A Game-Theoretic Model of Multi-Robot Interactions*

Wojciech Szynkiewicz[1]

[1] Warsaw University of Technology, Institute of Control and Computation Engineering
ul. Nowowiejska 15/19, 00-665 Warsaw, Poland, email: W.Szynkiewicz@ia.pw.edu.pl

Abstract. This paper presents a formal model of multi-robot interactions based on dynamic game theory. The application of dynamic game theory involves a sequential decision process evolving in (continuous or discrete) time with more than one decision maker (in our case autonomous robot), one or more performance criteria (cost functionals), and possibly having access to different information. The whole range of robot control problems arising from different levels of "cooperation" between robots can be precisely described using different branches of game theory. If the robots have a common goal and one performance objective they act as a team, then team theory is relevant. A noncooperative game refers to the case in which robots have different goals and independent performance objectives. Also another very important issue in action planning i.e. interaction with unknown or partially unknown environment can be viewed as a game against nature.

1 Introduction

A multi-robot system involves a group of robots working in the same workspace. The use of multi-robot systems can enhance the utilization of robots, may increase the robustness of the system by taking advantage of parallelism and redundancy, and improve the versatility in handling different applications. Sometimes cooperation between multiple robots is an essential requirement for the successful completion of the task. Many other reasons for using multiple robots have been given in the literature e.g. [2, 11, 14]. However, the use of multiple robot systems is not without its challenges. One of the ultimate goals is to develop control strategies for multiple robots sharing the same space or working as a *team*. To solve this problem a research effort is needed to develop methods and algorithms to enable reliable, safe, and cooperative operation of free ranging autonomous robots performing tasks in dynamic semi-structured or unstructured (physical) environments.

There are many approaches to coordination of multiple robots that range from hierarchical planning solutions, in which a central coordinator designates tasks for individual robots [9,14], to purely reactive groups in which there is no central unit and cooperation is inherent in the implemented rules that control the robots [2,11]. In general, approaches to multiple robot control can be categorized into two groups: *centralized* and *decoupled* strategies. The centralized approach consists of treating all the robots as if they were

*This work was supported by the Polish Ministry of Science and Information Technology grant: 4T11A 003 25

one composite robot with several body elements that are not necessarily connected to each other. The centralized planning has the advantage that it allows for complete planners, which are guaranteed to find a solution whenever one exists. A main drawback to this approach is that it often leads to exploring large-dimensional space, which may be computationally very expensive.

A decoupled approach consists of planning activity for each robot individually while ignoring the existence of the other robots, and then considering the interactions between the robots. However, the decoupled techniques are not complete, for example, they may fail to find the solution if one exists. Of course, it is impossible to ensure optimal activity when all decisions are to be made locally. The suitability of one approach over the other is determined by the trade-off between computational complexity associated with a given problem, and the amount of completeness that is lost.

Game theory ideas can be useful as a powerful mathematical framework for both the development and analysis of multi-robot coordination algorithms [7]. Using game theory one can model different aspects of multi-robot systems, such as an amount and type of cooperation between robots, one or more performance criteria, uncertainty in environment sensing and state predictability. However, game theory is rather rarely used to model real multi-robot problems, in part due to its high computational complexity [7, 13]. In [5] partially observable stochastic games are used as a model for decentralized robot teams. The proposed algorithm transforms the original problem into a sequence of smaller Bayesian games that are computationally tractable. The application of non-cooperative game equilibria to the design of a collision free movement of a team of mobile robots in a dynamic environment is proposed in [12]. In [9] a cooperative game is formulated for two robots manipulating a single object to distribute generalized force, and the objective of the game is to minimize the worst case interaction force between robots and object.

In general, one can distinguish the following types of multi-robot systems: cooperative systems, competitive systems, adversarial systems (e.g., pursuit/evasion systems). In our concept, the degree of "cooperation" between robots (agents) becomes an important factor for determining which branch of game theory should be used (Fig. 1).

