

A Mechanism for Forming Long-Term Coalitions of Customer and Vendor Personal Agents

Silvia Breban and Julita Vassileva*

Computer Science Department, University of Saskatchewan

1C101 Engineering Bldg. 57 Campus Drive

Saskatoon, SK, S7N 5A9 CANADA

jiv@cs.usask.ca

Abstract

Recently coalition formation has been explored in the area of electronic marketplace as temporary buying groups. We extend this concept to long-term coalitions that are formed of both customer and vendor agents after evaluating their relationships with other agents in the system. We propose a coalition formation mechanism designed at microscopic (agent) level as a decision problem and we analyze the effect of this mechanism at both microscopic and macroscopic (system) levels. Our results show that forming coalitions is beneficial for both the system (it reaches an equilibrium state) and for the agents (they have high gain increase over time).

Introduction

Most existing electronic markets like eBay and OnSale offer only limited trading mechanisms - fixed price and auctions - and no support for negotiation or grouping. Most research studies in game theory, DAI, and even electronic commerce take a global perspective and study protocols that optimize the local behavior of the markets. Only a few studies of automated negotiation [1, 2] take a microscopic perspective and focus on the reasoning of an individual agent. Our work aims at bridging the gap by developing a reasoning mechanism for each agent and analyzing the effect on a global level when many agents with this reasoning interact. Our focus is on self-interested agents forming coalitions in an electronic marketplace.

Some electronic markets, like LetsBuyIt exploit the idea of selling products to groups of people for discounted price. The larger the number of customers that wish to purchase the product, the lower the price becomes. This idea reflects some recent studies in the area of electronic markets [3, 4, 5] which deploy the formation of groups to improve the coordination and cooperation among self-interested agents as well as to increase their financial benefits. The main goal is to form customer coalitions for open agent societies on the Internet, where the set of agents in the system has a large size with possible high variations in time. The mechanisms developed in this area are inspired from previous approaches in the areas of game theory [6, 7] and

distributed AI [8] where the application domain is closed societies with a given finite sets of agents. The economic incentives that arise behind the formation of such temporary coalitions are presented in [3]. The authors describe a general scheme for coalition formation that can be used in pre-negotiation and post-negotiation coalitions. A physics-inspired mechanism that treats agents as randomly moving, locally interacting entities is proposed in [4]. This work is important since its authors are the first to consider the realistic case when agents can leave a coalition once they join it, behavior that is proved to be beneficial for the system and for the agents. Yamamoto and Sycara [5] propose a buyer coalition scheme based on a reverse auction where sellers bid discounted prices to sell large quantities. A buyer group is formed for a category of items, not for a particular item. The paper proves that the scheme can accommodate large numbers of buyers and provide them with increased benefits and higher chances to obtain the desired items.

A limitation of all coalition or group formation mechanisms proposed so far is that they model coalitions that last only one transaction. Another limitation is that each customer agent has to decide what coalition to join without memory of previous experiences. At microscopic level, the search for suitable coalitions and the decision of 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. It leads to increased dynamics of the system - high variation in the number of coalitions and in the size of each coalition - not desired in a large-scale multi-agent system.

We propose a new approach to forming long-term coalitions for the electronic market. We see coalition formation as a decision problem for the individual agent: to increase its long-term utility the agent has to decide in each epoch whether to join a coalition, form a new one, leave a coalition, or remain in the same one. Our primary goal is to provide the agent with a reasoning mechanism about coalition formation that takes into consideration its long-term utility, its preferences, and the relationships it has with its partners. Secondly, the proposed coalition formation mechanism is designed to accommodate large numbers of agents (thousands and millions) due to minimized

communication between agents and reduced complexity. Finally, for system stability and predictability reasons our approach has two more objectives: to reduce the dynamics of the agents (i.e. their movement between coalitions) and to stabilize the number of coalitions in the system.

