

Long-term Coalitions for the Electronic Marketplace

Silvia Breban, Julita Vassileva
University of Saskatchewan
Saskatoon, Canada
{svb308, jiv} @cs.usask.ca

ABSTRACT

This paper presents a new coalition formation mechanism for the electronic marketplace that extends the existing transaction-oriented coalitions to long-term ones based on nurturing customer-vendor relationships. Because trust is an important factor in any form of commerce and it has been an elaborated issue in Electronic Commerce applications in the last few years, we use trust based relationships between agents to model other agents and to help an agent when faced with the decision making problem of joining or leaving a coalition. Microscopic (agent level) description of the proposed coalition formation mechanism based on trust relationships between agents is provided.

Keywords

coalition, trust, multi-agent systems

1. INTRODUCTION

Coalition formation in multi-agent systems is the mechanism of grouping agents that agree to cooperate to execute a task or achieve a goal. The goal can be common to all agents in the group in the case of group or social rationality or it can be specific to each agent in the case of individual rationality.

There are two perspectives on coalition formation: a microscopic and a macroscopic one. The microscopic perspective has as central unit the agent and its reasoning mechanism about joining a coalition. In the case of individual rationality the agent will join a coalition that maximizes its own utility. In the case of existing different other preferences such as group rationality or a long-term utility, however, the decision problem becomes more challenging.

The macroscopic perspective has as fundamental unit the coalition. It describes the system of agents as an entity,

while the microscopic perspective describes the system as composed of small entities – the agents. Research on the macroscopic perspective of coalition formation mechanism studies how agents should be divided into a set of disjoint coalitions existing at a certain moment in time in the system. The partition of agents into disjoint coalitions is called coalition structure.

Several problems arise when trying to define such a mechanism: finding the most profitable coalitions for (all) agents, determining how the joint reward should be divided among them, and choosing a stable division of agents into coalitions.

Coalition formation has been studied in the AI literature from three different approaches: game theory, distributed artificial intelligence (DAI), and, most recently, electronic marketplace. Sometimes aspects of different approaches are found in the same work.

Game theory research focuses typically on the macroscopic perspective, preferred because it uses only few variables that make it easier to control and it lends itself to formalization more easily. The macroscopic perspective directly describes and allows investigating global properties of the system such as the number and the size of coalitions and how these properties change in time. Research from this area provides rigorous formal analysis concerning issues of solution stability [1], fairness, and payoff disbursements and constitutes the basis for all other approaches. However, the approach is centralized and some of the underlying assumptions of the developed algorithms do not necessarily hold in real-world multi-agent systems.

DAI research studies coalition formation mechanisms in the context of group rationality and seeks a proof that these mechanisms are beneficial. DAI approaches aim at improving coordination and cooperation among agents in multi-agent systems by using coalition formation mechanisms. DAI researchers adopted some of the game theoretical concepts and developed upon them distributed coalition formation algorithms to be used at the agent level

within multi-agent systems. These algorithms concentrate on distribution of the computations, complexity reduction, efficient task allocation [2], and communication issues [3].

Recent research brings the coalition formation process into electronic market environments as a mechanism for grouping customer agents with the intent of getting a desired discount from the vendor agent in large size transactions. The definition of the term “coalition” is extended to a group of self-interested agents that are better off as parts of the group than by themselves. The electronic marketplace is an open and free environment in which the agents have individual rationality and no social or group rationality requirements. In this context coalition formation mechanism is proved to be beneficial for both customer and vendor agents. Customer agents know that being part of a coalition allows them to buy large quantities that they might not need individually and this brings them a discount from vendors, so they would want to join a coalition that maximizes this discount. From the vendor’s point of view selling in bulk is also advantageous, since it helps decrease the advertisement, production, and distribution costs. Selling large quantities with a certain discount brings the vendor at least the same profit as retail selling.

