

Agent Reasoning Mechanism for Long-term Coalitions Based on Decision Making and Trust

Julita Vassileva, Silvia Breban, Michael Horsch

University of Saskatchewan, Computer Science Department,

Saskatoon, Saskatchewan S7N 5A9, Canada

{jiv, svb308, horsch}@cs.usask.ca

Abstract. We address long-term coalitions that are formed of both customer and vendor agents. We present a coalition formation mechanism designed at the agent level as a decision problem. The proposed mechanism is analyzed at both system and agent levels. Our results show that the coalition formation mechanism is beneficial for both the system – it reaches an equilibrium state – and for the agents – their gains highly increase over time.

Keywords: long-term agent coalitions, agent trust model, agent decision making, evaluation

1. Introduction

Group formation and activity be it in the context of coalitions, teams or congregations, is an important way for establishing close cooperation among agents. Research on the formation of agent groups and on their coordination and activity has received an increasing amount of attention in recent years, as the potential benefits of such close agent cooperation have been recognized. Only a few works in the area of multi-agent systems for the electronic marketplace (Lermann and Shehory 2000, Tzvetovat and Sycara 2000, Yamamoto and Sycara 2001) deploy the formation of groups to improve the coordination and cooperation among self-interested agents as well as to increase their financial benefits. The mechanisms developed in this area are inspired from previous approaches in the areas of game theory (Shehory and Kraus 1996) and distributed AI (Shehory and Sycara 1997) where the application domain is closed societies with a given finite sets of agents.

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 a 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 other preferences such as group

rationality or a long-term utility, however, the problem which coalition to choose becomes more challenging. In contrast, the macroscopic perspective has as its fundamental unit the coalition. 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 to ensure some global system optimality and/or stability criteria. The partition of agents into disjoint coalitions is called a coalition structure.

Most research studies that deploy inter-agent protocols for multi-agent systems (such as coalition formation or negotiation) take either a global perspective (Shehory and Sycara 1997, Tzvetovat and Sycara 2000) or a local perspective (Mudgal and Vassileva 2000) without studying the effects of the existing interaction between the two perspectives. Our work aims at bridging the gap by developing a reasoning mechanism for the individual agent and analyzing the global effects of a system formed of interacting agents that use this reasoning mechanism.

A limitation of all coalition formation mechanisms proposed so far (Lermann and Shehory 2000, Tzvetovat and Sycara 2000, Yamamoto and Sycara 2001) is that they model coalitions that last only one transaction. Another limitation is that each agent in the system has to decide each time it wants to buy/sell a product what coalition to join without memory of previous experiences, i.e. it does not learn from its experience. At the 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 (e.g. high variation in the number of coalitions existing in the system and in the size of each coalition) that is not desired in a large-scale multi-agent system.

In general, coalitions present a loose organizational structure with an informal contract among agents. While in formal contracts there is an implicit formal trust in the structure and the regulations of the system that needs no explicit specification, in the context of informal contracts each agent in the group should be able to trust explicitly the other agents. This is especially important in long-term coalitions, where the agents can learn from their previous experience what to expect in their future interactions. None of the existing coalition formation mechanisms deals with this issue, although some authors (Tzvetovat and Sycara 2000) mention it.

Therefore, our goal is to provide a reasoning (micro-level) mechanism for agents taking into consideration the long-term utilities of individual agents, their preferences, and the trust they have in their partners. Secondly, in order to be evaluated on a macro-level the proposed mechanism should scale up, so that it could be tested with large number (thousands) of interacting agents, each of them using the proposed reasoning mechanism. To achieve this scalability, the reasoning mechanism requires minimized communication between agents and a reduced complexity. Finally, for reasons

of system stability and predictability our approach has to reduce the dynamics of the agents and stabilize the number of coalitions in the system.

2. The Domain and the Assumptions

We focus on self-interested agents forming coalitions in an electronic market. Our agents trade a particular type of goods (e.g. books) on the Internet on behalf of their users and try at any time to maximize their own long-term utility. Before an item is traded between a vendor and a customer, the agents go through a bilateral negotiation phase to agree on a price. We name this process an *interaction*. We assume that the agents have a long lifetime of repeated interactions with other agents in the system and that the electronic market provides system matchmaker agents that are responsible for finding suitable partners to interact when a need arises. Our agents may have different interests in the items being traded (e.g. science fiction, romance, history books) and may belong to different economic categories (e.g. a customer that affords to buy books between \$20 and \$60 and a vendor that sells for more than \$100 belong to different categories).