Long-Term Coalitions

We address an electronic market composed of personal agents representing human users. They have fixed roles of either customer or vendor. They trade the same class of goods (e.g. books) on the Internet and use the same reasoning mechanism. We allow the agents to form coalitions to increase their individual benefits. A coalition also has the benefit of providing a smaller and more familiar environment that reduces search time in the routine interactions between agents. Hopefully, in this way coalitions will help to improve the efficiency of the market. In the design of the proposed coalition formation mechanism we make the following assumptions:

- Agents have individual rationality. They try at any time to maximize their own long-term utility. A vendors' utility is represented by its sales over time. A customer's utility consists of the amount gained from discounted transactions.
- Agents have a long lifetime of repeated interactions with other agents in the system and no interdiction to interact with agents outside of their coalition.
- Agents may have different interests in the goods that being traded (e.g. science fiction, romance, history books). Agents may belong to different economic categories (e.g. a customer that affords to buy books between \$20 and \$60 and a vendor that sells only for more than \$100 belong to different categories and can never engage in successful trade).
- Coalitions may have a long lifetime once created. They are disjoint, i.e. one agent cannot belong to two coalitions at the same time. While this assumption seems quite restrictive it is not unrealistic, assuming that a coalition reflects a range of agent interests and goods that can be traded. In our future research we will relax this condition.
- A coalition is created when an agent wants to form a new coalition with another agent. The latter always agrees, since there is no cost for joining a coalition. This may seem as unrealistic assumption too, if we consider real people. However, we are talking about software agents that follow a particular protocol defined here. A coalition dissolves when composed of only one agent.
- An agent can join or leave a coalition at any moment. However, a penalty might have to be paid by agents leaving a coalition. This reflects the costs that a vendor agent has to incur by losing his clientele, and the higher search costs a client has to pay to find suitable agents for trading in a new coalition.

- Agents in the same coalition receive a specific fixed discount for each transaction executed. Thus the agents prefer to be part of the coalition with the agents with whom they expect to have most frequent transactions future.
- The coalition structure (the partition of agents into coalitions) is global knowledge. This is also an assumption that does not seem realistic in real-life markets. However, in a market formed by agents following the described protocol, this can be achieved.

A vendor agent enters a coalition to increase its sales. It prefers to be part of the same coalition as customer agents with whom it has most transactions and it agrees with a certain discount for each transaction inside its coalition. Figure 1 presents the idea visually: instead of selling in bulk at one moment to several customers, as proposed in previous approaches on coalition formation, the vendor uses a policy that leads to selling a large quantity to the same customer after repeated transactions at different times.

Thus a vendor joins coalitions to get closer to customers that have compatible preferences, and thus to nurture vendor-customer relationships. The concept is similar to the established practice in real-life markets like Safeway or Sears that give a minimal discount to members of their clubs. This policy known in economics as Customer Relationship Management (CRM) [9] and is motivated by the fact that establishing a friendly and trustworthy relationship with clients promises vendors more transactions for the long run and retention of customers.




Figure 1: Discounted transactions for:

A) all customers in a one-transaction coalition;

B) one customer in a long-term coalition at different times.

The advantage for customers of being members of long-term coalitions in this approach is that they are sure to get a discount at any time, while in previous approaches the discount depends entirely on finding other agents interested in the same product at the same time, agreeing to buy the necessary amount of goods, and successfully negotiating with vendors.

The Decision Problem

From the point of view of a decision-making agent with finite computational resources, it makes sense to view the coalition formation as a problem of decision making in the face of uncertainty. An agent must decide at each epoch whether or not a change its coalition according to its best interests. It may choose to form a new coalition, to join an

existing coalition (possibly by leaving an existing coalition first), to leave a coalition (and remain independent), or, finally, to remain in its current coalition.