Our work concentrates on this direction of self-interested agents that trade goods or knowledge in a large-scale multi-agent system. We use coalitions as a means of grouping agents to improve the cooperation and coordination among them and to increase their individual benefit. The paper aims to present a new approach to coalitions of agents for the electronic marketplace. The explosion of the Internet communities and the increasing number of real life market clubs based on nurturing vendor-customer relationships lead us to the idea of extending the existing transaction - oriented coalitions to long-term coalitions based on trust relationships between agents. We present in the next section related work on coalitions for the electronic marketplace and on a similar concept named congregations. In the third section we continue with an argument for extending temporary coalitions to long-term ones. In section four we present an existing model of subjective trust that an individual has after a history of previous interactions and experiences with the subject of trust. We briefly describe in the fifth section how this model of subjective trust can be used in the proposed coalition formation mechanism to help an agent decide whether to join or leave a coalition. We conclude and present directions and future work in the sixth section.

2. COALITION FORMATION FOR THE ELECTRONIC MARKETPLACE

Sycara and Tsvetovat [4] illustrate the economic incentives behind formation of “buying clubs” and achievement of effect on large-scale economies within temporary agent coalitions. The coalition formation mechanism includes

several stages. The first one is *Negotiation* in which a leader or representative of the coalition negotiates with one or more suppliers to provide the goods or service needed. Then follows a *Coalition Formation* stage, in which the coalition leader solicits new members to join its coalition, based on a set of admission constraints. Next the members elect a coalition leader or cast direct votes for or against certain bids in the *Leader Electing/Voting* stage. *Payment Collection* follows: the coalition leader or third party collects the payment from the coalition members and is responsible for conveying the full amount to the supplier. In the final stage - *Execution/Distribution* - the transaction is executed, the purchased goods arrive, and they are distributed to the members of the coalition.

In real-world environments the formation and administration of coalitions, as well as the distribution of purchased goods to the members are time and resource consuming. A coalition is viable only if the increase in the group’s total utility from wholesale purchases is greater than the cost of creating and running such coalitions.

The general coalition model described above is complex and requires expensive communication among agents that makes it hard to implement or scale up. Improvements have been proposed by Lermann and Shehory [6] that introduce a physics-inspired mechanism for coalition formation. It treats agents as randomly moving, locally interacting entities. The microscopic (agent level) model is simple and tries to minimize the communication between agents: an agent joins a coalition by placing an order to purchase a product; it can leave the coalition by withdrawing the order for the product. This model requires no global knowledge; it can accommodate a large number of agents and still provide a good enough performance in terms of agent benefits and consumption of computational resources. The macroscopic (coalition level) model is expressed mathematically as a set of first-order differential equations that describe how the number of coalitions of different sizes evolves in time. Two cases are considered: one in which agents are not allowed to leave a coalition once they join it and the more realistic case in which detachment is allowed. Introducing even a very small detachment rate allows the system to reach an equilibrium steady state in number of coalitions of different sizes and the increase in the global utility gain is more than twice than in the no-detachment case. The designer can predict the final distribution of coalitions even for large systems. One big contribution of this paper is to illustrate how cooperative behavior (e.g. coalition formation) emerges from interactions among many simple self-interested agents. It is the first approach with no negotiation for forming the coalition and the first approach that allows agents to leave a coalition, behavior that is proved to be beneficial both for the agents (they gain in utility) and for the system (it reaches an equilibrium steady state).

The two models presented above are complex because they involve four or five stages that require expensive communication on behalf of the agents. They address temporary coalitions that last only one transaction. This makes the mechanisms impractical because at each step an agent has to decide what coalition to join without memory of previous experiences. At the microscopic level, searching for suitable coalitions and deciding what coalition to join is time and resource consuming. At the macroscopic level, forming and running new coalitions at each step is also computationally expensive and decreases the stability of the system. The mechanisms lead to high dynamics of the system (i.e. high variance in the number of coalitions existing in the system at a certain moment and high variance in the size of coalitions) that is not desired.

Another limitation of existing coalition formation mechanisms is that they do not address the issue of trust among agents in a coalition. In the context of formal contracts among agents in a group there is a formal trust in the structure and the regulations of the system that needs no explicit specification, but in the context of informal contracts each agent in the group should be able to trust other agents. According to Sycara [4] unless the group is formed by a number of individuals who already know each other there has to be an explicit leader selection/verification mechanism, or a mechanism for collective negotiation. Such mechanisms seem not feasible for open-ended environments, since they involve a lot of interaction and knowledge on behalf of the agents.

