Coalition Formation Problem: New Multi-Agent Methods with Preference Models

Samir Aknine[1], Suzanne Pinson[2] and Melvin F. Shakun[3]

[1]LIP6, Université Paris 6
8, rue du Capitaine Scott
75015 PARIS Cedex 15, France
Samir.Aknine@lip6.fr

[2]LAMSADE, Université Paris Dauphine
Place du Maréchal De Lattre De Tassigny
75775 Paris Cedex 16, France

[3]Leonard, N. Stern School of Business,
New York University
44, West 4th Street, New York, NY 10012-1126, USA

From: AAAI Technical Report WS-02-04. Compilation copyright © 2002, AAAI (www.aaai.org). All rights reserved.

Abstract

Current multi-agent coalition formation methods present two major problems. First, some of these methods can be applied only to cooperative multi-agent systems; second, the algorithms proposed may fail in the formation of the coalitions. This article proposes two methods for problems of agent coalition formation in cooperative and non-cooperative multi-agent systems. These methods are based on agent preference models and on preference aggregation using the Choquet integral.

1 Introduction

Coalition formation is one of the fundamental research issues in distributed artificial intelligence and multi-agent systems. Current methods present two major limitations. First, some of these methods can be applied only to cooperative multi-agent systems; second, the algorithms proposed do not necessarily guarantee the formation of coalitions in all coordination situations [Aknine 00]. This article presents two methods for the coalition formation problem in cooperative and non- cooperative multi-agent systems in which agents need to come together in order to complete their tasks. We consider that in a cooperative multi-agent system the agents exchange information. The agents can also agree to carry out the tasks of the other agents without asking for compensation. In a competitive multi-agent system, information is not exchanged among agents or exchange is limited. Here agents maximize their own preference function, not of the multi-agent system. In order to take into account the interactions between criteria and the dependencies between agents, we have chosen the Choquet integral for preference aggregation. One of the methods based on the ESD (Evolutionary System Design) methodology for restructuring the multi-agent problem [Shakun 96] is suggested for competitive multi-agent systems. The other method is suggested for cooperative multi-agent systems.

This article is structured as follows. In section 2, previous work is analyzed in detail and the principal limitations of the existing solutions are examined. Section 3 formalizes the problem of coalition formation following our approach. Section 4 presents our methods for solving this problem, and gives the results of the experiments done using these methods. Section 5 Briefly summarizes the research.

2 Different approaches for solving the problem of coalition formation

Much research work has been done on multi-agent coalition formation (Ketchpel 94; Sandholm and Lesser 97; Shehory et al. 97, 98; Zlotkin and Rosenschein 96). Sandholm and Lesser define the process of coalition formation by a distribution of the agents in exhaustive and disjoined coalitions [Sandholm et al. 97]. Shehory and Kraus have extended this definition by allowing the formation of overlapping coalitions (i.e. an agent can belong to several coalitions at the same time) [Shehory et al. 98]. The analysis of the algorithms (Ketchpel 94; Klusch and Shehory 96; Sandholm et al. 97, 99; Shehory et al. 97) shows that these approaches do not necessarily guarantee the formation of the coalitions in all cases. All these methods have been discussed in (Aknine 99; Aknine 00).

In his work, Ketchpel has proposed two algorithms for coalition formation. Ketchpel's first algorithm [Ketchpel 94] is centralized. As for his second one, it does not solve the problem of coalition formation in all the cases. Zlotkin and Rosenschein have proposed a mechanism for coalition formation that uses cryptography techniques for sub-additive task-oriented domains. This mechanism is based on a Shapley value which is the expected utility that each agent will have from such a random process [Zlotkin et al. 96]. However this mechanism can only be applied to small sized multi-agent systems because of its combinatorial complexity due to the calculation of all possible coalitions.