An agent is assumed to know the current coalition structure with certainty before having to take any action. As well, the agent has memory of previous interactions with other agents. However, the agent does not know how the other agents will act in the current epoch. There are basically three options. First, the agent could consider all possible actions of all other agents, and try to determine a suitable steady state; this we discard, since we assume that the agent cannot determine this state with the computational resources available to it. Second, our agent could treat the behavior of other agents as stochastic, by providing a probability distribution over possible coalition structures in the next epoch. However, there are too many possible coalition structures to consider by enumeration. Third, the agent can assume that the coalition structure will remain static, except for the action it takes. For rationally bounded agents this approach has some psychological relevance and is feasible. The agent's action directly affects the newly formed coalition structure, and this affects the agent's utility. The new coalition structure is different from the old one only regarding the current agent. Also the agent does not know with what coalition it will have most transactions in the future – that depends on the new coalition structure. If we represent the decision problem about coalition formation that the agent faces as an influence diagram the solution can be found by maximizing the agent's utility; the action associated with the maximum utility is the most beneficial for the agent.

To directly solve the agent decision problem about coalition formation using an influence diagram would involve expensive computation since a couple of choice nodes may vary over all possible coalition structures. As mentioned by Sandholm et al. [10] the total space of enumerating (or searching) possible coalition structures is exponentially proportional with the number of agents in the system.

Therefore, we decided to use an intermediate utility function that depends on the past transactions of the agent and on its future expectations. This function characterizes each coalition from the coalition structure. It reflects the agent's opinion about having common interests and preferences with agents from that coalition that would promise increased profit within the coalition. We name this utility function the agent's *relationship with a coalition*. The strength of this relationship reflects the agent's expectation to have fruitful future transactions within the coalition. When the agent knows the strengths of all its relationships with existing coalitions it can solve its decision problem by finding and joining the coalition with which it is most strongly related. For this the agent maintains a simplified model of its past experience with any other agent in the system that we call *inter-agent relationship*. When the agent needs to decide which coalition is most profitable it calculates its relationship with each coalition as a function of its inter-agent relationships

with agents from the coalition. We present two different definitions of this function in section 4.

We represent the strength of inter-agent relationships using the trust model proposed by Jonker and Treur in [11]. Given a set of experience classes E and a set of predefined *inter-agent relationship strength* quantifications S , a mapping for the transition from one strength value s to another $strength(e, s)$ can be defined as:

$$strength : E \times S \rightarrow S$$

$$strength(e, s) = (1 - d) * e + d * s$$

We consider the case in which an experience can take any value in the interval $E = [-1, 1]$. If an experience e is evaluated as a positive one it is assigned a positive value from E^+ ; if e is a negative experience it takes a negative value (from E^-). We consider the set of predefined strength quantifications $S = [-1, 1]$. The parameter $d \in [0, 1]$ is an inflation rate used to model the fact that older experiences become less important over time, while the most recent experience is the most relevant (since the agent preferences may change in time). In this function after each new experience e , the existing strength of the relationship s is multiplied by d and the impact of the new experience e is added, normalized so that the result fits in the desired interval S .

Based on this representation and on the set of discrete time values when experiences take place $Time = \mathbb{N}^+$ (the set of natural numbers), a evolution function $evol$ is inductively defined in [11]. This function is used by the agent when it has to update its relationship in another agent at each step from the $Time$ set:

$$evol : E \times Time \rightarrow S$$

$$evol(e, 0) = 0$$

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

The definition of the evolution of inter-agent relationship strength specifies that the initial strength for step 0 is set to a neutral value 0 . At each step $i+1$ the strength is updated based on the previous value (from step i) and the current experience e_i according to the strength function defined above. We use this formal model to represent inter-agent relationships.

Coalition Formation Mechanism

Scenario

We consider a system of multiple agents trading books on behalf of their users in an open electronic market. Before a transaction between a customer and a vendor is executed, the two agents go through a bilateral negotiation phase to agree on a certain price using a previously developed negotiation mechanism [1]. The negotiation consists of an iterative process in which the two agents make offers and counteroffers based on the preferences of their users and on the reply of their opponent. The agents negotiate according to preferences set up by their users, which define the minimum acceptable price for vendors and the maximum