A last limitation that we want to highlight is that none of the mechanisms discussed above focuses on the agent's reasoning mechanism about coalition formation in detail. However, a microscopic level description is crucial when developing a practical application that we address.

A related work by Brooks and Durfee [5] presents the notion of long-term groups of self-interested agents under the name of congregations. The authors present congregating both as a metaphor for describing and modeling multi-agent systems and as a means for reducing coordination costs. Agents are expected to have long lifetime during which they take on different roles, perform different tasks, and interact with different agents. Congregating becomes a multi-agent learning problem in which agents have to search for "suitable partners" according to a compatibility criterion not explicitly defined in the paper. Agents are characterized by individual rationality expressed by a long-term utility function, the capability to voluntarily join or leave a congregation, long-term existence of repeated interactions with other agents, and congregation dependent satisfaction. Even under the long lifetime assumption, the congregating problem is reduced to an initial search for the "right" agents to associate with. This should help the agent in future interactions by devoting initial resources and time to find

what congregation to join. In an open environment like the Internet the search problem becomes exponentially difficult as the number of agents increases. To reduce the complexity the authors introduce *labelers* and *congregators*. Labeler agents correspond to producers in a market environment; they have the role to label their congregations so that they can attract congregator agents. Congregator agents are similar to consumers; they can be part of only one congregation at any moment and can interact only with agents from the same congregation. The introduction of labels transforms the problem from one in which each congregator must make a decision as to what congregation to join into a problem (with supposedly lower complexity) where each labeler must decide which label to offer.

The congregations proposed by Brooks and Durfee [5] are similar to coalitions of agents. The different term is used to distinguish them from the task-oriented coalitions existing in DAI research, but it brings religious connotations. Interesting is how labels are used as an advertisement mechanism for producers. In real world markets producers use attractive simple names counting on their resonance and the impact they have on consumers. This is not applicable to agents and the solution brought by this paper is to incorporate advertisement information in the congregation names (labels). One possible problem with this mechanism is that it can lead to long labels. It is not clear how information is extracted from these labels or what can happen if two producers produce exactly the same goods and they try to use the same label. However, assumptions such as long lifetime of agents and repeated interactions as well as the liberty to join or leave a group at any moment are realistic and suitable for an open electronic marketplace that we address in the next section.

3. LONG-TERM COALITIONS

The exponential growth in worth and size of the electronic marketplace in the last years is due to the attractive open environment that the Internet offers to its users. Existing electronic markets provide only limited trading mechanisms: fixed price and auctions, with no negotiation or grouping supported. Customer coalitions try to bring classical formal concepts on coalition formation mechanism into practical real-world environments. Coalition formation is found to be profitable for both customer agents and vendor agents and the numerous Internet communities show that such a behavior has a big potential in creating large-scale economies among similar minded customers.

To overcome limitations presented in the previous section of existing coalition formation mechanisms (i.e. complex and expensive, low stability of the system, no trust addressed, no microscopic description) we propose a model of long-term coalitions that last more transactions and have

long lifetime. The concept is similar to the congregations from [5] in which agents are supposed to have long lifetime of repeated interactions with other agents. What differentiates our groups from congregations is that the reason for being part of a group is not some compatibility criterion, but the discount that members of the same group give to each other. Two other differences are: first, agents are not constrained to interact only with agents in the same group; secondly, finding suitable partners to interact with is not an initial search problem, but a continual evaluation of interactions with other agents. These make our group concept closer to that of a coalition and we will use this term in the future.

A coalition is formed by peer agents that can play both customer and vendor roles; an agent can switch from one role to another depending on the goal it has at the current moment. Agents know that being part of a coalition will bring them discounts from members of the same coalition in future interactions and they have no interdiction to interact with agents outside of their coalition. They can join or leave a coalition at any moment in time, but they can be part of only one coalition at a moment. Agents have individual rationality and try at each moment to maximize their long-term utility function by being part of a coalition with agents from whom they expect to get the largest discount.

