning and execution for pass, attack, and defense in a multi-robot soccer team. In: Proceedings of AAMAS'14, the Thirteenth International Joint Conference on Autonomous Agents and Multi-Agent Systems (2014)
2. Browning, B., Bruce, B., Bowling, M., Veloso, M.: Stp: Skills, tactics and plays for multi-robot control in adversarial environments. In: Proceedings of the Institue of Mechanical Engineers, Part I: Journal of Systems and Control Engineering. pp. 219(1):33–52 (2004)
3. Brumitt, B., Stentz, A.: Dynamic mission planning for multiple mobile robots. In: Proceedings of the IEEE International Conference on Robotics and Automation. vol. No. 3, pp. 2396-2401 (1996)
4. Castelpietra, C., Iocchi, L., Nardi, D., Piaggio, M., Scalzo, A., Sgorbissa, A.: Coordination among heterogeneous robotic soccer players. In: Intelligent Robots and Systems, 2000.(IROS 2000). Proceedings. 2000 IEEE/RSJ International Conference on. vol. 2, pp. 1385–1390. IEEE (2000)
5. Farinelli, A., Scerri, P., Tambe, M.: Building large-scale robot systems: Distributed role assignment in dynamic, uncertain domains. In: Proceedings of Workshop on Representations and Approaches for Time-Critical Decentralized Resource, Role and Task Allocation. vol. 102 (2003)
6. Galindo, C., Fernández-Madrigal, J.A., González, J., Saffiotti, A.: Robot task planning using semantic maps. In: Robotics and Autonomous Systems. vol. 56 Issue 11 (2008)
7. Hunsberger, L., Grosz, B.J.: A combinatorial auction for collaborative planning. In: MultiAgent Systems, 2000. Proceedings. Fourth International Conference on. pp. 151–158. IEEE (2000)
8. Janssen, R.: Centralized learning and planning: for cognitive robots operating in human domains (2014)
9. Jensen, R., Veloso, M.: Obdd-based universal planning: Specifying and solving planning problems for synchronized agents in non-deterministic domains. In: Lecture Notes in Computer Science. vol. No 1600, pp. 213-248 (1999)
10. Kuhn, H, W.: The hungarian method for the assignment problem. In: Naval Research Logistics. pp. 83–97. IEEE (2006)
11. Parker, L.E.: Designing control laws for cooperative agent teams. In: Proceedings of the IEEE International Conference on Robotics and Automation. vol. pp. 582-587 (1993)
12. Svestka, P., Overmars, M.: Coordinated path planning for multiple robots. In: Robotics and Autonomous Systems. vol. Vol 7, pp. 83-124 (1997)