Long-term Coalitions for the Electronic Marketplace

Silvia Breban
Department of Computer Science
University of Saskatchewan
57 Campus Drive
Saskatoon, SK. S7N 5A9
Canada
email: svb308 @cs.usask.ca

Supervisor: Julita Vassileva

Abstract

This paper presents a detailed overview on coalition formation mechanisms developed in the areas of game theory, distributed artificial intelligence, and electronic commerce. Further it proposes a concept of long-term coalitions for the electronic marketplace based on nurturing customer-vendor relationships. Because trust is an important factor in any form of commercial interactions we use trust-based relationships between agents to model other agents and to help an agent when faced with the decision of joining or leaving a coalition. Detailed microscopic (agent level) description of the proposed coalition formation mechanism based on trust relationships between agents is provided.

1. Introduction

Cooperation and coordination are important issues in agents' behavior in the context of multi-agent systems (MAS). One powerful approach to these issues is to group agents that have similar interests or common goals to improve the behavior of the whole system or of each individual agent. Grouping agents brings a new dimension in a multi-agent system - the social behavior defined according to S. Ossowski [Ossowski99] as consisting of two descriptions: the social structure and the social function. The first is a static description of how individuals are arranged in institutionally controlled or defined relationships and roles. The social function is a dynamic description of the functioning of the structure in time.

There are two different types of groups that appear in the AI literature depending on the context and the problem they have to solve. The first type is based on a formal fixed contract among agents in the group and the second - on an informal flexible contract among agents. Examples for the first type of groups are the teams that appear in the ROBOCUP project or STEAM (A Shell for TEAMwork) model, the groups formed in the ContractNet protocol, and the organizations of virtual enterprises that simulate real ones. These types of groups have to be described in detail at the social structural level. The bonds between agents are created according to their roles in the group and the obligations, permissions, and interdictions imposed by these. This approach imposes rigidity to agents' actions and possible sanctions if they do not obey a predefined social functioning. The second type of groups is based on informal flexible contract among agents that leaves more freedom at the agent level. The social structure is minimized, but the social functioning becomes complex due to uncertainty in the behavior of agents. In this case cooperative group behavior emerges from the local interactions among simple agents. Examples of such groups are coalitions, congregations [Durfee2000], and buying clubs [Sycara2000]. The groups based on formal contracts among agents are useful in the context of group or social rationality where the benefit of the entire system is more important than the benefit of each individual agent, while the informal groups are more suitable in systems formed of self-interested agents. In the last case the social functioning of the system becomes more complex and hard to control and ensuring the stability of the system is an important issue.

Our work concentrates on 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 so that they can individually benefit. The paper presents a new approach to build long-term 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 concept of task oriented (transaction oriented) coalitions to long-term coalitions based on trust relationships between agents. In the second section we present an overview on existing coalitions and the different approaches in the literature: game theory, distributed AI, and electronic commerce. We continue with an argument for extending temporary coalitions to long-term ones in the third section. 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 object of trust. In the fifth section we present guidelines and a brief description of a new coalition formation mechanism that takes into account the modeled trust between agents. We conclude and present directions and future work in the sixth section.

2. Coalition Formation in Multi-Agent Systems

Coalition formation in multi-agent systems is the mechanism of grouping agents that agree to cooperate to execute a task or achieve a (common) goal. There are several problems that 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.

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 making 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 a coalition structure.

In the following subsections we present three different approaches to coalition formation mechanism originating from the areas of game theory, distributed AI, and electronic commerce. These three approaches on coalitions are not entirely separated; sometimes aspects of different approaches are found in the same work.

2.1 Game Theory Research on Coalitions

Game theory research focuses typically on the macroscopic perspective, preferred because it uses only a 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 size of coalitions and how these properties change in time.