The coalition formation method proposed by Shehory and Kraus [Shehory et al. 98] can only be applied for cooperative multi-agent systems in which the agents are able to exchange their information. In order to reduce the complexity of their algorithm, Shehory and Kraus authorized the agents to carry out certain collective and shared processing by exchanging information. This model cannot be applied to non-cooperative multi-agent systems in which the agents cannot exchange their information. The algorithm of Sandholm et al. [Sandholm et al. 99] extends that one. However, the principal problem of the multi-agent approach, i.e. distribution of centralized algorithm control on several autonomous agents, has not been addressed in this method.

It is worth noting that the work of [Derk et al. 92] on game theory was the first to propose coalition formation methods based on building a lattice of coalition structures. Even if this work does not focus on algorithmic aspects, it provides an analysis and a thorough formalization of the coalition formation problem. The authors were mainly concerned with the properties of the structures of coalitions, the evolution of these structures, the stability of the coalitions and the utility of the agents in these coalitions.

In more recent work, [Tsvetovat et al. 00] proposed an algorithm based on the principle of electing a leader for coalition formation. This algorithm has been applied to electronic commerce processes. This approach is similar to the one proposed in [Aknine 99]. Lerman et al. have proposed an alternative, physics-motivated mechanism for coalition formation that treats agents as randomly moving, locally interacting entities [Lerman et al. 00]. They consider that a new coalition may form when two agents meet randomly, and it may grow when a single agent randomly meets the coalition. The aim of this work was to define a mathematical model, formalized as a series of differential equations. These equations have steady state solutions that describe the equilibrium distribution of coalitions. But the authors have not given any details of the autonomous agent behaviors and how they concretely use this mathematical model. The algorithmic specifications have not been proposed and the convergence of this model has not been addressed.

3 Formal Description of our approach for multi-agent coalition formation

3.1 Motivations

All the existing coalition formation models for multi-agent systems are based on the assumption that the agents seek to maximize a global utility function which is directly defined in the agents [Shehory et al. 98]. They will then distribute to each other the profits generated by the execution of their collective tasks. This involves both strong complexity of the algorithms, coordination and negotiation problems around the social utility functions and final distribution of the profits. Our decentralized approach for solving the problem of coalition formation is motivated by the decentralized behaviors of agents and by economic reality: usually the problem involves too many agents.

Our approach considers the preferences of the agents and not a global utility function. We do not try to satisfy a global utility function, which is difficult to build, but we consider the individual points of view of the agents represented by their preference models. We have chosen the preference model for several reasons: (1) a preference model is easier to build than a global utility function; (2) the preferences can be calculated according to various criteria; (3) when a coalition is formed, the agents can consider the aggregated preferences of the coalition.

We propose two types of protocols and coalition formation methods according to whether or not there is sharing of agent preferences. Sharing the preferences allows a more effective solution in a cooperative situation. On the contrary, not sharing the preferences makes it possible to solve non-cooperative situations. The algorithms that we propose can form coalitions in all cases.

To illustrate our approach and the set of concepts that will be defined in this article, let us consider the example of collector agents in a computer manufacturing center. This type of agent intervenes in the tasks of loading and storing hardware material resulting from a manufacturing process. On each computer is an indication of its weight. Some computers cannot be carried out by a single agent, given its limited physical capacity, which means that several agents must necessarily join to complete this task, hence the necessity of a dynamic agent coalition formation. We suppose that there is no hierarchy among the agents, i.e. there is no central coordinator of the agent society.

3.2 Formalizing the multi-agent coalition formation problem

In this section, we present some notations and definitions of the concepts in order to help understand our methods. They underly the distributed process of coalition formation detailed in section 4.

Let S be a multi-agent system, $A = \{a_1, a_2, .., a_n\}_{|n \in N}$ be the set of agents, $T=\{t_1, t_2, .., t_p\}_{|p \in N}$ be the set of tasks in S.

3.2.1 Definition of agent preferences

We consider that a necessary condition for the integration of an agent in a coalition is that it is accepted by the members of the coalition and that it wishes to join the coalition. We formalize this desire by the preferences of the agent for the members of the coalition. The agent preference model may be mono-dimensional or multi-dimensional (Roubens et al. 85 ; Chandon et al. 81).