affordable price for customers, the subjective importance that money has for the users, the urgency of the current goal of selling or purchasing a certain product, the risk-attitude, and the time constraints for executing the current transaction [1]. These preferences play a crucial role in the result of the negotiation and influence the compatibility or incompatibility between the agents (and correspondingly, between their users). Therefore, negotiations that result in failure are interpreted by the agent as negative experience. If the negotiation has completed with success, the actual transaction is executed, and the result is evaluated by the user. This evaluation may include many factors, e.g. if the good was delivered in time, whether the quality was acceptable, if the payment was received in time, etc. The user then informs his/her agent, about the final result of the evaluation, if it was *unsuccessful* or *successful* transaction. The agent interprets this information as negative or positive experience, correspondingly. Positive experiences are assigned values from the positive subset of experience classes E^+ while negative experiences are evaluated in the negative subset of experience classes E^- (see Figure 2).




Figure 2: The coalition formation mechanism.

For the purposes of our experiment, we limited the scope of evidence for the strength of the relationship to the outcome of the negotiation between the agents, because modeling the actual transaction involving payment and delivery of goods with variable quality in the simulation would have added complexity that would have made it hard to interpret the results of the experiment. In addition, it would not have contributed to the evaluation of the coalition formation mechanism. So for the purpose of evaluating this coalition mechanism, we assume that once a negotiation between agents has resulted in success, the goods were always delivered in time and were of satisfactory quality. Note, however, is not a limitation to the mechanism per se, but just a simplification for the purpose of evaluation. When the new experience is evaluated, the strength of the relationship that the two agents have in each other is updated according to the evolution function defined in the previous section. Each agent stores a representation of its

relationships with all agents in the system with whom it has ever interacted. The agent's relationships with the other agents have a strength of *null*. An agent represents the relationship with another agent by a tuple, consisting of the other agent's name and a relationship strength value from S .

Each change in the representation of the agent's relationships triggers the coalition reasoning mechanism shown in Figure 2. The agent makes use of its relationships with individual agents to calculate its relationships with the currently existing coalitions in the system. The agent's relationship with a specific coalition is calculated as a function of the agent's relationships with the individual agents from that coalition. There are many possible ways to define this function. We implemented and evaluated two different ones. The first function (called *soc1* strategy) defines the strength of an agent's relationship with a coalition by calculating the sum of the strength of the inter-agent relationships that the agent has with the agents from that coalition. The second one (called *soc2* strategy) simply computes the *number of all agents* in the coalition with whom the agent has positive relationships. Interestingly, our experiments showed that though these two functions appear to be similar, they lead to different macro- and microscopic results.

Microscopic Description

To express formally the problem, let us assume that the multi-agent system consists of n agents that form at the current moment m coalitions. A coalition is identified by a label that is a number between 1 and m . We represent the special case of an agent outside of coalitions as the agent belonging to a special coalition with label 0. We assume as global knowledge in the system the set of agents $A = \{A_1, A_2, \dots, A_n\}$ and the set of their corresponding coalitions $C = \{C_1, C_2, \dots, C_m\}$, where C_i is the label of the coalition to which A_i belongs (remember the assumption that each agent can belong to only one coalition at a time). Agent A_i has relationships with k agents (a subset of A) that we represent as vector $R(A_i)$:

$$R(A_i) = (A_{i1}, A_{i2}, \dots, A_{ik})$$

where $A_{ij} \in A \setminus \{A_i\}$ for $1 \leq j \leq k$. We denote by $S(A_i)$ the corresponding vector of relationship-strengths that A_i has in each of these agents:

$$S(A_i) = (S_{i1}, S_{i2}, \dots, S_{ik})$$

where S_{ij} is the strength of the relationship that A_i has with A_{ij} . $S_{i1}, S_{i2}, \dots, S_{ik}$ take values from the set of predefined quantifications S . The corresponding coalitions of these k agents form the set of coalitions $C(A_i)$:

$$C(A_i) = (C_{i1}, C_{i2}, \dots, C_{ik})$$

where C_{ij} is the label of the coalition to which A_{ij} belongs. The current coalition of A_i is denoted by C_i . Coalitions $C_i, C_{i1}, C_{i2}, \dots, C_{ik}$ have values of either 0 (if the agent is outside of all coalitions) or between 1 and m (i.e. the coalition label, if the agent is in a coalition).

