

Using Inter-Agent Trust Relationships for Efficient Coalition Formation

Silvia Breban and Julita Vassileva

University of Saskatchewan, Computer Science Department,
Saskatoon, Saskatchewan S7N 5A9, Canada
{svb308, jiv}@cs.usask.ca

Abstract. We address long-term coalitions that are formed of both customer and vendor agents after evaluating their trust relationships with other agents in the system. We present a coalition formation mechanism designed at agent level and analyzed at both system and agent levels. Evaluation has been conducted to compare agent strategies (individual vs. social) and to analyze the system behavior under different circumstances. 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 increase exponentially in time).

1. Introduction

Coalition formation in multi-agent systems has been seen in the game theory [1] and distributed AI [2] as 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. Recent research brings the coalition formation process into the electronic marketplace as a mechanism of grouping customer agents with the intent of getting desired discounts from the vendor agents in large size transactions. In this context the definition of the term *coalition* means a group of self-interested agents (with no social or group rationality requirements) that are better off as parts of the group than by themselves.

Coalition formation mechanisms have been proven to be beneficial for both customer and vendor agents by several studies [3, 4]. The already existing Internet communities like newsgroups, chat-rooms, and virtual cities constitute examples for the potential of creating large-scale economies among similar minded customers that can be explained by the high value (or usefulness) of networks that allow group formation. Such groups are known as Group-Forming Networks (GFN). In general, the value of a network is defined [5] as the sum of different access points (users) that can be connected for a transaction for any particular access point (user) when the need arises. There are three categories of values that networks can provide: a linear value, a square value, and an exponential value. The linear value is derived from the Sarnoff's Law that specifies that the power of a broadcasting network is linearly increasing in proportion to the number of its users. It characterizes networks that provide services to individual users like TV programs or news sites. The square value is derived from

the Metcalfe's Law that states that the value of a peer-to-peer network equals the square of the number of its users. It is applicable to networks that facilitate transactions such as commercial sites or telephony systems and it has been used by economists as an explanation for the fast growth of the Internet. Reed [5] finds that the networks that allow group affiliation are even more powerful. According to the GFN Law that he promotes, the value of such networks grows exponentially with the number of users. As a consequence, networks that allow group formation among its components (users or agents) are expected to bring the highest benefit in the future.

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. No prior work addresses the issue of trust among members of the same coalition. Therefore, we propose a new approach to forming coalitions that have a long lifetime and are based on trust among agents. Our primary goal is to provide a coalition formation mechanism that takes into consideration the long-term utilities of individual agents, their preferences, and the trust they have in their partners. Secondly, the proposed mechanism is designed to accommodate large numbers of agents (thousands and millions) due to minimized communication between agents and reduced complexity. Finally, for reasons of system stability and predictability our approach has two more objectives: to reduce the dynamics of the agents and to stabilize the number of coalitions in the system.

The remainder of this paper is structured as follows: the next section describes the concept of long-term coalitions that are based on trust relationships between agents. It also presents the existing formal model that we use for representing trust at the agent level. Section 3 shows how this model is integrated in the coalition formation mechanism. Detailed description of the mechanism from an agent's perspective is provided. Section 4 presents the evaluation we conducted at both system and agent levels. We draw conclusions and present future directions in the last section.

2. Long-term Coalitions Formed on Trust Relationships

We address a multi-agent system composed of selfish agents that have fixed roles of either customer or vendor and trade books on the Internet. Agents can form coalitions with the intent of increasing their individual benefits. This improves their cooperation and coordination and, as a result, the efficiency of the trading system.

We make use of several assumptions in our design:

- Agents have individual rationality and no social or group requirements. They try at each moment to maximize their individual long-term utility function.
- 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.
- An agent can join or leave a coalition at any moment. This is a direct consequence of the realistic assumption that an agent's preferences can change over time.
- Agents may have different interests in books being traded (e.g. science fiction, romance, history). They may belong to different economic categories (e.g. a customer agent that can afford to buy only books between \$20 and \$60 and a vendor that would not sell for less than \$100 belong to different categories).