In some cases, using one criterion to build an agent preference model is not sufficient to assign the agent to a coalition. Our algorithm automatically restructures the problem so as to use other criteria [Shakun, 96]. Consequently, an agent can transform its mono-dimensional preference model defined with one criterion into a multi-dimensional preference model.

Definition 1.1 (*Multi-dimensional agent preference model*). A multi-dimensional preference model of an agent a_i on a set A is a matrix denoted $\prod_{ai}^{p}$ whose elements are the preferences of agent a_i for each agent $a_j \in A$ regarding criteria $D_{k|k=1..p}$ (also called dimensions), and whose i^{th} column corresponding to agent a_i is equal to zero.

Each line of this matrix concerns a specific criterion defined by the designer.

To build the aggregated preference model of an agent based on its multi-dimensional preference model, the weighted sum is often used as an aggregation operator of the preferences but it has several limitations [Grabisch 95]. To overcome these limitations, we propose another aggregation operator: the Choquet integral. It was defined by Grabisch [Grabisch 95]. The Choquet integral is used to:

- aggregate the preferences of a single agent for another agent regarding several criteria;
- aggregate the preferences of a single agent for several agents and to obtain its preference for a coalition;

- aggregate the preferences of several agents for one agent.

Using the Choquet integral as a preference aggregation operator, an agent computes its preference vector $\prod_{ai} = (x_{ij})$ as follows:

$$x_{ij} = \sum_{1 \le k \le p} (y^k_{ij} - y^{k+1}_{ij})\, \mu(E_k) = C_\mu (y^1_{ij}, \ldots, y^p_{ij})$$

$$\forall a_i, a_j \in A, \forall y^k_{ij} \in \prod_{ai}^{p} \quad (i, j = 1, \ldots, n)\ (k = 1, \ldots, p) \text{ with } y^{p+1}_{ij} = 0, y^k_{i,i} = 0$$

E_k is a combination of criteria, $E_k \subseteq D$, such that $|E_1| = 1, \ldots, |E_p| = p$. $\mu(E_k)$ is the weight of a subset of criteria E_k given by the designer. It can be computed with a utility function μ which respects to the constraints specified by the Choquet integral [Grabisch 95].

This definition means that while computing the aggregated preference model of an agent according to a set of criteria, the formula progressively takes into account these criteria and their dependence relation. It starts with one criterion, then takes two criteria, etc. until all the criteria have been considered.

D represents the set of criteria $(D_i)_{i=1..p}$, $\mu(\emptyset) = 0$ and $\mu(D_1, \ldots, D_p) = 1$ means that applying μ on all the criteria gives 1.

In order to understand the rest of our method, it is important to distinguish between the following concepts: the preference of an agent, the collective weight of a set of agents and the individual weight of an agent. The individual weight of an agent is a general value that the designer of the multi-agent system grants to each agent, i.e. independently of the application on which the agent is working or will work. It is a value which is computed according to the features of the agent and the features of the other agents sharing its environment. In our application, for instance, the weights of the agents are defined according to their size, physical weight, etc, without taking into account the execution state of the application. A collective weight is given by the designer for a set of agents. This value is computed according to the features of the set of agents. The fact that two agents a_i and a_j, taken collectively, can perform tasks that they would not be able to perform separately for some reason such as the physical constraints of the agent, implies that the designer defines a higher collective weight for the two agents than their respective individual weights. The only preference aggregation operator which is able to take into account these collective and individual weights is the Choquet integral. This is why we have used it in order to compute the different parameters of the coalitions.

3.2.4 Definitions of some coalition formation concepts

All agents in a coalition are represented by a global preference model. This model aggregates the individual preference models of each agent of the coalition according to their respective weights.

Definition 3.1 (*Coalition preference model*). Let $\prod_i$ be the preference model of coalition C_i composed of m agents. The preference of C_i for agent a_k, denoted $\prod_{ik}$, is such that:

$$\prod_{ik} = \sum_{1 \le j \le m} (x_{j,k} - x_{j+1,k})\, \mu(A_j) = C_\mu (x_{1k}, \ldots, x_{mk})$$

$\mu(A_j)$ is either an individual weight of agent a_j or a collective weight of a subset of agents $A_j \subseteq C_i$.