As explained in the previous section, the agent knows the coalition structure, since it is public knowledge. As a result it knows the coalition it belongs to C_i and the coalitions

$C(A_i)$, in which agents from $R(A_i)$ belong. It also knows the strengths of its own relationships with other agents $S(A_i)$. Based on these, it calculates the strength of its relationships in each coalition. The new vector is denoted by $SC(A_i)$:

$$SC(A_i) = (SC_{i0}, SC_{i1}, \dots, SC_{im})$$

where SC_{ij} represents the strength of the relationship that A_i has with coalition j . In the *soc1* strategy the strength of the relationship of agent i with coalition k is calculated as the sum of the strength of relationships between agent i and the agents in the coalition:

$$SC_{ik} = \sum_j (S_{ij}),$$

where agent A_{ij} is in coalition with label k (i.e. $C_{ij} = k$), j varies from 1 to $|C_{ij}|$, and k varies from 1 to m . In the *soc2* strategy the same relationship strength is defined as the number of agents from the coalition with whom the agent has relationships with positive strength, or:

$$SC_{ik} = |R'(A_i)|,$$

where $R'(A_i)$ is a subset of $R(A_i)$ defined as the set of agents from coalition k with whom A_i has positive relationships. An agent A_{ij} is in $R'(A_i)$ if and only if $C_{ij} = k$ and $S_{ij} > 0$. For both strategies we consider the special case of agents that are outside of coalitions as being members of coalition 0. The rule-based algorithm is the same for the *soc1* and *soc2* strategies:

Build SC(A_i) - the vector of relationships in coalitions

Find coalition k with highest relationship strength

if (A_i in a coalition) AND (k is not the coalition of A_i)

A_i leaves its current coalition

if (k is not coalition 0) then

A_i joins coalition k

else

A_i forms a new coalition with the agent(s) outside of coalitions with whom it has highest relationships

We evaluate and compare the effect of the two agent strategies (*soc1* and *soc2*) in the next section.

Evaluation

We have developed a simulation prototype of the proposed coalition formation mechanism in Java. We ran 18 sets of experiments with different parameter configurations over several platforms: Windows 2000, Windows NT, HP Unix, Sun Solaris, and Linux. Each set of experiments consisted of 100 trials over which the results were averaged.

Our goal was to evaluate the mechanism at *macroscopic* and *microscopic* levels. For the first part we investigated the number of coalitions in the system, the overall dynamics, and how these factors evolve over time. The evolution of the number of coalitions is relevant for reasons of predictability while the system dynamics (calculated as the sum of the number of coalitions visited by each agent) is important in establishing whether the system reaches an equilibrium state or not. For the microscopic evaluation we focused on the individual gains of the customer agents (calculated as the average of the sum of benefits obtained from all discounted transactions by each customer). The

experiments were intended to compare the different agent strategies described in the previous section (*soc1* and *soc2*). Some variables involved in the design of the mechanism were set constant for all experiments: the inflation rate of the relationship strength ($d = 0.5$), the evaluation of positive experiences (0.2), the evaluation of negative experiences (-0.2), and the discount rate (5% of the price). The parameters under investigation are summarized in Table 1.

# customers	100; 1000
# vendors	100
# interactions	1; 100; 1000; 10,000; 100,000;
agent strategy	<i>soc1</i> ; <i>soc2</i>
# categories	1; 5; 10

Table 1: Simulation parameters and their values

Note that the number of vendor agents was set to 100, while the number of customer agents was varied to 100 and 1000. A significant parameter for the evolution of different factors over time is the number of interactions between agents (each corresponding to an epoch of the mechanism). Another simulation parameter used in our evaluation was the number of categories of agents. A category is meant to predefine the classes of interests and preferred ranges of prices that are set up by the human users. Vendor and customer agents from different categories never reach an agreement since their preferred ranges of prices do not overlap.