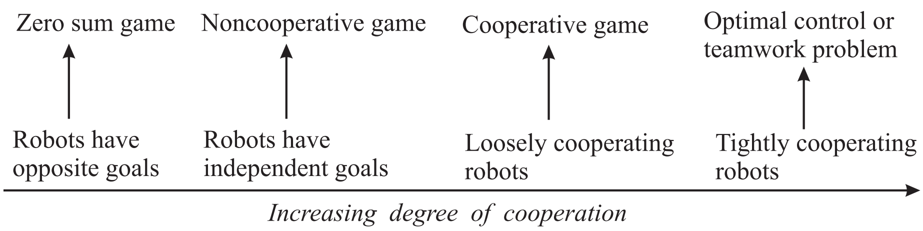

Figure 1. Game corresponding to the degree of robot cooperation.

A clear distinction exists between two-player zero-sum games and the rest [1]. *A zero-sum game* might be used to describe the situation, in which two players (in our case robots or robot teams) have totally opposite goals i.e. no cooperation between the players is allowed.

A noncooperative game refers to the case in which robots have different goals and independent performance measures [1,8]. Without cooperation the robots choose actions that take into account interests that conflict with the other robots.

If the robots act in unison but each has different decision objective, the *multiple objective optimization problem* is obtained [4]. A situation in which some part of the players wish to act in unison (i.e. they cooperate), in order to obtain a mutually beneficial outcome, can be described as a *cooperative game* [10]. This case can be referred to as *loosely cooperating* robots.

If all robots have a common goal and one performance objective, and they act as *a team* or even *a single mechanism*, then *team theory* or *optimal control theory* can be applied to solve the problem [14,16]. In this situation we have *tightly cooperating* robots.

Also another very important issue in multi-robot systems i.e., interaction with unknown or partially unknown environment can be viewed as *a game against nature*. In the real world, robots must always deal with uncertainty.

Dynamic games require an unambiguous and strict mathematical description of the problem to be solved. The class of related dynamic games problems that can be solved directly is fairly restrictive. As the models of the multi-robot system become more realistic the solution of the mathematical problem involved becomes more and more difficult, time consuming and expensive. One promising way to overcome this problem is based on combining Artificial Intelligence (AI) or especially Distributed Artificial Intelligence (DAI) i.e., multi-agent models and techniques with concepts of Dynamic Game Theory [3]. It seems to be a promising direction because the representation of a multiple robots coordination problem (essentially a multi-stage decision making process) by a decision tree is equivalent to a game in extensive form [1]. The pruning of such a tree is one of the basic tasks of heuristic search techniques used in most AI algorithms. Those techniques provide a way of solving real-life problems for which no more direct approach is available. Using AI techniques for the analysis of multi-robot systems allows a qualitative description of the problem, which is extremely important in complex situations. Rather than providing empirical results, in this paper we focus on the formal and theoretical aspects.

2 Problem formulation

Let $\{\mathcal{R}^1, \ldots, \mathcal{R}^N\}$ denote a collection of N robots operating in a shared workspace. Each robot may be an articulated robot (such as a manipulator) or a rigid robot (such as a mobile robot). The motion planning and control problem can be stated as a dynamic game that is played with one or more *players* also called *agents* or *decision makers* (in our case controllable robots). Sometimes we have an additional player, called "nature", whose actions influence the evaluation of the state of the game, leading to uncertainties.

We are interested in feedback control strategies (policies) that optimize one or more performance indices in an appropriate game-theoretic sense (such as a Nash equilibrium or Pareto optimality), which can be applied to different multiple robot coordination problems.

Let us consider a mathematical framework for the general motion planning problem formulated as N-person discrete-time infinite dynamic game in the extensive form [1]. The game (in deterministic case) involves the following components:

- A set, $\mathbf{N} = \{1, 2, \ldots, N\}$, of N *players* or, equivalently, *agents* (independent robots $\{\mathcal{R}^1, \ldots, \mathcal{R}^N\}$), and a set, $\mathbf{K} = \{0, 1, \ldots, K\}$ denoting the *stages* of the game, that correspond to moments at which decisions are made. K is the maximum number of the game stages.