The agents form coalitions to increase their individual benefits. Customer agents receive a fixed percentage discount when they buy from vendor agents in their coalition. Vendor agents benefit in time because of an increased number of interactions due to having regular customers from the coalition. From the global point of view of the market system this is beneficial since it helps to bring closer compatible agents. This facilitates the search job of the matchmaker agents since they have to search in a smaller space (the coalition of the agent) first. As a result, the efficiency of the trading system is also increased.

We assume that coalitions are disjoint. Once created, they may have a long lifetime. A coalition is created when an agent wants to form a new coalition with another agent. The latter always agrees to forming a new coalition, since there is no cost for forming, joining, or leaving a coalition. A coalition dissolves when composed of only one agent. Any agent can join or leave a coalition at any moment. Agents in the same coalition agree with a specific fixed discount for each transaction executed. We assume the coalition structure (i.e. the partition of agents into coalitions) is global knowledge.

In our market a vendor's utility is represented by the amount of sales over time. Thus, it enters a coalition to increase these sales and it prefers to be part of the same coalition as customer agents with whom it has most transactions. As a consequence, a vendor joins a certain coalition to get closer to customers that have compatible preferences, by nurturing 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 (web-CRM) – 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.

A customer's utility consists of the amount gained from discounted transactions. It knows that being part of the same coalition as some vendor agent promises discounts from that vendor in future transactions. As a consequence it prefers to belong to the same coalition as vendor agents with whom it has had most transactions.

Belonging to long-term coalitions is beneficial for both vendors and customers since it helps increase their long-term utility. However, the most important problem for an individual agent is to find the most profitable coalition to be part of. This is a dynamic process since the agents' preferences and compatibilities may change over time. We model the individual agent's reasoning about coalition formation as a decision problem of finding the best coalition to be part of. For this we make use of the agent's past interactions with other agents in the system. The agent evaluates these interactions and updates the probability of having interactions ending with transactions with those agents.

The remainder of this paper is structured as follows: section 3 presents the general scheme of the proposed coalition formation mechanism; section 4 describes in more detail the decision problem that the agent faces about belonging to a certain coalition and our approach to solving this problem. Section 4 presents the experimental evaluation. Section 5 describes related work on coalitions. We draw conclusions in the last section.

3. The Decision Problem

This section presents the decision problem the agent faces when it has to choose the most promising coalition. For this it has to take its previous experiences into account, i.e. to learn. In the next two sub-sections we present two perspectives on the decision-making problem, a probabilistic formal one and a trust-based one and we show how trust can be related to formal models of learning.

3.1. Probabilistic Perspective

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. In this section, we discuss the assumptions behind a bounded-rational agent involved in the process of coalition formation. As mentioned above, an agent must decide at each epoch whether or not a change in coalition is to its best interests. It may choose to join an existing coalition (possibly by leaving an existing coalition first), to form a new coalition (with an agent not already in a coalition), 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. We also discard this approach since 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. We take the latter approach.

The agent's action directly affects the coalition structure, and the coalition structure affects the agent's profits. We can model the future effects of an agent's actions by modeling a number of future epochs. At one extreme, we could model only the current decision, and approximate the future by modeling the agent's expectations for the future. The other extreme would be to model all future epochs, and treat the situation as a stationary, infinite horizon decision problem. A compromise between these two extremes would be to choose a finite number of epochs, and treat the problem as a finite horizon multi-stage decision problem; the number of epochs modeled would depend on the computational resources available to the agent. However, it is not clear that a larger model will reflect her situation more accurately, as any model that does not model the other decision making agents' choices is already a fairly loose model! Our model is the "one-shot" decision model.

The value of being in a certain coalition depends on the nature of the future interactions of the agent. There is value to be gained from discounts available to buyers in the coalition, and to long-term or bulk sales for sellers. A simple model for the future is to assume a constant expected net gain for all successful future interactions. Therefore the agent's expected future reward depends only on the probability of successful interactions with other agents. Let p_{ij} be the probability of a successful interaction between agents i and j . By "successful interaction" we mean an interaction consisting of a successful negotiation (i.e. the agents reached an agreement about the price) and resulting in a satisfactory transaction (i.e. satisfactory quality of goods / service, no delays in payment or delivery). We assume that the success of an interaction between two agents does not depend on the outcome of previous interactions. A simple counting argument supports an estimate of:

$$\hat{p}_{ij} = \frac{n_{ij} + 1}{N_{ij} + 2} \quad (1)$$

In equation (1) n_{ij} is the number of successful interactions between agents i and j , out of a total of N_{ij} interactions between these agents (successful and unsuccessful). This estimate is also the maximum likelihood estimate for p_{ij} , based on a $Beta(1,1)$ prior, i.e., a uniform prior distribution over the success of an interaction. The estimate for \hat{p}_{ij} assumes that there are no changing success probabilities. In other words, agents don't change during their lifetimes, and the

probability of a successful outcome is unknown, but constant. Note that we can express \hat{p}_{ij} incrementally, using the previous estimate \hat{p}'_{ij} and the current event e_{ij} which is 1 if the current interaction was successful, and 0 otherwise:

$$\hat{p}_{ij} = \hat{p}'_{ij} \cdot \frac{N_{ij} + 1}{N_{ij} + 2} + \frac{e_{ij}}{N_{ij} + 2} \quad (2)$$

This formula in fact expresses the agent's learning how to estimate the probability of successful interactions with another agent based on its past experiences and evidence from the most recent interaction. It can be related to our current approach based on trust.

3.2. Trust-based Decision Making

An alternative way to solve the decision problem is to use an intermediate utility function that depends on the past transactions of the agent and on its future expectations. This function reflects the agent's opinion about the possibility of having successful interactions with another agent and we call it the "trust" which one agent has in another agent (Breban & Vassileva, 2002a, 2002b). The notion of trust has been an object of continuous interest in economics, sociology, and more recently in AI research. Trust relationships between agents in multi-agent systems are analyzed and modeled in task-delegation problems and in electronic commerce. For a more detailed survey of trust-based approaches, see (Breban & Vassileva, 2001). We represent trust function using the model proposed in (Jonker and Treur, 1999). The trust function can take a numerical value from a set of predefined quantifications. A new experience between two agents has a value from a predefined set of experience classes – i.e. it is evaluated to be either positive or negative with a particular strength. Each new experience leads to an update of the agents' trust in the other agent. We briefly present here the formal model described in (Jonker and Treur 1999).

Given a set of experience classes E and a set of predefined trust quantifications T , a mapping for the transition from one strength value t to another $trust(e, t)$ can be defined as:

$$\begin{aligned} trust : E \times T &\rightarrow T \\ trust(e, t) &= (1 - d) * e + d * t \end{aligned} \quad (3)$$

We consider the case in which an experience e 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 $T = [-1, 1]$ of predefined quantifications for trust t . 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 t is multiplied by d and the impact of e is added, normalized so that the result fits in the desired interval T .

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 (or learning) function $evol$ is inductively defined in (Jonker and Treur 1999). This function is used by the agent when it has to update its relationship with another agent at each step from the $Time$ set:

$$\begin{aligned}
 evol &: E \times Time \rightarrow T \\
 evol(e, 0) &= 0 \\
 evol(e_0 e_1 \dots e_i, i+1) &= trust(e_i, evol(e_0 e_1 \dots e_{i-1}, i)) \quad (4)
 \end{aligned}$$

The definition of the evolution of 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 the trust that each agent has in each other agent it has interacted with.

Interestingly, formula (3) in (Jonker and Treur 1999) is similar to (2) from our probabilistic decision model. In (Jonker and Treur 1999), d is a fraction in $[0,1]$ which limits the “memory” of the trust parameter. In their model trust is a measure in $[-1, 1]$. Without loss of generality, the scale of this measure can be normalized to the range of probabilities, $[0,1]$. If we take e to be 1 for a successful interaction, and 0 otherwise, we can see that the formulae for \hat{p}_{ij} and t are equivalent iff

$$d = \frac{1}{N_{ij} + 2} \quad (5)$$

In (Jonker and Treur 1999), d is taken to be constant, and in the range $[0, 1]$. However, the probabilistic formula has d decreasing with the number of interactions between agents i and j , and is in the range $[0, 0.5]$. Thus the probabilistic account puts more emphasis on the history of the agent, and the $trust$ account places more emphasis on the current interaction and the recent past interactions.

To make a decision about which coalition to join, the agent has to maximize an utility function that reflects the agent’s trust in a whole coalition, as a set of individual agents in whom the agent may have established a trust already. The trust of an agent in a coalition 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 as a function of its trust with the individual agents in this coalition. There are many possible ways to define this function. We implemented and evaluated three different functions.