Configuration		Peak values		Lowest values	
		<i>soc1</i>	<i>soc2</i>	<i>soc1</i>	<i>soc2</i>
100V/ 100C	1 ctg	32	34	5	1
	5 ctg	19	39	14	5
	10 ctg	8	31	7	10
100V / 1000C	1 ctg	88	87	7	1
	5 ctg	75	74	30	5
	10 ctg	72	80	80	10

Table 2: Number of coalitions (peak and lowest values)

Configuration		Dynamics		Agent Gains	
		<i>soc1</i>	<i>soc2</i>	<i>soc1</i>	<i>soc2</i>
100V/ 100C	1 ctg	44	1	3	13
	5 ctg	20	0	1	2.6
	10 ctg	2	0	1.2	1.3
100V / 1000C	1 ctg	50	3	0.2	1.3
	5 ctg	23	3	0	0.2
	10 ctg	15	2	0	0.1

Table 3: Highest dynamics and agent gains (in thousands)

We plot in Figure 3 the results that show the evolution of the number of coalitions, system dynamics, and the individual agent gain. Graphs A and D represent the number of coalitions for different agent strategies: *soc1* and *soc2*; graphs B and E – the system dynamics for *soc1* and *soc2*; graphs C and F – the individual agent gain for *soc1* and *soc2*. On the X-axis the number of interactions is

represented on a logarithmic scale from 1 to 1,000,000; the Y-axis shows the investigated parameter.

points with 1 million interactions) in Table 3. We note that for *soc1* the values are higher and for *soc2* they are lower.