To compute the preference of a coalition C_i for an agent a_k, the agents of this coalition order their preferences for a_k and build a decreasing sequence. This formula indicates that the preference of the coalition for an agent a_k is equal to the Choquet integral of the individual preferences of each agent in this coalition for agent a_k.

To formalize the desire of an agent to join a coalition if it is not a member of it and its desire to remain in the coalition if it is a member, we define the preference of an agent for a coalition

Definition 3.2 (*An agent preference for a coalition*). An agent preference for a coalition is a global evaluation made by the agent of all the members of this coalition using the individual preferences associated with each agent of this coalition. Formally:

Let $\prod_{ak}$ be the preference model of an agent a_k. The preference of a_k for a coalition C_i composed of m agents, denoted $\Re(a_k, C_i)$, is such that:

$$\Re(a_k, C_i) = \sum_{1 \le j \le m} (x_{kj} - x_{kj+1})\ \mu(A_j) = C_\mu (x_{k1}, \ldots, y_{km})$$

$\mu(A_j)$ represents the individual weight of agent a_j or the collective weight of a subset of agents in $A_j \subseteq C_i$. To compute the preference of agent a_k for a coalition C_i, a_k orders its preferences for each member of the coalition C_i and builds a decreasing sequence. This formula indicates that the preference of agent a_k for a coalition C_i is computed using the Choquet integral on its individual preferences for each agent in this coalition C_i.

When agents try to form coalitions, each agent looks in the subset of possible partners for those which it prefers and but also which prefer it. In our approach, this concept is called *bilateral attraction* of the agents and is described in detail below. In order to start the coalition formation process, we have introduced the concept of unilateral preference attraction to identify the preferred agent from among the other agents. This unilateral preference attraction can be used to solve the coalition formation problem and is defined below.

Definition 3.3 (*Agent unilateral preference attraction*). The unilateral preference attraction defines the force which enables an agent to convince the agents which are not yet in coalitions to choose it as a partner. Formally:
Let a_i be an agent of S. The unilateral preference attraction of a_i, defined on the set S of n agents, denoted U-Attraction (a_i), is such that:

$$\text{U-Attraction}(a_i) = \sum_{j=1(j \neq k)}^{n} (x_{j,i} - x_{j+1,i})\, \mu(A_j) = C_\mu(x_{1i}, \ldots, x_{ni})$$

$\mu(A_j)$ represents an individual weight of agent a_j or a collective weight of a subset of agents in $A_j \subseteq \{a_1, \ldots, a_n\}$.

The process used to form coalitions also considers the desire of an agent to join the coalition which attracts it. This is explained in the following definition of agent bilateral preference attraction.

Definition 3.4 (*Agent bilateral preference attraction*). Let a_i be an agent of S. The bilateral preference attraction of a_i for a coalition C_k, denoted B-Attraction (a_i, C_k), is such that:

$$\text{B-Attraction}(a_i, C_k) = \prod_{ki} \times \Re(a_i, C_k)$$

Having presented some fundamental definitions, we can now present our methods for agent coalition formation.

4 Our coalition formation methods

In this section we present two methods of coalition formation. The first method, based on ESD (Evolutionary System Design) methodology defined by M. Shakun [Shakun 96], is useful for competitive multi-agent systems (where the agents cannot exchange their knowledge), whereas our second method is powerful for cooperative multi-agent systems (where the agents can exchange their knowledge).