The first function (called *soc1* strategy) defines the agent's trust in a coalition by calculating the sum of the trust that the agent has in each of the individual agents from that coalition. From a probabilistic perspective, the agent chooses to join the coalition that has the highest probability of successful interaction. The second one (called *soc2*) computes the number of all agents in the coalition with whom the agent has positive relationships. From a probabilistic perspective, the agent chooses to join the coalition that contains the highest number of agents with a probability of successful interactions greater than 0.5. The third one (called *ind*) finds the coalition where the most trusted agent is. From the probabilistic perspective, this strategy is based on the consideration of maximizing the probability of successfully interacting with any single agent in all coalitions. Interestingly, our experiments showed that though these two functions appear to be similar, they lead to different agent behaviors.

For the case when the agent uses *soc1* strategy let p_{iC} be the probability of a successful interaction between agent i and an agent who is a member of coalition C . We can estimate this parameter with the following:

$$\hat{p}_{iC} = \frac{\sum_j n_{ij}}{N_{iC}} \quad (6)$$

where N_{iC} is the number of interactions that have taken place between agent i and agents in coalition C . The agent chooses the coalition with the highest probability calculated using (6).

For the case when the agent uses *ind* strategy, it has to maximize

$$\hat{p}_{iC} = \frac{\max_j n_{ij}}{N_{iC}} \quad (7)$$

The trust-based approach is not significantly different from the probability estimate presented in the section 3.1. The significant difference is only in the choice of values for d and e in (3). Both approaches are based on counting previous successes and creating a normalized representation of the past experience. Either of these two approaches can be used in the design on the decision-making algorithm of the agent. In the next section we present the algorithm for decision-making using the trust-based mechanism described above.

3.3. Trust-based Decision Algorithm

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 $C = \{C_1, C_2, \dots, C_n\}$, where C_i is the label of the coalition to which A_i belongs. An agent A_i has trust relationships with k agents (a subset of A) that we represent as vector $R(A_i)$:

$$R(A_i) = (A_{i1}, A_{i2}, \dots, A_{ik}) \quad (8)$$

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

$$T(A_i) = (T_{i1}, T_{i2}, \dots, T_{ik}) \quad (9)$$

where T_{ij} is the trust of the agent A_i in A_{ij} . $T_{i1}, T_{i2}, \dots, T_{ik}$ take values from the set of trust quantifications T . The corresponding coalitions of these k agents form the set of coalitions $C(A_i)$:

$$C(A_i) = (C_{i1}, C_{i2}, \dots, C_{ik}) \quad (10)$$

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 (if it is in a coalition).

The decision making phase of the Coalition Reasoning Mechanism assumes that an agent knows the coalition it belongs to (C_i), its trust relationships with other agents ($T(A_i)$), and the coalitions in which these agents belong ($C(A_i)$). Based on these notation we can describe in pseudocode a rule-based algorithm for the decision problem associated with the *ind* strategy:

Find A_{ip} - the most trusted agent by A_i
if (A_i in a coalition) AND (A_{ip} in a different coalition) then
 A_i leaves the current coalition
if (A_{ip} is in a coalition) then A_i joins this coalition
 Else A_i forms a new coalition with A_{ip}

In the socially oriented strategies (*soc1* and *soc2*) the agent has to build a representation of its trust in each coalition based on the vector of trust relationships with individual agents ($T(A_i)$). The new vector is denoted by $TC(A_i)$:

$$TC(A_i) = (TC_{i0}, TC_{i1}, \dots, TC_{im}) \quad (11)$$

where TC_{ij} represents the trust that A_i has in coalition j . In the *soc1* strategy trust in coalition k is calculated as:

$$TC_{ik} = \sum_j (T_{ij}) \quad (12)$$

where agent A_{ij} is in coalition k ($C_{ij} = k$) and k varies from 1 to m . In the *soc2* strategy the same trust is:

$$TC_{ik} = | R'(A_i) | \quad (13)$$

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 trust relationships. An agent A_{ij} is in $R'(A_i)$ if and only if $C_{ij} = k$ and $T_{ij} > 0$. 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 *soc1* and *soc2* strategies:

Build TC(A_i) - the vector of trust relationships in coalitions

Find coalition k most trusted by A_i

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 group of agents from outside of coalitions that it trusts most

We analyze in the next section the proposed coalition formation mechanism by comparing the results of these three agent strategies (*ind*, *soc1*, and *soc2*). We cannot compare the proposed mechanism with any other one, since there is no previous work directly comparable. No prior coalition formation mechanism provides a description at the agent level.