Game theory literature such as [Rapoport84] describes which coalitions will form in N-person games under different settings and how the players will distribute the benefits of cooperation among themselves. Another important concern is to prove that the generated coalition structure is stable. Different stability criteria are explored: Kernel stability (k-stability), Nash equilibrium, and Pareto optimality. A coalition structure is k-stable if there is equilibrium between pairs of individual agents in the same coalition, in the sense that they cannot claim a part of the other's payoff. Nash Equilibrium is defined as a set of strategies and payoffs associated with the property that no agent can benefit by changing its strategy while the other agents keep their strategies unchanged. A coalition structure is Pareto optimal if there is no agent that would be better off as part of another coalition without making another agent worse off.

One of the first works on coalitions in multi-agent systems is [Ketchpel94]. The author presents a utility distribution mechanism designed to perform in situations where there is uncertainty in the utility that a coalition obtains. The mechanism is based on assigning to one of the agents the responsibility for managing the group. The manager agent represents the coalition, pays the members, and earns a profit if the coalition's global utility is bigger then the members' payments. A polynomial model for coalition formation and payoff distribution is described and analyzed in [Shehory96]. The paper provides a coalition formation algorithm based on negotiation, called CNA. In each step at least one coalition makes an attempt to improve the payoffs of its members by making a coalition formation proposal to another coalition. The algorithm stops when the agents reach a stable state according to the polynomial k-stability defined in the paper. Recent research such as [Sandholm99] and [Sen2000] focuses on finding the optimal division of agents into disjoint coalitions so that the total sum of the payoffs of the coalitions is maximized. The context is characteristic function games (CFG) in which the value of each coalition is given by a characteristic function and the value of a coalition structure is defined as the sum of the values of the coalitions that it contains. Sandholm presents in [Sandholm99] a deterministic anytime algorithm for searching the optimal coalition structure, while Sen describes in [Sen2000] an order-based genetic algorithm (OBGA) for the same problem.

In summary, game theory research on coalitions has developed methods for rigorous formal analysis concerning issues of solution stability, fairness, and payoff disbursements. However, the methods are centralized and some of the underlying assumptions of the developed algorithms do not necessarily hold in real-world multi-agent systems.

2.2 Coalition Formation in DAI Research

Historically, research in Distributed AI (DAI) focused first on Distributed Problem Solving (DPS), i.e. decomposing complex problems in goals and subgoals that can be solved by distributed processors /agents and on the coordination / scheduling problems arising. In the last 10 years, the focus of these studies expanded to deal with more general problems in the new domain of Multi-Agent Systems (MAS), including the problem of coalition formation among cooperative agents. DAI researchers approached coalition formation by adopting game theoretical concepts and developing 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 [Zlotkin94], and communication issues [Sycara96].

Shehory [Shehory97] presents the implementation of distributed coalition formation algorithms within a real world multi-agent system. The paper reports on the RETSINA (REusable Task-based System of Intelligent Networked Agents) multi-agent system designed to integrate information gathering from web-based sources and decision support tasks. The agents are distributed and decide how to decompose tasks and whether to pass them to others, what information is needed at each decision point, and when to cooperate with other agents. Coalitions appear when groups of agents work jointly in order to accomplish their tasks. Agents have group rationality: they join a coalition only if they benefit as a coalition at least as much as the sum of their personal benefits outside of the coalition. Introducing the coalition formation mechanism is proved to increase the efficiency of group task execution. The coalition formation algorithm takes into consideration requirements and constraints arising from the dynamic nature of the environment in which the system operates. This approach to coalitions is closer to the ContractNet protocol because of the group rationality assumption. The reason for this is that the problems they try to solve are similar: task allocation and effective task execution in multi-agent systems.

DAI research on coalition formation has two important contributions: to prove that the mechanism is beneficial in the context of group rationality and to improve coordination and cooperation among agents in multi-agent systems.

2.3 Coalitions of Customers for the Electronic Marketplace

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 big 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 social structure is minimized, but the social functioning is complex due to the big number of interactions among agents. The agents have individual rationality and no social or group rationality requirements. Customer agents know that being part of a coalition allows them to buy big quantities that they might not need individually and this may bring them a discount from vendors, so they will 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 big quantities with a certain discount will bring the vendor at least the same profit as retail selling.