4.1 Presentation of two coalition formation methods

4.1.1 ESD coalition formation method for competitive multi-agent systems

When experimenting the well-known coalition formation methods, we have observed that these methods do not guarantee the formation of coalitions in all cases (cf. section 2). We have therefore a new method based on the Evolutionary System Design methodology, ESD [Shakun 96]. ESD is a general formal modeling/design framework for multi-agent problem solving and negotiation that can be applied to define and solve specific problems (see [Shakun 96] for more details). When negotiation solutions are not forthcoming, problem restructuring is a key approach. ESD considers that to enable the agents to reach an agreement during a negotiation process when they fail, we must allow them to restructure the problem by defining new negotiation criteria for the agents. If the agents conclude that it is impossible to form coalitions with the agent preference models based on only one criterion (for example, the time of arrival of the agents at the site of task execution), they rebuild their preference models by gradually introducing new negotiation criteria. In our approach, we assume that the designer of a multi-agent system defines a priori several negotiation criteria as well as their weights. The dimensions are introduced lexicographically. In our application, the first criterion is the arrival time on the site where the task should be executed. The second one is the remaining energy (battery life) available in an agent (robot) before it goes flat. The new agent preference model built on these two criteria is obtained by aggregating the two mono-dimensional models, i.e. the time model and the energy model according to formula given in sections 3.2.1 and 3.2.2. Additional criteria can be used.

The algorithm of the behavior of an agent in the coalition formation process uses communication primitives and operators whose semantics are presented in [Aknine 00]. It is based on the principle of exchanging offers with the other agents. When the coalition formation algorithm is activated, if an agent has not yet obtained the result it is expecting, the problem is restructured. To do so, the agents use all the criteria defined for the problem to be solved. Each time a new preference model is computed, the agent reactivates the algorithm with the new preference values.
The advantage of this method is that an agent does not need to know the preferences of the other agents in the system in order to form coalitions. This is better because it saves time and it guarantees the confidentiality of the knowledge.

4.1.2 Coalition formation method for cooperative multi-agent systems

We have defined a second method for cooperative multi-agent systems in which the agents can exchange their information, the Coalition Formation Algorithm with Shared Preference Models, that guarantees the formation of coalitions in all situations. It is based on the principle of circulating the coalition formation procedure among the different agents of the system. The agents coordinate their activities so that only one agent may form its coalition at a time. For each agent, we have defined a unilateral preference attraction which evaluates the preference of all the other agents for this agent (cf. definition 3.3). By choosing the most preferred agent to start its coalition, we guarantee that the preferences of all the agents will be respected. An agent starts the coalition formation process if it has the greatest unilateral preference attraction. The process of coalition formation passes from the most preferred agent to the less preferred ones as long as there are agents that are not yet in coalitions.

The algorithm that we propose is made up of two phases. During the first phase, each agent a_i builds its preference model as defined in section 3.2. It sends its model successively to all the agents of the system and waits to receive their preference models. Each agent compares its preference model with the models of the other agents. If it has the greatest unilateral preference attraction, it forms a coalition with the agents whose preference models are equivalent (cf. definition 3.3) or so as to maximize the bilateral preference attraction of the agents in the coalition. This is measured by the product of preferences (cf. definition 3.4). Then, the agent waits for the first phase of the process to be over and the agent with the next highest unilateral preference attraction can begin forming its coalition. If it has been chosen in a coalition, the agent accepts and finishes the coalition formation process. If the agent has formed its coalition, it waits for confirmation from all of its partners in order to release the coalition formation process.

The analysis of this algorithm shows that its computational complexity is less than the complexity of the other algorithms, in particular, [Shehory et al. 98]. There are at most $[n^2/k]$ necessary rounds for building coalitions of 'k' cardinality with 'n' agents in the system.

4.2 Implementation results

To evaluate the performance of our two methods, we have implemented a multi-agent system. At each run, we have tested it on various sets of agents by increasing the cardinality of the formed coalitions. Several tasks were performed by the agents at each run, each task needing the same number of agents in order to be performed. Therefore, in all experiments all the coalitions have the same cardinality or size at a given time. The results obtained are summarized in figures 1 and 2.

Several experiments have been carried out with a number of agents varying from 20 to 200. The coalitions formed by the agents have cardinalities which vary from 2 to 100. For each set of agents and each coalition cardinality, we have carried out several series of experiments with different preference values at each time. The results presented in the following figures are obtained by calculating the average of the results of all tests.