4. Evaluation

We have developed a simulation prototype of the proposed trust based coalition reasoning mechanism in Java. We ran 9 sets of experiments with different configurations of parameters. Each set of experiments consisted of 100 trials over which the results were averaged. Our goal was to evaluate the mechanism at macroscopic (system) and microscopic (agent) 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 level evaluation we focused on the individual gains of the customer agents (calculated as the average of all benefits obtained from discounted transactions by each customer). The experiments were intended to compare the different agent strategies described above (*ind*, *soc1*, and *soc2*).

Some variables involved in the design of the mechanism were set constant for all experiments: the inflation rate of trust ($d = 0.5$), the evaluation of positive experiences (0.2), the evaluation of negative experiences (-0.2), and the discount rate (5% of the value of the books being traded). The parameters under investigation are summarized in Table 1. The number of vendor agents was set to 100, while the number of customer agents was varied to 100, 1000, and 10000. A significant parameter for the evolution of different factors over time is the number of interactions between agents. We can control it easier than the time periods in which a certain number of interactions happen. This depends entirely on the characteristics of the market.

number of customers	100; 1000; 10000
number of vendors	100
number of interactions	1; 100; 1000; 10,000; 100,000; 1,000,000
agent strategy	<i>ind; soc1; soc2</i>

Table 1. Simulation parameters and their values

We plot in Figure 2 our results: the evolution of the number of coalitions is shown in graph A, the evolution of the system dynamics – in graph B, and the evolution of the individual gain – in graph C. On the X-axis of each graph the number of interactions is represented on a logarithmic scale from 1 to 1,000,000; the Y-axis represents the parameter under investigation. In Figure 2A we use a scale from 1 to 100 coalitions; in Figure 2B we show the system dynamics on a scale from 0 to 200 thousands; in Figure 2C the individual gain varies from 0 to 15 thousands.

The results in Figure 2A show that as we increase the number of interactions among agents in the system, the number of coalitions first grows, it reaches a peak, and then it starts to decrease. The meaning of this behavior is that at the beginning 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 (reached after 1 million interactions) that reflect the merging rate of the coalitions. Both values are shown in Table 2. We note that the number of coalitions has a predictable and controllable evolution over time. It is limited by the number of vendor agents in the system. There are small differences among the three agent strategies. Significant is that with *soc2* the number of coalitions drops faster to 1, with *ind* it reaches 1 in fewer cases, while with *soc1* it never reaches 1, although the number of coalitions seems to stabilize at small values.

Configuration	Peak #Coalitions			Lowest #Coalitions			System dynamics		
	<i>ind</i>	<i>soc1</i>	<i>soc2</i>	<i>ind</i>	<i>soc1</i>	<i>soc2</i>	<i>ind</i>	<i>soc1</i>	<i>soc2</i>
100C/100V	34	32	33	2	5	1	10	44	0
1000C/100V	89	88	87	11	7	1	23	52	3
10000C/100V	98	98	98	35	93	1	128	82	30

Table 2. Results from the macroscopic analysis.

Our second evaluation factor – shown in Figure 2B – is the system dynamics defined as the sum of coalitions visited by each agent. Generally, the dynamics is insignificant for small values of the number of interactions (from 1 to 1000) and has a slow increase for bigger values (up to 100,000). When the number of interactions grows further (to 1 million) the dynamics presents either a higher increase (in several cases) or stabilization (in the other cases). The higher increase is observable in configurations using the *ind* and the *soc1* strategies. In configurations using the *soc2* strategy the dynamics stabilizes between 100,000 and 1 million interactions. For easier comparison, we show in Table 2 the highest dynamics (reached after 1 million interactions). Comparing the values of the system dynamics for different agent strategies, we observe that *soc2* results in the lowest dynamics, followed by *ind* and, lastly, by *soc1* (1 exception applies).

C)

Figure 2. Results showing: A) the number of coalitions; B) the system dynamics; C) the individual gain.

The individual gain – evaluated for the agent level of our analysis and shown in Figure 2C – calculates the average of the sum of all discounts that each customer receives. The general trend of the curves is to grow very slowly from 1 to 10,000 interactions, after which they rise slowly until 100,000 interactions and highly until 1 million interactions. Comparing the values of the individual gains for different agent strategies we note that they are higher for *soc2*, medium for *ind*, and lower for *soc1* (some exceptions apply). An explanation for this is that when the agents use the *soc2* strategy the number of coalitions drops faster to its lower limit, fact that increases a customer's chances to interact with a vendor from the same coalition and get a desired discount. The results of the individual gain reveal once again that *soc2* is the most beneficial strategy, while *soc1* brings lower benefits for customers.