- An infinite set $\mathcal{X} = \mathcal{X}^1 \times \mathcal{X}^2 \times \cdots \times \mathcal{X}^N$, where $\mathcal{X}^i$ is a state space of the i-th robot, called the *state space* of the game. The state of the game x_k belongs to $\mathcal{X}$ for all $k \in \mathbf{K}$ and $k = K+1$. The discrete-time representation induces a discretisation of the state space.
- An infinite set U_k^i, defined for each $k \in \mathbf{K}$ and $i \in \mathcal{X}$, which is called *control (action) space* of the i-th robot at stage k. Its elements are admissible *controls* u_k^i of $\mathcal{R}^i$ at stage k.
- A finite or infinite set Θ, which denotes the *action (control) space* of the additional $(N+1)$ player, *nature*, at stage $k \in \mathbf{K}$. Any admissible action θ_k of nature at stage k is an element of Θ. For example, future positions of the mobile robot are not completely predictable due to slipping of the wheels. Thus, the game-theoretic interpretation of this situation is that a nature player interferes with motion commands.
- A function f_k,
$$f_k : \mathcal{X} \times U_k^1 \times \ldots \times U_k^N \longrightarrow \mathcal{X},$$
defined for each $k \in \mathbf{K}$, so that
$$x_{k+1} = f_k(x_k, u_k^1, \ldots, u_k^N) \tag{1}$$
for some $x_0 \in \mathcal{X}$ which is the initial state of the game.
- stochastic case,
$$f_k : \mathcal{X} \times U_k^1 \times \ldots \times U_k^N \times \Theta \longrightarrow \mathcal{X},$$
defined for each $k \in \mathbf{K}$, so that The difference equation (1) is the state transition equation.
- A set, $\mathcal{Y}_k^i$, defined for each $k \in \mathbf{K}$ and $i \in \mathbf{N}$, and called the *perception* or *observation space* of $\mathcal{R}^i$ at stage k, to which the observation y_k^i belongs at stage k.
- A set, Ω^i, defined for each $k \in \mathbf{K}$ and $i \in \mathbf{N}$, which is called the *sensing action space for nature* at stage $k \in \mathbf{K}$. The sensing action ω_k^i, at stage k is an element of Ω_k^i, and $\omega_0^i, \omega_1^i, \ldots$ is the sequence of measurement errors, or measurement noise.
- A function h_k^i, $h_k^i : \mathcal{X} \to \mathcal{Y}_k^i$, defined for each $k \in \mathbf{K}$ and $i \in \mathbf{N}$, so that
$$y_k^i = h_k^i(x_k), k \in \mathbf{K},\ i \in \mathbf{N} \tag{2}$$
The equation (2) is the *observation equation* of $\mathcal{R}^i$ concerning the value x_k. The function h_k^i is a system output function.
- A finite set, η_k^i, defined for each $k \in \mathbf{K}$ and $i \in \mathbf{N}$ as a subset of $\{y_0^1, y_1^1, \ldots, y_k^1; \ldots; y_0^N, y_1^N, \ldots, y_k^N; u_0^1, u_1^1, \ldots, u_{k-1}^1; \ldots; u_0^N, u_1^N, \ldots, u_{k-1}^N\}$, which determines the information gained and recalled by $\mathcal{R}^i$ at stage k of the game. It worth knowing that information plays an integral role in choosing actions and this results in two different classes of games: games of imperfect information and games of incomplete information. In the former class all the aspects of the game are known to every player. While a player might not know what specific actions other players will take, it does have access to all the information upon which those decisions will be made. Extensive-form and normal-form games fall into this category. With this definition, each robot might have access to the actions and observations that where made by the other robots.