[Sycara2000] illustrates the economic incentives behind formation of such "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 impractical, hard to implement or scale up. Improvements have been proposed by [Lerman2000]. A physics-inspired mechanism for coalition formation 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 captures the dynamics of the coalition formation process in a phenomenological way, not directly derived from the microscopic theory. It 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. The measure for the efficiency of the system used in the paper is the global utility gain that represents the price discount all agents receive by being members of coalitions. The authors conclude that 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).

In summary, the models presented above do not provide a detailed microscopic description of the mechanism. The authors address only temporary coalitions that last one transaction. They talk about the need for trust among members of the same coalition, but they do not find solutions for this issue.

2.4 Congregating in Multi-Agent Systems

The notion of long-term groups of self-interested agents appears in Durfee [Durfee2000] 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. Congregating agents are expected to have long lifetime during which they take on different roles, perform different tasks, and interact with different agents. An "affinity group" is defined as a set of agents that share some common trait or preference and want to congregate with other agents that have the same trait and avoid agents that do not have it. 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. Congregating 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 of the agents . 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, while congregator agents are similar to consumers. Labelers have the role to label their congregations so that they can attract congregator agents; congregators can be part of only one congregation at a 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 Durfee [Durfee2000] are similar to coalitions of agents. The different term is used to distinguish them from the task-oriented coalitions existing in DAI research; and since it brings some religious connotations. It is interesting 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 a 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 offer only limited trading mechanisms - fixed price and auctions - and no support for negotiation or grouping. It may be advantageous to bring coalition formation mechanism into real-world electronic marketplaces, since coalition formation is profitable for both customers and vendors. The already existing Internet communities, like newsgroups, chat-rooms, and virtual cities hold a promise to create large-scale economies among similar minded customers.

Coalition models presented in the previous section refer to 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. Searching for suitable coalitions and deciding what coalition to join is time and resource consuming. Forming and running new coalitions at each step is also computationally expensive. To overcome these limitations we propose a model of long-term coalitions that last many transactions and have long lifetime. The concept is similar to the congregations proposed by [Durfee2000] in which agents are supposed to have long lifetime and repeated interactions with other agents in the coalition. 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, that 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 agents that can play both customer and vendor roles. 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. Previous work has not focused on the agent reasoning mechanism about coalition formation in detail. However, a microscopic level mechanism is crucial when developing a practical application. From the perspective of an individual agent coalition formation can be viewed as a decision making 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 some important factors such as different (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 a common experience. 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 [Vassileva98]. 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 [Ganzaroli99] and [Jonker99]. 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 must be able to trust the other agents. According to Sycara [Sycara2000] 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. We believe that modeling trust relationships among agents and using them in the agent reasoning about creating, entering or leaving groups of agents that know and trust each other is more suitable for practical coalition formation. We use a model inspired by Ganzaroli [Ganzaroli99] and Jonker [Jonker99] to represent trust in the internal state of an agent.

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 [Crispo99] or electronic commerce [Ganzaroli99]. [Crispo99] analyses the security of the existing delegation mechanisms and describes 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 [Ganzaroli99] 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 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. The interacting independent agents together with the trust relationships established among them form a network. An important characteristic of the network is its density, calculated as the number of existing relations among agents in the network divided by the number of all possible relationships in the network. Because of technological and time limitations, agents are bounded in the number of relationships they can maintain and it is reasonable to assume that the density of the network is negatively related to the size of the network, i.e. the larger the network the lower the density we can expect. If the network becomes too large then it will have no longer influence on the behavior of the agents. The idea of community based trust and the notion of network of agents linked by their trust relationships fit well our concept of coalitions based on trust and nurturing relationships between agents.