We observe that the individual gain is directly related to the merge of coalitions. It is inversely proportional to the number of coalitions. In all strategies the gains become higher only after the coalitions start to merge (after 10,000 interactions). Three distinct behaviors are apparent:

- In all cases with *soc2* and several with *ind* the number of coalitions stabilizes to 1 the system dynamics stabilizes and the individual gain highly increases. This is a direct consequence of agents finding compatible partners faster and stabilizing in certain coalitions.
- In the remaining cases (all with *soc1* and the rest with *ind*) the number of coalitions is dropping slowly to its lower limit, the dynamics is highly increasing, and the individual gain is increasing slower. This means that it is harder for the agents to find compatible partners and most profitable coalitions.

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 for all strategies. It also brings stability to the system since the overall dynamics reaches an equilibrated state for the *soc2* strategy. Although the system dynamics highly increases for the other two strategies (i.e. *ind* and *soc1*), 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 exponential increase in the individual gains of customers. *Soc2* is the best strategy for reasons of stability (least dynamics) and utility (best gain). *Ind* is better than the *soc1* strategy. More details about the evaluation as well as the full set of experimental results are given in (Breban, 2002).

6. Conclusion

In this paper, we proposed a coalition formation mechanism for software agents on an electronic marketplace. We discussed two perspectives of the decision making problem that an individual agent faces, a probabilistic and a trust-based one and we implemented and evaluated the trust based one. The mechanism is described on a microscopic level (i.e. the reasoning of an individual agent) and it has been evaluated on a macroscopic level by running a simulation with over 10000 agents using this reasoning mechanism. We showed that this mechanism brings stability to the system (in the number of coalitions and in the overall dynamics) and provides the customer agents with increased benefits over time. The mechanism uses reduced communication between the agents that makes it scalable for large numbers of agents and interactions. Future work includes investigation of the proposed coalition formation mechanism under more realistic circumstances such as allowing more types of goods to be traded in the system and giving the agents the liberty to belong to more coalitions at a time.

REFERENCES

- Breban S. Long-Term Coalitions for the Electronic Marketplace. M.Sc. Thesis, University of Saskatchewan, February 2002. (available from <http://bistrica.usask.ca/madmuc/>)
- Breban S., J.Vassileva (2001) Long-Term Coalitions for the Electronic Marketplace, in B. Spencer (ed.) Proceedings of the E-Commerce Applications Workshop, Canadian AI Conference, Ottawa, pp. 6-12.
- 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.
- Breban S., J.Vassileva (2002b) A Coalition Formation Mechanism Based on Inter-Agent Trust Relationships. To appear in L. Johnson & C. Castelfranchi (eds.) Proceedings of the First Conference on Autonomous Agents and Multi-Agent Systems, Bologna, 17-19 July 2002, ACM Press.

- Jonker C., J. Treur (1999) Formal Analysis of Models for the Dynamics of Trust based on Experiences Autonomous Agents, Proceedings of the Workshop on Deception, Fraud and Trust in Agent Societies, Autonomous Agents'99 Conference, Seattle, pp. 81-94.
- Lermann K., O. Shehory (2000) Coalition Formation for Large Scale Electronic Markets. Proceedings of the Fourth International Conference on Multiagent Systems ICMAS'2000, Boston, pp. 216-222.
- Mudgal C., J. Vassileva (2000) 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, pp. 107-118.
- Shehory O., S. Kraus (1996) A Kernel-oriented model for coalition formation in general environments: Implementation and Results. Proceedings of the Thirteenth National Conference on Artificial Intelligence, Portland, pp. 130-144.
- Shehory O., K. Sycara (1997) Multi-agent Coordination through Coalition Formation. Proceedings of International Workshop on Agent Theories, Architectures, and Languages (ATAL97), Providence, RI, pp 135-146.
- Tsvetovat M, K. Sycara (2000) Customer Coalitions in the Electronic Marketplace. Proceedings of Autonomous Agents'2000, Barcelona, pp. 263-264.
- Web-CRM: Customer Relationships Management, available on line at <http://www.crm.com>
- Yamamoto J., K. Sycara (2001) A Stable and Efficient Buyer Coalition Formation Scheme for E-marketplaces. Proceedings of Autonomous Agents'2001, Montreal, pp.549-556.