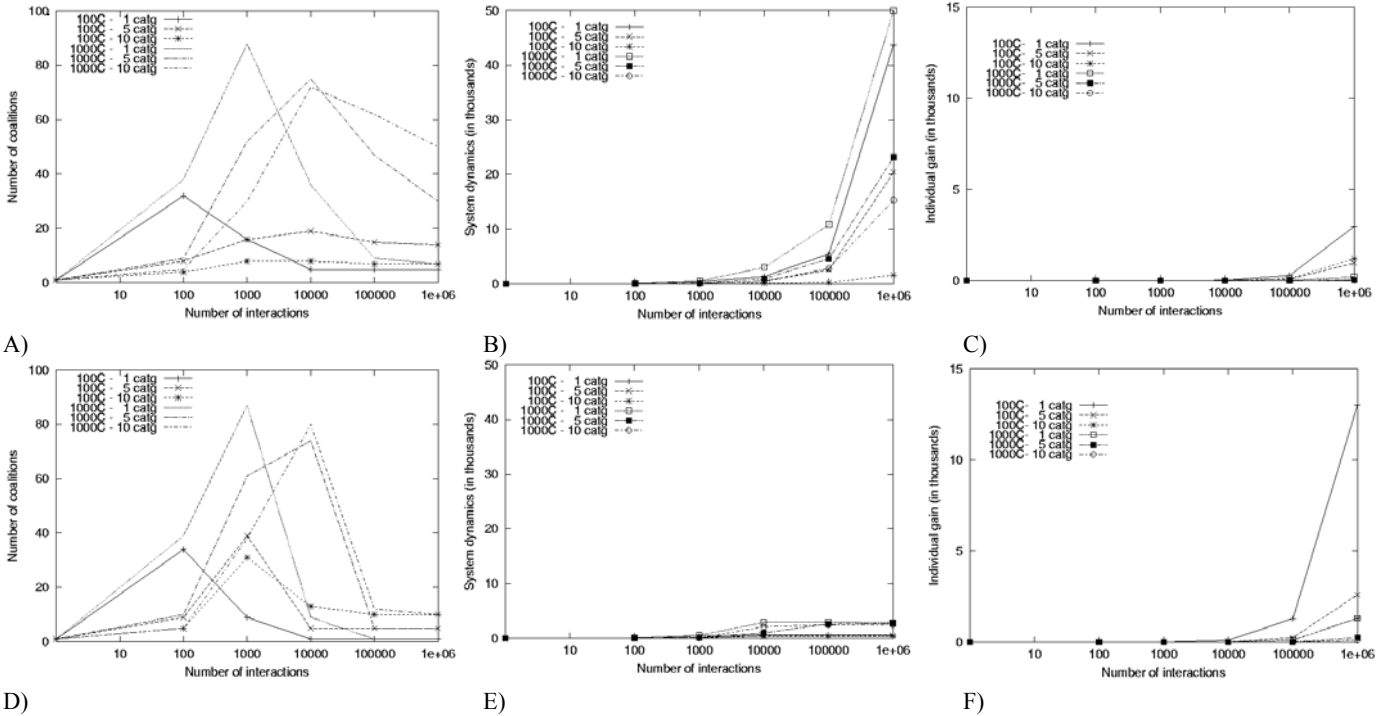


Figure 3: A) Number of coalitions for *soc1*; B) System dynamics for *soc1*; C) Individual gain for *soc1*; D) Number of coalitions for *soc2*; E) System dynamics for *soc2*; F) Individual gain for *soc2*.

The number of coalition, shown in Figures 3A and 3D, ranges from 1 to 100. Our results show that as the number of interactions among agents increases, the number of coalitions first grows, it reaches a peak, and then it starts to decrease. At the beginning the agents form coalitions and after a while they start merging. We focus on analyzing the peak values that reflect the formation of coalitions and the lowest values (after 1 million interactions) that reflect the merging rate of the coalitions. The difference between these two values shows the speed of coalition merging. Both the peak and the lowest values are shown in Table 2.

The number of coalitions has a predictable and controllable evolution over time that depends on the number of vendors (this is its higher limit); on the number of customers (it has different merging behaviors for 100 and 1000 customers); on the number of agent categories (this is its lower limit); and on the agent strategy (in *soc2* the drop is faster, while in *soc1* a pronounced decrease in the peak and drop values is noticed when more categories are considered). These results can be explained by the fact that it becomes harder for the agents to find proper coalitions when more customers or agent categories are considered. With the *soc2* strategy the number of coalitions drops faster

Our second evaluation factor is the system dynamics – represented in thousands in Figures 3B and 3E. The sum of the number of coalitions visited by all agents varies from 0 to 50,000. We show the peak values (reached in sample

The individual agent gain – the average of the sum of benefits obtained from all discounted transactions by each customer is shown in Figures 3C and 3F to range from 0 to 15,000. The highest values of individual gain (reached after 1 million interactions) are presented in Table 3.

Overall, our results show that the proposed coalition formation mechanism is beneficial for the customer agents and for the system. It ensures high benefits over time for the customers in all strategies. The mechanism leads to a predictable behavior of the system since the number of coalitions drops quickly to small values (limited by the number of agent categories) for all strategies. It also brings stability to the system since the overall dynamics reaches an equilibrium state for the *soc2* strategy (note that it remains nearly constant for 900,000 interactions). Although the system dynamics increases for *soc1*, the increase is linear with the number of interactions (note the logarithmic scale of axis X), and we expect it to stabilize after a larger number of interactions. The explanation of this behavior is that when most agents belong to the same coalition as the partners with whom they share similar interests and preferences, they stop moving from one coalition to another. This leads to stabilization in the number of coalitions and in the system dynamics as well as to high increase in the individual gains of customers.

We explored several factors that influence the time in which this equilibrium is reached: the agent strategy, the number of customer agents in the system, and the number of categories of agents. The strategy where an agent joins

the coalition with the highest number of strongly related agents (*soc2*) proved to be more beneficial than the strategy where an agent joins the coalition with the highest summative strength of relationships (*soc1*). When the number of customer agents is equal to the number of vendor agents, the merging rate of coalitions is smaller compared to the case with larger number of customers. More predefined agent categories, i.e. preference diversity among the agent, leads to a delay in the formation of coalitions and in their merge rate that has impact on the system dynamics. The categories also reduce the individual gains. This seems to confirm the intuition that agents are better off in homogenous systems where all agents are compatible (i.e. in 1 category) than in fragmented systems. Further experiments have been carried out to evaluate the proposed mechanism, investigating other functions for defining strength of relationship between an agent and coalition [12], and considering a cost for leaving coalitions [13]. A detailed description of the mechanism and experiments is available in [14].