- A set $\mathcal{I}_k^i$, defined for each $k \in \mathbf{K}$ and $i \in \mathbf{N}$ as an appropriate subset of $\{(\mathcal{Y}_0^1 \times \mathcal{Y}_1^1 \times \ldots \times \mathcal{Y}_k^1) \times \ldots \times (\mathcal{Y}_0^N \times \mathcal{Y}_1^N \times \ldots \times \mathcal{Y}_k^N) \times (U_0^1 \times U_1^1 \times \ldots \times U_{k-1}^1) \times \ldots \times (U_0^N \times U_1^N \times \ldots \times U_{k-1}^N)\}$, compatible with η_k^i. Set $\mathcal{I}_k^i$ is called the *information space* of the i-th decision maker at stage k, induced by his information η_k^i. The information space can be considered as a replacement for the state space in the case of imperfect state information.
- A prespecified class Γ_k^i of mappings $\gamma_k^i : \mathcal{I}_k^i \to U_k^i$, which are the *strategies* (or, equivalently, decision rules) available to $\mathcal{R}^i$ at stage k. The combined mapping $\gamma^i = \{\gamma_1^i, \ldots, \gamma_K^i\}$ is a strategy for $\mathcal{R}^i$ in the game, and the class Γ^i of all such mappings γ^i is the *strategy space* of $\mathcal{R}^i$. A simultaneous specification of the strategy for each $\mathcal{R}^i$ is called the *game strategy*, and denoted by γ. The space of game strategies is denoted by $\Gamma = \Gamma^1 \times \cdots \times \Gamma^N$.
- A real-valued functional $Q^i : (\mathcal{X} \times U_0^1 \times \ldots \times U_0^N) \times (\mathcal{X} \times U_1^1 \times \ldots \times U_1^N) \times \ldots \times (\mathcal{X} \times U_K^1 \times \ldots \times U_K^N) \to \mathbb{R}$ (where $\mathbb{R}$ denotes the space of real numbers) is defined for each $i \in \mathbf{N}$, and called *cost functional* (or *loss functional*) of $\mathcal{R}^i$ in the game of fixed duration.

The preceding definition presents an extensive form description of a dynamic game, since the evolution of the game, the information gains and exchanges of the robots throughout the decision process, and the interactions among the robots are explicitly displayed in such a formulation.

The general task is to determine strategies that optimize the cost functionals in some appropriate sense. The selection of a game strategy (robot control strategy) depends on the amount of cooperation that occurs between robots. For any given N-tuple of strategies $\{\gamma^i \in \Gamma^i;\ i \in \mathbf{N}\}$ the motion strategies (control laws) of the robots are determined by the relations

$$u^i = \gamma^i(\eta^i), \qquad i \in \mathbf{N} \tag{3}$$

where η^i denotes the information set of $\mathcal{R}^i$. A several possible information structures can be distinguished [1] (e.g. open-loop pattern, closed-loop patterns, feedback patterns, etc.). We are interested in closed-loop patterns, particularly in the state feedback pattern i.e. $\eta_k^i = \{x_k\}, k \in \mathbf{K}$. We assume that for each k, the state x_k is available so that admissible control strategies are of the form

$$u_k^i = \gamma_k^i(x_k), \qquad i \in \mathbf{N} \tag{4}$$

We will select the specific robot motion strategies depends on the amount of cooperation that occurs among robots.

3 Coordination strategies

Let $Q^i(\gamma)$ denotes the cost associated with γ, to bring the robot from given initial position $x_{initial}$ to a final position x_{final}. In general, there may be many strategies in Γ that produce equivalent costs. The natural partial ordering, $\preccurlyeq$, can be defined on the space of game strategies, Γ. For two strategies $\gamma, \gamma' \in \Gamma$ we say that

$$Q^i(\gamma) \preccurlyeq Q^i(\gamma') \text{ iff } Q^i(\gamma) \leqslant Q^i(\gamma')\ \forall i\ \in \mathbf{N}$$

The *minimal game strategies* with respect to $\preccurlyeq$ are better or equal to all other strategies in Γ [8]. The main goal is to find the strategy γ which minimize $Q^i(\gamma), i \in \mathbf{N}$. Very often some additional constraints should be taken into account, and generally those constraints can expressed as $\mathbf{g}(\mathbf{x}) \leq \mathbf{0}$.