At the agent level, trust is defined in [Jonker99] as an attitude of an agent who believes that another agent has a given property. In our coalitions for the electronic marketplace trust depends on an agent's own experiences with the subject of trust and on its expectations for future interactions. Given the agent's long lifetime assumption, trust presents an evolution over time called the dynamics of trust. 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 loose 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 like "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. 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 given as an example for the framework presented above in [Jonker99]. Given a set of experience classes E and a set of trust qualifications T, we can define a trust function as:

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

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

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 $d \in [-1, 1]$ 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, 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:

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

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

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

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 previously defined.

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 trade goods in an open electronic market environment and 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 [Mudgal2000]. Negotiation can end either with a Rejection or with an Agreement between the two agents (and, in the latter case, with a 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 account 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^- (in our example [-1, 0]).

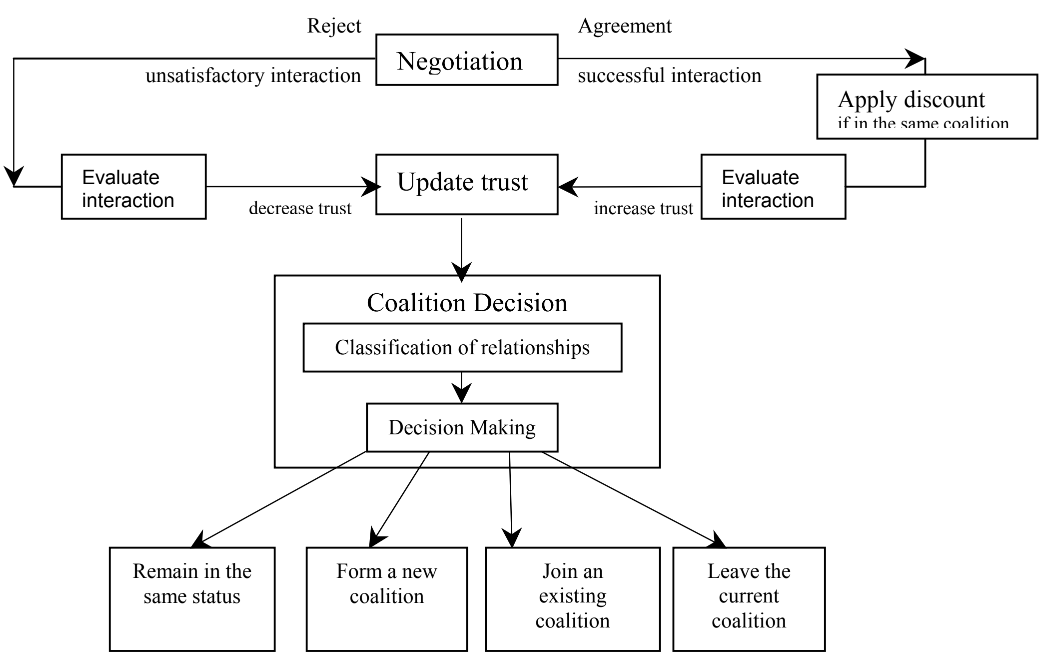

Figure1 Agent level reasoning description of the proposed coalition formation mechanism

Each agent stores a representation of all its trust relationships with other agents in the system with whom it 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 an agent's trust representation is changed, it triggers a Coalition Decision reasoning mechanism as shown in Figure1.

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, its trust relationships with other agents, and the coalitions in which these agents are. 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 approach an 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 trusted agent is outside of any coalition, they other agent form a new coalition. In the socially oriented classification the 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 score. If this coalition is different from the agent's current coalition, it leaves it and joins the new one.

6. Directions

This paper reviews existing coalition formation mechanisms in the AI literature with emphasis on the most recent research in electronic commerce. In the context of this new form of commerce, trust is an important issue that needs to be developed at an individual level and at a community level. We also present an existing formal model of trust based on the personal experience of the agent and we show how this model can be integrated in a 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. The motivation for this approach is to save computational resources and to gain system stability. 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. Another improvement that the proposed model brings is to minimize the last two stages of the general coalition formation mechanism: Payment Collection and Execution/Distribution. Detailed microscopic description of this mechanism is provided in the paper.