Our focus in this work is on the microscopic (agent oriented) level of the coalition formation mechanism, i.e. the reasoning of an individual agent whether to join, leave, or form a coalition. From the perspective of an individual agent, coalition formation can be viewed as a decision problem: at each moment an agent faces a decision of whether to form a new coalition, remain in the same coalition, or leave the current coalition for a better one. The decision should take into account important factors such as the (long-term) goals of the agent, knowledge about others, and global knowledge about the system. The decision should maximize the agent's long-term utility. An agent decides what action to take based on its previous interactions with other agents from the same coalition and outside of it and on its expectations. To model other agents we propose to use the relationships that are established between agents after common experiences. In general, relationships between individuals can reflect different aspects of their interaction: the roles they play in the interaction, the goals they have, the importance that the interaction has for each of them, and the trust they have in one another [11]. In the absence of a formal contract between the agents, the most appropriate aspect of a relationship for coalition formation for the electronic marketplace is the trust that agents have in each other. In the context of formal contracts among agents in a group there is an implicit trust in the structure and the regulations of the system that needs no explicit specification. However,

in the context of informal contracts each agent in the group should be able to trust the other agents.

Trust between agents is an important issue in the Electronic Commerce. We use a trust model inspired from [7] and [8] to represent the mental state of an agent when faced with the problem of deciding whether it has sufficient trust to engage in an action of joining or leaving a coalition and present the model in the next section.

4. TRUST RELATIONSHIPS

The notion of trust has been the object of continuous interest in economics, sociology, and more recently in the AI research. Trust relationships between agents in multi-agent systems are analyzed and modeled in task-delegation problems [9] or Electronic Commerce [7]. Crispo and Christianson [9] analyze the security of the existing delegation mechanisms and describe a new protocol based on trust relationships between client and agent, agent and service provider, and client and service provider. The authors promote the idea that in commercial and financial environments a particular entity of the system should not be trusted a priori, but the principle of the least trust should be applied until otherwise proved. Ganzaroli [7] presents a generic model of trust for Electronic Commerce that extends the single agent perspective to a social one defined in terms of institutional context, moral context, and network structure. The institutional and moral contexts refer to the trust in a functional system, its norms and values, and the ability to be controlled. Electronic Commerce lacks both of them because of the specifics of the Internet as a new technological medium and the moral differences among diverse cultures. One solution to increase the trust in Electronic Commerce is the application of community-based trust models. In an open environment the responsibility to guarantee trust within the community is distributed among all members. Each member of the community is able to verify the trustworthiness of the other members and it is responsible for the trust that it generates in other members.

We address the concept of community-based trust that is established among agents that share common experiences. In general, trust can be seen as a form of absolute reliability in the subject of trust. This form is known as *objective trust* and is similar to the trust that we have in institutions or uniforms. For example, we trust doctors or policemen because of the uniforms they wear and the institutions they represent. The other form of trust, known as *subjective trust*, depends on the moment in time and the experiences accumulated between the truster and the subject of its trust. We use this latter form of trust that presents an evolution over time called the dynamics of trust [8]. Each event that can influence the degree of trust is interpreted by the agent to be either a trust-negative experience or a trust-positive one. In the former case the agent will lose trust to some

degree and in the latter case the agent will gain trust to some degree. The degree to which the trust is changed depends on the trust model used by the agent. This implies that the trusting agent performs a form of continual verification and validation of the subject of trust over time.

A trust representation has several characteristics. It can be described using specific qualitative labels such as “unconditional trust” or “no trust” in a qualitative description or using numbers in a quantitative description. The description, either qualitative or quantitative, has to be stored in a set of trust qualifications. It has to specify a value for the initial trust assigned to an unknown agent and a representation of how trust evolves in time. Another important characteristic of a trust representation is future independency that refers to the fact that trust only depends on past experiences, not on future expectations. A trust representation should reflect a distinguishable past: the time moments when experiences happen is meaningful in the sense that recent experiences are more important than older ones. There are two possible approaches for the design of a trust representation in the agent’s mind. A first approach is to represent sequences of past experiences and to calculate the corresponding trust at each moment. This is space and time consuming. In the other approach the agent does not need to build a representation of the past experiences, but only of trust. A new experience will instantly lead to an update of the trust representation, with no record of the experience itself. Each experience has to be evaluated and classified in the set of predefined experience classes.