To simplify notation we restrict our further considerations to the case of the system consisting of two robots.

3.1 Robots have independent goals

Suppose that each mobile robot $\mathcal{R}^i, i = 1, 2$ has individual goal and cost functional, Q^i, to be optimized. Each robot is interested in local results and any cooperation is allowed. Without cooperation the robot chooses actions that conflict with the other robot. In a deterministic case this problem can be stated as follows

$$\min_{u^i \in U^i} Q^i(x^i, u^1, u^2),\ i = 1, 2 \tag{5}$$

This is a deterministic game for which the simplest "solution" is a *Nash equilibrium* [1] i.e., a pair (u^{*1}, u^{*2}) satisfying

$$Q^1(u^{*1}, u^{*2}) \leq Q^1(u^1, u^{*2}), \forall u^1 \in U^1 \tag{6}$$

$$Q^2(u^{*1}, u^{*2}) \leq Q^2(u^{*1}, u^2), \forall u^2 \in U^2 \tag{7}$$

where $(u^{*1}, u^{*2}) \in U^1 \times U^2$ is obtained as

$$u^{*1} = \arg \min_{u^1 \in U^1} Q^1(u^1, u^{*2}) \tag{8}$$

$$u^{*2} = \arg \min_{u^2 \in U^2} Q^2(u^{*1}, u^2) \tag{9}$$

It should be noted, that generally, for noncooperative dynamic games there may exist an uncountable number of *informationally non-unique Nash equilibria.* However, here, we deal with a more restrictive class of Nash equilibrium solutions – the so called *feedback Nash equilibrium* – which is devoid of any non-uniqueness.

Another way to eliminate informational non-uniqueness is to formulate the dynamic game in a stochastic framework. Under an appropriate stochastic formulation, every strategy has a unique representation [1]. In this case we introduce an additional player, *nature*, and an additional set Θ, which denotes the *action (control) space* of the additional player, at stage $k \in \mathbf{K}$. Any admissible action θ_k of nature at stage k is an element of Θ. For example, future positions of the mobile robot are not completely predictable due to slipping of the wheels. Thus, the game-theoretic interpretation of this situation is that a nature player interferes with motion commands.

Example 1: In the situation, in which robots have independent goals, a typical task is to simultaneously move each robot $\mathcal{R}^i$ from a given initial state $x^i_{initial} \in \mathcal{X}^i$ to some final state $x^i_{final} \in \mathcal{X}^i$ while avoiding collisions with other robots and obstacles. In this case the *discrete state trajectory* for an individual robot is represented as $x^i : [0, K+1] \to \mathcal{X}^i$. To evaluate performance of each robot in such a task, a cost functional can be of the form

$$Q^i(\gamma^i) = \sum_{k=1}^{K} \left[q^i_k(x^i_k, u^i_k) \sum_{j \neq i} c^{ij}_k(x) \right] + p^i(x^i_{K+1}) \tag{10}$$

in which q_k^i represents a cost function, which is a standard term used in a discrete-time dynamic optimization problems. Penalty term c_k^{ij} represents the interactions between the robots, and function p_k^i represents the goal in terms of performance [8].

The most extreme case when the robots have totally opposite goals i.e. $Q^1 \equiv -Q^2$, can be described as a zero-sum dynamic game. Two robots (or robot teams) playing football might be given as the typical example of such a game Fig. 2.

Figure 2. Robots playing football [6].

3.2 Loosely cooperating robots

A situation in which some subsets of robots or even all robots wish to act in unison but they have different cost functionals can be described as a *multiobjective optimization problem* or *cooperative game problem*. In this case the strategies, known as a *Pareto optimal solution* [10], stands out as a reasonable equilibrium solution, since it features the property that no other joint decision of the robots can improve the performance of at least one of them, without degrading the performance of the other. In a deterministic case the action coordination problem of the two "loosely cooperating" can be written as follows

$$\min_{u^i \in U^i} Q^i(x^1, x^2, u^1, u^2),\ i = 1, 2 \tag{11}$$

There are many methods to find Pareto optimal solution [4, 10].