In contrast with the results using the Shehory and Kraus method, the ESD coalition formation methodology forms coalitions not only whenever the Shehory and Kraus method does, but also in many cases in which the Shehory and Kraus method fails (cf. figure 1). However, the ESD methodology fails to solve some cases, but the number of failures using the ESD methodology is definitely lower than the number of failures using the Shehory and Kraus method. In figure 1, we plot the time for coalition formation in milliseconds (ms) vs. coalition size (cardinality). The absence of a plotted point at a given cardinality indicates a coalition formation failure. We can observe the discontinuity of the curves describing the evolution of coalition formation time according to the number of agents in the system and the cardinality of the coalitions.

For example, using the ESD methodology, the agents start with one criterion. If a coalition of cardinality k (k=2, 3, 4, .., 100) does not form, the agents try to use two criteria. If it still fails, they try successively 3, 4, 5, 6 and 7 criteria and stop introducing criteria when coalitions are formed. For efficiency reasons, we have limited the maximum number of criteria used by the agents to seven. If seven criteria are not enough, the agents consider this coalition formation case as a failure.

Time for coalition formation (ms)

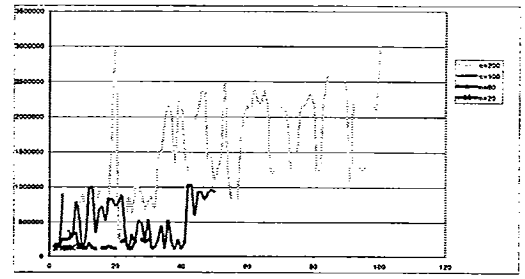

Cardinality

Figure 1. Experimental results of our first distributed coalition formation method using the evolutionary system design (ESD) methodology.

With our second method (Coalition Formation Algorithm with Shared Preference Models), we have observed that the multi-agent system forms coalitions in all cases including the cases in which the other methods fail. The results of this second method (cf. figure 2) show that the time necessary for coalition formation depends on the preference models of the agents. The time is proportional to the number of agents in the system, i.e. higher the number of agents, the higher the negotiation time. As for the cardinality of the coalitions, it influences the time of coalition formation. For instance, in a multi-agent system composed of 200 agent (cf. figure 2, n=200) the agents need 28500ms to form coalitions of cardinality k=20 and 17650ms to form coalitions of cardinality k=80. In our second method, the coalition formation time is less than the time required in both ESD and Shehory and Kraus methods.

We have observed that the time for coalition formation decreases. This result is due to the principle of circulating the coalition formation procedure in this method. It is clear that the higher the cardinality of the coalitions, the more slowly the coalition formation process circulates. Consequently, the length of the coalition formation process of the agents will decrease with cardinality.

As far as the experimental results obtained using the ESD methodology are concerned, the curves do not have the same shape because this methodology is based on the principle of restructuring the negotiation problem. The length of the coalition formation process is proportional to the time that the agents spend to reach a consensus each time they restructure the problem. This explains the independence relation between the cardinality of the coalitions and the global time for coalition formation.

Our second method has the advantage of forming coalitions in all cases. However, as noted before, in this method agents need to know the preference models of the other agents. Thus, our second method is limited to situations where the sharing of agent preference models is acceptable, which is the case in cooperative multi-agent systems. On the contrary, the ESD methodology, which does not require preference sharing, can be applied to non-cooperative multi-agent systems as well as to cooperative multi-agent systems.

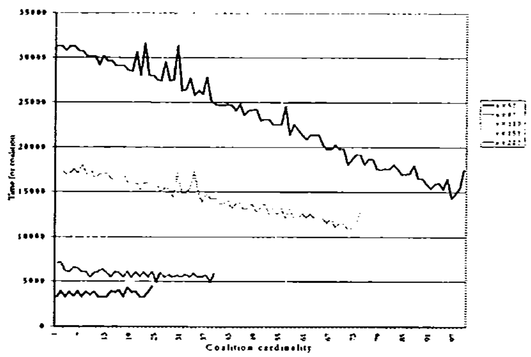

Figure 2. Experimental results of our second method, Coalition Formation Algorithm with Shared Preference Models.