We use the quantitative model proposed in [7] as an example for the framework presented above. Given a set of experience classes E and a set of trust qualifications T , we can define a trust function as:

$$\text{trust} : E \times T \rightarrow T$$

$$\text{trust}(e, t) = d * t + (1 - d) * e$$

We consider the case in which $E = [-1, 1]$ meaning that an experience can take any value in the assigned interval and $T = [-1, 1]$ - trust qualifications can have any value bigger than or equal to -1 and less than or equal to 1 . Parameter $d \in \hat{I} [-1, 1]$ is an inflation rate used to model the distinguishable past characteristic of a trust representation. In this trust function after each new experience the existing trust value t is multiplied by d and the impact of the new experience e is added, normalized in such a manner that the result fits in the desired interval $T = [-1, 1]$.

Based on this representation and the set of discrete time values $Time$, a trust evolution function $evol$ can be inductively defined to be used by an agent when it has to update its trust in another agent at each step from $Time$ set:

$$\text{evol} : E \times Time \rightarrow T$$

$$\text{evol}(e, 0) = 0$$

$$\text{evol}(e_0 e_1 \dots e_i, i + 1) = \text{trust}(e_i, \text{evol}(e_0 e_1 \dots e_{i-1}, i))$$

The definition of this evolution function specifies that the initial trust for time step 0 is set to a neutral value 0 . At each time step $i + 1$ the trust is updated based on the previous trust (from time step i) and the current experience e_i according to the trust function defined above.

The formal model presented in this section has all the characteristics required by a trust representation. Several problems arise when trying to integrate it in our coalition formation model. One of them is how to express the current interaction in terms of experience classes E . Another problem is how to use the model of evaluated trust to help the agent make a decision of joining or leaving a coalition. We answer these questions in the next section.

5. PROPOSED COALITION FORMATION MECHANISM

We refer to a system of multiple agents that can play the role of vendor or the role of customer in a transaction. The agents are peers that trade goods in an open electronic market environment. They are helped by a system matchmaker agent to find each other when interested in buying/selling a specific product. Before a transaction between a customer and a vendor is executed, the two agents go through a negotiation phase to agree on a certain price. We use the negotiation protocol described in [10]. Negotiation can end either with a Rejection or with an Agreement between the two agents (and, in the latter case, with the price they agreed on). In the case that the two agents belong to the same coalition a certain discount is applied to the price they agreed upon. We will use the term *interaction* for any attempt to make a transaction between a vendor and a customer agent. If the negotiation phase ends with an Agreement the interaction is *successful*, otherwise we call it *unsatisfactory*. After the interaction has finished, both the customer and the vendor have to evaluate it, considering it a positive experience if the interaction is successful and a negative one if the interaction is unsatisfactory. A positive experience is evaluated taking into consideration the quality of the product and the delivery time with a value from the positive subset of experience classes E^+ (in our example $[0, 1]$). A negative experience is evaluated in the negative subset of experience classes E^- ($[-1, 0]$ in our example). The reason for considering an unsatisfactory interaction as a negative experience and a satisfactory interaction as a positive experience is that an agent’s trust reflects its expectations to have similar interests and money amounts with the agents in the same coalition, because this leads to more interactions and discounts, so more profit.