Example 2: One of the possibly tasks in which robot cooperation can increase execution time is foraging. Foraging consists of searching the environment for objects (referred to as attractors) and carrying them back to central location. Mobile robots performing this task would potentially be suitable for garbage collection or specimen collection in a hazardous environment.

3.3 Tightly cooperating robots

If the robots are working to accomplish a single task and there is a common cost functional $Q \equiv Q^1 \equiv Q^2$ we have a *team problem*

$$\min_{u^i \in U^i} Q(x^1, x^2, u^1, u^2),\ i = 1, 2 \tag{12}$$

The robots have no local goals and they try to optimize the same cost functional. The resulting solution (u^{*1}, u^{*2}) is known as the *team-optimal solution*, and it requires satisfaction of a single inequality

$$Q(u^{*1}, u^{*2}) \leq Q(x^1, x^2, u^1, u^2), \forall\, u^1 \in U^1, u^2 \in U^2 \tag{13}$$

In fact, deterministic team problems are not different from optimal control, since all equilibrium solutions are different representations of the same N-tuple of strategies which is associated with the global minimum of a single objective functional.

Let us now consider the other extreme case – two robots act as a single mechanism (a composite robot). Of course, the robots have a common goal and one performance objective.

Example 3: The system consists of two robot manipulators manipulating a common object shown in Fig.3.

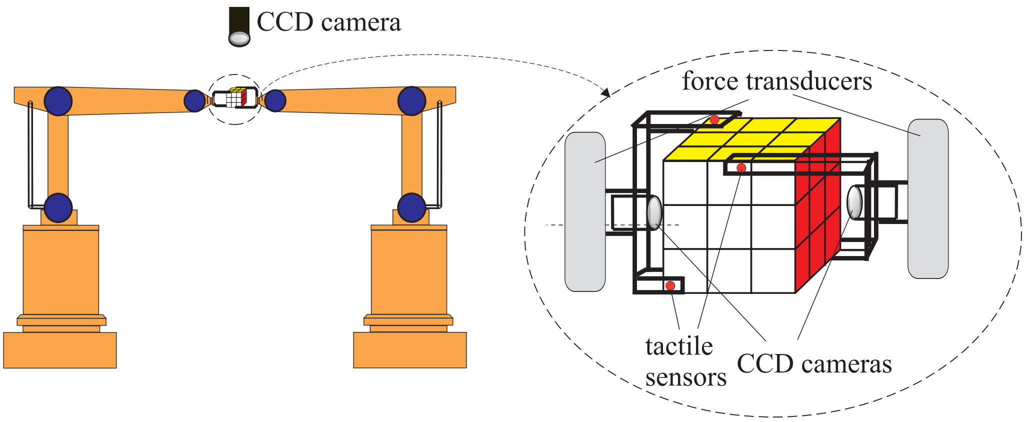

Figure 3. Two robot manipulators manipulating a common object.

The goal is to find the admissible trajectories of two cooperating robot arms in order to manipulate a common movable object [14]. This problem can be expressed as an optimal control problem [15]

$$\min_{(\mathbf{u}^1, \mathbf{u}^2) \in U^1 \times U^2} Q(\mathbf{x}^1, \mathbf{x}^2, \mathbf{u}^1, \mathbf{u}^2), \tag{14}$$

subject to loop closure constraint (to preserve closed kinematic chain)

$$\mathbf{h}(\mathbf{x}^1, \mathbf{x}^2) = \mathbf{0} \tag{15}$$

and some additional constraints

$$\mathbf{g}(\mathbf{x}^1, \mathbf{x}^2) \leq \mathbf{0} \tag{16}$$

which are mechanical constraints due to mobility limitations of the joints, and constraints caused by workspace occupancy conflicts between the object and the robot arms. Now, $\mathbf{x}^1, \mathbf{x}^2$ and $\mathbf{u}^1, \mathbf{u}^2$ are vectors, whose dimensions are equal to degrees of freedom of the robot manipulators. Of course, in general, due to high nonlinearities this optimization problem cannot be solved analytically. To obtain a numerical solution, the infinite-dimensional optimization problem can be approximated by the finite-dimensional problem in mathematical programming [15].