5 Summary

In this article we have developed two new coalition formation methods for multi-agent systems. First, we have shown the existence of several limitations in the existing work. To overcome these practical and theoretical limitations, we have proposed two solutions to the coalition formation problem. Our first method (ESD Coalition Formation) methodology solves the cases in which the well-known existing method [Shehory et al. 98] succeeds in coalition formation; but ESD also solves many other cases where the Shehory and Kraus method fails. Our second method (Coalition Formation Algorithm with Shared Preference Models) solves all coalition formation cases with a shorter coalition formation time than that using the Shehory and Kraus or ESD methods. Each of our two methods offers advantages and limitations. The ESD methodology works for both cooperative and non-cooperative multi-agent systems but does not, however, guarantee the formation of coalitions in all cases. Our second method requires that agents share their preference models. It is powerful for cooperative multi-agent systems and performance is good thanks to the exchanged preference models.

References

Aknine, S., Pinson, S., Shakun, M.F. "Coalition Formation Methods for Multi-Agent Coordination Problems", *Group Decision and Negotiation*, Glasgow, July, 2000.

Aknine, S. Multi-agent Coordination Models and Methods, PhD. Thesis (in French Language), Paris-Dauphine University, France, 2000.

Aknine, S. "An Algorithm for Multi-agent Coalition Formation", *14th IEEE International Symposium on Intelligent Control, Intelligent Systems*, Kokar, M. (ed.), Cambridge, Massachusetts, USA, 142-147, September, 1999.

Chandon, J.L., Vincke, P. "La Modélisation des Preferences", Batteau, P., Jacquet-Lagrèze, E., Monjardet, B. (eds.), Analyse et Agrégation des Preferences, *Economica*, Paris, 1981.

Derk, J.M., Gilles, R.P. "Hierarchical Organization Structures and Constraints on Coalition Formation", *International Journal of Game Theory*, Vol. 95, 147-163, 1995.

Grabisch, M. "Fuzzy Integral in Multi-criteria Decision Making", *Fuzzy Set and Systems*, Vol. 69, 1995.

Ketchpel, S. "Forming Coalitions in the face of uncertain rewards", *Nat. conf. on AI*, AAAI, 1994.

Klusch, M., Shehory, O. "A Polynomial Kernel-Oriented Coalition Formation Algorithm for Rational Information Agents", Tokoro, M. (ed.), *International Conference on Multi-agent Systems*, 1996b.

Lerman, K., Shehory, O. "Coalition Formation for Large-Scale Electronic Markets", *ICMAS*, 2000.

Roubens, M., Vincke, P. Preference Modelling, *Lecture Notes on Economic and Mathematical Systems*, 250, Springer Verlag, 1985.

Sandholm, T., Lesser, V. "Coalition Formation Among Bounded Rational Agents", *IJCAI*, 1995.

Sandholm, T., Lesser, V. "Coalitions among Computationally Bounded Agents", *Artificial Intelligence*, 1997.

Sandholm, T., Larson, K., Andersson, M., Shehory, O., Tohmé, F. "Coalition Structure Generation with Worst Case Guarantees", *AI*, 111(1-2), 1999.

Shakun, M.F. "Modeling and Supporting Task-Oriented Group Processes: Purposeful Complex Adaptive Systems and Evolutionary Systems Design", Shakun, M.F. (ed.), *Group Decision and Negotiation Journal*, 5 (4-6), 1996.

Shehory, O., Kraus, S. "Methods for Task Allocation via Coalition Formation, *Artificial Intelligence*, Elsevier Science, 1998.

Shehory, O., Sycara, K., Jha, S. "Multi-agent Coordination Through Coalition Formation", *Intelligent Agents IV*, ATAL, M. Singh, A. Rao, M. Wooldridge, (eds.), Springer, Lecture Notes in Artificial Intelligence, 1997.

Tsvetovat, M., Sycara, K., Chen, Y., Ying, J. "Customer Coalitions in the Electronic Marketplace", *Int. Conf. On Autonomous Agents*, Barcelona, Spain, 2000.

Zlotkin, G., Rosenschein, J. "Mechanisms for Automated Negotiation in State Oriented Domains", *Journal of Artificial Intelligence Research*, 5, 1996.