Some issues remain open for the moment, such as the evaluation of agent interactions and the solution for the decision-making problem of joining or leaving a coalition. The decision-making problem can be solved using an influence diagram or rule-based algorithms. Currently a Java-based multi-agent system for the electronic marketplace is being implemented. Agents trade goods using a negotiation mechanism described in [Mudgal2000]. The implementation of the new coalition formation mechanism is under development. Verification and analysis at agent level, coalition level, and system level are also planned.

References

[Ossowski99] S.Ossowski. Coordination in Artificial Agent Societies. Springer Verlag: Berlin-Heidelberg. 1999

[Rapoport84] J.P. Kahan, A.Rapoport. Theories of Coalition Formation. Lawrence Erlbaum Associates Publishers. 1984

[Ketchpel94] S. Ketchpel. Forming coalitions in the face of uncertain rewards. Proceedings of the Twelfth National Conference on Artificial Intelligence, Seattle, July 1994, pp. 414--419

[Zlotkin94] 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

[Shehory96] O. Shehory, S. Kraus. A Kernel-oriented model for coalition formation in general environments: Implementation and Results, Proceedings of the Thirteenth National Conference on Artificial Intelligence, Portland, 1996

[Shehory97] O. Shehory, K. Sycara Multiagent coordination through coalition formation. Proceedings of Agents Theories, Languages, and Methodologies, ATAL 97, 1997.

[Sandholm99] T. Sandholm, K. Larson, M. Andersson, O.Shehory, F. Tohme Coalition structure generation with worst-case guarantees. Artificial Intelligence vol.111, no. 1—2, p. 209—238, 1999

[Durfee2000] C.H. Brooks, E.H. Durfee, A. Armstrong An Introduction to Congregating in MultiAgent Systems. Proceedings of Fourth International Conference on MultiAgent Systems, Proceedings of ICMAS’2000.

[Sen2000] S. Sen, P.S. Dutta Searching for optimal coalition structures. , Proceedings of ICMAS’2000

[Lerman2000] K. Lerman, O. Shehory Coalition Formation for Large Scale Electronic Markets. Proceedings of ICMAS’2000

[Sycara2000] M Tsvetovat, K. Sycara Customer Coalitions in the Electronic Marketplace. Proceedings of the Fourth International Conference on Autonomous Agents 2000

[Sycara96] K. Sycara, K. Decker, A. Pannu, M. Williamson, D. Zeng Distributed intelligent agents. IEEE Expert, December 1996

[Sandholm99] T. Sandholm, K. Larson Anytime Coalition Structure Generation An Average Case Study. Proceedings of Autonomous Agents, Seattle, 1999

[Ganzaroli99] A. Ganzaroli, Y.Tan, W. Thoen The Social and Institutional Context of Trust in Electronic Commerce. Autonomous Agents, Deception, Fraud and Trust in Agent Societies, Seattle 1999

[Crispo99] B. Crispo, B. Christianson A Note About the Semantics of Delegation. Autonomous Agents, Deception, Fraud and Trust in Agent Societies, Seattle 1999

[Jonker99] C. Jonker, J. Treur Formal Analysis of Models for the Dynamics of Trust based on Experiences. Autonomous Agents, Deception, Fraud and Trust in Agent Societies, Seattle 1999

[Mudgal2000] C. Mudgal, J. Vassileva An influence diagram model for multi-agent negotiation. Proceedings of the International Conference on Multi-Agent Systems ICMAS'2000, Boston, July 2000

[Vassileva98] J. Vassileva Goal-Based Autonomous Social Agents Supporting Adaptation and Teaching in a Distributed Environment, In Proceedings of ITS'98, San Antonio, Texas.LNCS No1452, Springer Verlag: Berlin pp.564-573