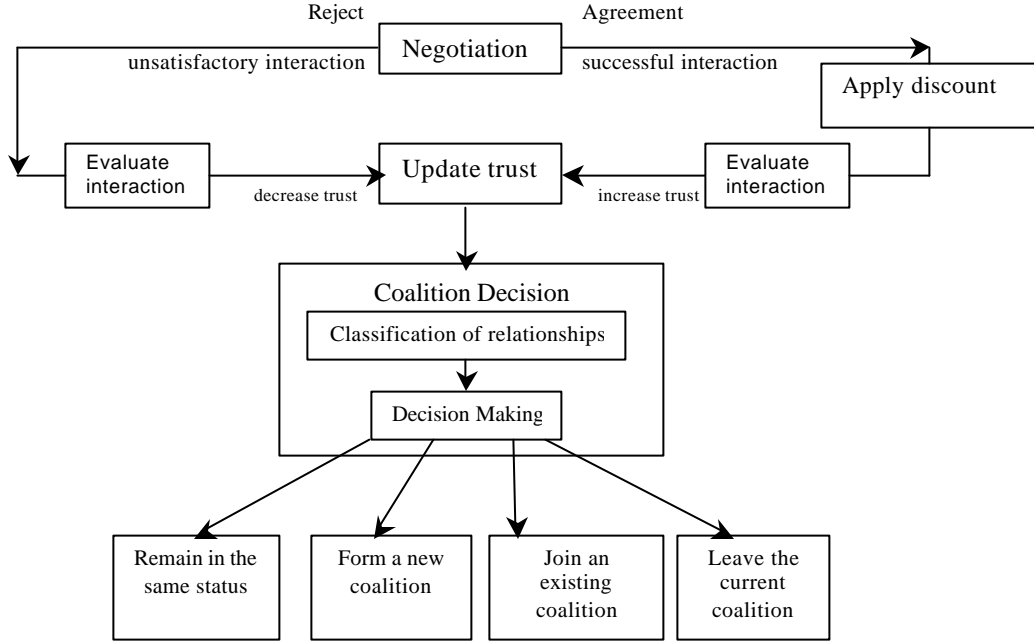


Figure 1 Agent level reasoning description of the proposed coalition formation mechanism

Each agent stores a representation of all its trust relationships with other agents in the system with whom it ever interacted with. A relationship is represented by the name of the agent to be trusted and a specific trust qualification from the set of general trust qualifications T . After an interaction ends and is evaluated, the trust qualification is updated according to the trust evolution function defined in the previous section. Each time the trust representation of an agent is changed, it triggers a Coalition Decision reasoning mechanism as shown in Figure 1.

The Coalition Decision reasoning mechanism consists of two parts: first the agent interprets all its relationships and second it has to make a decision of whether it has sufficient trust to engage in an action of joining or leaving a coalition or it remains in the same status as before. In order to decide what action is most profitable at each moment an agent has to know the coalition it belongs to at the current moment (C), its trust relationships with other agents (T_1, T_2, \dots, T_n), and the coalitions in which these agents are (C_1, C_2, \dots, C_n). The interpretation of an agent's relationships is in fact a classification of the relationships with all other agents it ever interacted with. Classifying trust relationships can be done in two different ways: individually oriented and socially oriented.

In the individually oriented classification the current agent prefers to be in the same coalition with the agent with whom it has the best relationship. First the agent finds an agent that it trusts most. If the two agents are in different coalitions, the agent leaves its own coalition and joins the coalition in which the most trusted agent is. If the most trusted agent is outside of any coalition, the agents form a new coalition.

In the socially oriented classification the current agent prefers the coalition in which it has most summative trust. Given its trust relationships in individual agents, it first calculates the trust it has in each coalition as a sum of the trust relationships in agents from that coalition. Then, it finds a coalition that maximizes this trust. If this coalition is different from the agent's current coalition, it leaves its own coalition and joins the new one.

Rule based algorithms for both individually oriented and socially oriented approaches can be described given the current status of an agent C , the set of trust relationships T_1, T_2, \dots, T_n in agents A_1, A_2, \dots, A_n with whom it interacted with, and the corresponding coalitions of these agents C_1, C_2, \dots, C_n . The current status of the agent (C) and the set of corresponding coalitions C_1, C_2, \dots, C_n take values between 1 and m (the number of existing coalitions in the system at the current moment) or 0 if the agent is outside of coalitions. We consider the coalition to which an agent belongs as global knowledge. Trust relationships T_1, T_2, \dots, T_n have values in the previously defined set of trust qualifications T .