Conclusions and Future Work

In this paper, we proposed and evaluated a mechanism for long-term coalition formation that takes into consideration the relationships between agents. The mechanism extends the existing work on temporary customer-only coalitions to long-term coalitions formed by customers and vendors. We showed that our coalition formation mechanism brings stability to the system (in the number of coalitions and in the overall dynamics) and provides the customer agents increased benefits over time. We compared two agent strategies. The results can be summarized as “go to the coalition where you have the highest number of friends (even if they are not your best friends)”. The mechanism uses reduced communication between the agents that makes it scalable for large numbers of agents and interactions. Our plans for future work include investigation of the proposed coalition formation mechanism under weaker assumptions, for example allowing differentiated goods to be traded in the system, allowing the agents to belong to more coalitions at a time, and to refuse entering coalitions.

References

[1] Mudgal C., Vassileva J. (2000) *Bilateral Negotiation with Incomplete and Uncertain Information: A Decision-Theoretic Approach Using a Model of the Opponent* Proceedings of the 4th Intl Workshop, CIA, Boston, Springer LNAI 1860, July 2000, pp. 107—118.

[2] Wong W. Y., Zhang D.M., Kara-Ali M. (2000) *Towards an Experience Based Negotiation Agent* Proceedings of the 4th Intl Workshop, CIA, Boston, Springer LNAI 1860, pp. 131—142.

[3] Tsvetov M., Sycara K. *Customer Coalitions in the Electronic Marketplace* Proceedings of the 4th International Conference on Autonomous Agents, Barcelona, 2000, pp.263—264.

[4] Lermann K., Shehory O. *Coalition Formation for Large Scale Electronic Markets* Proceedings of the Fourth International Conference on Multiagent Systems ICMAS'2000, Boston, July 2000, pp. 216—222.

[5] Yamamoto J., Sycara K. *A Stable and Efficient Buyer Coalition Formation Scheme for E-marketplaces* Proceedings of Autonomous Agents 2001, Montreal, 2001, pp. 576—583.

[6] Sen S., Dutta P.S. *Searching for optimal coalition structures* Proceedings of the International Conference on Multi-Agent Systems ICMAS'2000, Boston, July 2000, pp. 286—292.

[7] Sandholm T., Larson K. *Anytime Coalition Structure Generation An Average Case Study* Proceedings of Autonomous Agents, Seattle, 1999, pp. 40—47.

[8] Dutta P., Sen S. *Identifying partners and sustenance of stable, effective coalitions* Proceedings of Autonomous Agents 2001, Montreal, 2001, pp. 23—24.

[9] Customer Relationships Management (available on-line at <http://www.crm-expo.com/>, accessed April 21, 2001)

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

[11] Jonker C., Treur J. *Formal Analysis of Models for the Dynamics of Trust based on Experiences* Autonomous Agents, Deception, Fraud and Trust in Agent Societies, Seattle, 1999, pp.81—94.

[12] Vassileva, J. Breban, S. Horsch M. (to appear) *Agent Reasoning Mechanism for Long-Term Coalitions Based on Decision Making and Trust*, Computational Intelligence, Special Issue on Agent Mediated Electronic Commerce.

[13] Breban S., J. Vassileva (2002a) *Using Inter-Agent Trust Relationships for Efficient Coalition Formation*, to appear in R. Cohen and B.Spencer (eds.) Proceedings of the 13th Canadian Conference on AI, Calgary, 28-30 May, 2002, Springer Verlag.

[14] Breban S. (2002) *Long-Term Coalitions for the Electronic Marketplace*. M.Sc. Thesis, University of Saskatchewan, February 2002. (<http://bistrica.usask.ca/madmuc/>)

Acknowledgement: This work has been supported by an NSERC Individual Research grant.