4 CONCLUSION

A problem of coordination of multiple robots working in a shared environment has been described. The relationship between multiple robots coordination problems and game-theoretic issues has been showed. A dynamic game approach can provide a formulation of the optimal multi-robot control strategies that can be obtained. That framework can be served as a unified structure facilitates the comparison of different algorithms. It is obvious that, in some real-life situations, the approach based on game theory cannot be expected to directly yield a numerical solution, particularly if the dimension of the state space is very high (e.g. a lot of robots with many degrees of freedom). As the models of the multiple-robot system become more realistic the solution of the mathematical problem (i.e. corresponding dynamic game) involved becomes very difficult, time consuming and computationally expensive. One possible way to overcome this problem is based on combining DAI techniques (e.g. multi-agent models and algorithms) with concepts of Dynamic Game Theory. DAI techniques provide a way of solving real-life problems for which no more direct approach is available.

Bibliography

[1] T. Başar and G.J. Olsder. *Dynamic Noncooperative Game Theory, 2nd Ed.* Academic Press, London, 1995.

[2] T. Balch and R.C. Arkin. Behavior-based formation control for multirobot teams. *IEEE Trans. on Robotics and Automation*, 14(6):926–939, 1998.

[3] M. Bowling and M. Veloso. Existence of multiagent equilibria with limited agents. *Journal of Artificial Intelligence Research*, 22:353–384, 2004.

[4] W. Stadler (ed.). *Multicriteria Optimization in Engineering and in the Sciences.* Plenum Press, New York and London, 1988.

[5] R. Emery-Montemerlo, G. Gordon, J. Schneider, and S. Thrun. Game theoretic control for robot teams. In *Proc. of the IEEE International Conference on Robotics and Automation*, pages 1175–1181. 2005.

[6] The RoboCup Federation. htpp://robocup.org/.

[7] S.M. LaValle. Robot motion planning: A game-theoretic foundation. *Algorithmica*, 26(3):430–465, 2000.

[8] S.M. LaValle and S.A. Hutchison. Optimal motion planning for multiple robots having independent goals. *IEEE Trans. on Robotics and Automation*, 14(6):912–925, 1998.

[9] Q. Li and S. Payandeh. Multi-agent cooperative manipulation with uncertainty: A neural net-based game theoretic approach. In *Proc. of the IEEE International Conference on Robotics and Automation*, pages 3607–3612. 2003.

[10] G. Owen. *Game Theory, 2nd Ed.* Academic Press, New York, NY, 1982.

[11] L.E. Parker. Alliance: An architecture for fault tolerant multirobot cooperation. *IEEE Trans. on Robotics and Automation*, 14(2):220–240, 1998.

[12] K. Skrzypczyk. Control of a team of mobile robots based on non-cooperative equilibria with partial coordination. *Int. Journal of Applied Mathematics and Computer Sciences*, 15(1):89–97, 2005.

[13] W. Szynkiewicz. Game-theoretic approach to multi-robot motion planning and control. In *Proc. of the 6th IEEE Int. Symposium on Methods and Models in Automation and Robotics, Międzyzdroje, Poland*, pages 503–510, 2000.

[14] W. Szynkiewicz. Motion planning for multi-robot systems with closed kinematic chains. In *Proc. of the 9th IEEE Int. Conf. on Methods and Models in Automation and Robotics, Międzyzdroje, Poland*, pages 779–786, Aug. 25-28 2003.

[15] W. Szynkiewicz. Optimization-based approach to dual-arm manipulation planning. In *Proc. of the 8th National Conf. on Evolutionary Computation and Global Optimization*, pages 235–242, May 30–June 1 2005.

[16] M. Tambe. Towards flexible teamwork. *Journal of Artificial Intelligence Research*, (7):83–124, 1997.