We describe the rule-based algorithm for the individually oriented approach in pseudocode:

```

k = arg max ( Ti )
if ( C == 0 ) then // current agent is outside of coalitions
    if ( Ck != 0 ) then JoinCoalition ( Ck )
    else FormCoalition ( Ak )
else // current agent is in a coalition
    if ( Ck == 0 ) then
        LeaveCoalition ( C )

```

FormCoalition (A_k)
elseif (C_k != C) then
LeaveCoalition (C)
JoinCoalition (C_k)

For the socially oriented approach the rule-based algorithm is similar. The only difference is that first the current agent needs to construct a representation of its personal trust in each coalition as the sum of its trust in agents from the coalition with it had interactions.

6. DIRECTIONS

We reviewed in this paper existing coalition formation mechanisms in the AI literature with emphasis on the most recent research for the electronic marketplace. We argued for the importance of developing trust both at an individual and a community level in the new form of commerce that the Internet offers. An existing formal model of trust based on personal experience of agents is presented and integrated in the proposed coalition formation mechanism. The goal of the paper is to introduce a new model for coalition formation that extends the existing transaction-oriented coalitions to long-term ones by using trust relationships between customers and vendors. The motivation for this approach is to save computational resources and to gain system stability. We expect that most agents will be prone to remain in the same coalition for the next transaction and only a limited number will leave their coalition for better ones.

A limitation of our approach is the restriction imposed to agents to be part of only one coalition at a moment.

Currently a Java-based multi-agent system for the electronic marketplace is being implemented. Agents trade goods using a negotiation mechanism described in [10]. The implementation of the new coalition formation mechanism is under development. Methods for evaluating agents' interactions and benefits are being developed. Experiments for verification and analysis of the coalition structure stability and the improvement that the mechanism brings in the system's performance are planned.

7. REFERENCES

- [1] T. Sandholm, K. Larson, M. Andersson, O. Shehory, F. Tohme *Coalition structure generation with worst-case guarantees* Artificial Intelligence vol.111, no. 1—2, 1999, pp. 209—238
- [2] G. Zlotkin, J.S. Rosenschein *Coalition, cryptography and stability: Mechanisms for coalition formation in task oriented domains* Proceedings of the Twelfth National Conference on Artificial Intelligence, Seattle, July 1994, pp. 432--437
- [3] K. Sycara, K. Decker, A. Pannu, M. Williamson, D. Zeng *Distributed intelligent agents* IEEE Expert, December 1996, pp. 36--46
- [4] M Tsvetovat, K. Sycara *Customer Coalitions in the Electronic Marketplace* Proceedings of Autonomous Agents 2000, Barcelona, pp.263--264
- [5] C.H. Brooks, E.H. Durfee, A. *An Introduction to Congregating in MultiAgent Systems* Proceedings of Fourth International Conference on MultiAgent Systems, Boston, IEEE 2000, pp. 79--86
- [6] K. Lerman, O. Shehory *Coalition Formation for Large Scale Electronic Markets* Proceedings of the Fourth International Conference on Multiagent Systems, Boston, July 2000, pp. 216--222
- [7] A. Ganzaroli, Y.Tan, W. Thoen *The Social and Institutional Context of Trust in Electronic Commerce* Autonomous Agents '99, Workshops on Deception, Fraud and Trust in Agent Societies, Seattle 1999, pp. 65--76
- [8] C. Jonker, J. Treur *Formal Analysis of Models for the Dynamics of Trust based on Experiences* Autonomous Agents '99, Workshop on Deception, Fraud and Trust in Agent Societies, Seattle 1999, pp. 81--94
- [9] B. Crispo, B. Christianson *A Note About the Semantics of Delegation* Autonomous Agents '99, Workshop on Deception, Fraud and Trust in Agent Societies, Seattle 1999, pp. 55--64
- [10] C. Mudgal, J. Vassileva *Bilateral Negotiation with Incomplete and Uncertain Information: A Decision-Theoretic Approach Using a Model of the Opponent* Proceedings of the 4th International Workshop on Cooperative Information Agents (CIA IV), Boston, July 2000, pp. 107—118
- [11] J. Vassileva *Goal-Based Autonomous Social Agents Supporting Adaptation and Teaching in a Distributed Environment* Proceedings of ITS'98, San Antonio, Texas.LNCS No1452, Springer Verlag: Berlin pp.564-573