- A coalition is automatically created when an agent wants to form a new coalition with another agent; it dissolves when it is composed of only two agents and one of them leaves the coalition.
- Coalitions are formed of both vendors and customers.
- Coalitions have a long lifetime once they are created.
- Agents in the same coalition agree with a specific discount for each transaction executed.
- Agents may have to pay a penalty (cost) for leaving a coalition.
- Agents prefer to be part of the coalition in which they expect to have most future transactions.
- Membership to coalitions is global knowledge.
- Coalitions are disjoint.

The main reason for forming coalitions is to bring agents with compatible preferences closer, by nurturing vendor-customer relationships. The concept is similar to the established practice in real-life markets like SafewayTM or SearsTM that give a minimal discount to members of their clubs. For customers, the member card – and belonging to the members' club – represents an appreciation of the trustworthiness of the store as well as a preference for dealing with that particular merchant, based on positive previous trades and satisfaction with goods purchased, their quality and price. The vendor agrees to give a specific discount to its club members as a sign of appreciation of their long-term relationship. Such a policy 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. It is known in the field of Economics as Customer Relationship Management (CRM) [6]. CRM is also called relationships marketing or customer management. It is concerned with the creation, development, and enhancement of individualized customer relationships with carefully targeted customers and customer groups that result in maximizing their total customer lifetime value. The large adoption of CRM in real-life market systems is being fuelled by a recognition that long-term relationships with customers lead to improved customer retention and profitability. We use it as the basis for motivating agents and their users to enter long-term coalitions.

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. A customer agent knows that being part of the same coalition as some vendor agent will bring it 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 because this promises to bring it most discounts and increased profit. We impose to agents the restriction to be part of one coalition at a moment for reasons of decreased complexity of the mechanism. In the alternative case of allowing an agent to belong to more than one coalition, any customer agent would prefer to belong to all existing coalitions in the system. Thus, all customers would want to be part of each coalition. The existence of coalitions in such a case would become futile: the effect of creating a small, comfortable environment that brings compatible agents closer to each other to have more frequent interactions and, therefore, increased benefits for both customers and vendors would be lost. For vendor agents giving discounts to all customers would also be unacceptable. The CRM policy cannot function when a preferred group of customers is not distinguished from the others.

From the point of view of an individual agent (either customer or vendor), we see the coalition formation mechanism as a decision problem: at each moment the agent faces a decision of whether to remain in the same status, form a new coalition, or leave the current coalition to join a different one. The decision should maximize the agent's long-term utility function after taking into account important factors such as its long-term goals, incomplete knowledge about other agents, and global knowledge about the system. To help an agent make the right decision our approach uses the relationships established between the agent and other agents in the system after sharing common experiences. When faced with the decision problem of whether to join or leave a coalition, the agent has to evaluate all these relationships using some evaluation criterion. Based on this evaluation and on the natural assumption that an agent expects to have more profit from compatible agents with whom it has the best relationships, the agent prefers to be part of the same coalition as its partners with whom it has stronger relationships.

In general, relationships between individuals can model different aspects of their interaction: the roles they play in the interaction, their goals, the importance that the interaction has for each of them, and the trust they have in one another [7]. In the absence of a formal contract between the agents, we find that the most appropriate aspect of a relationship is the trust that agents have in each other. In the context of formal contracts among agents in a group, there is an implicit trust in the structure and the regulations of the system that needs no explicit specification. In the context of informal contracts, each agent in the group should be able to explicitly trust the other agents.

Trust has been thoroughly studied in e-commerce applications over the last years. Different definitions of trust can be found in [8, 9]. In our view the trust of an agent in another agent represents its belief that the other agent has similar preferences and this will lead to many successful transactions between them in the future. For instance, a customer that has had satisfactory transactions with a certain vendor trusts that vendor to promise beneficial transactions for the future. A vendor that is satisfied with the purchases of a customer also believes that it will have positive transactions with that customer in the future. To represent trust at the agent level we use the model proposed in [8] that assigns to each trust relationship of an agent a numerical value from a set of trust quantifications. A new experience between the truster and the trusted 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. It leads to an update of the agents' trust according to a transition function defined between different states of trust.

We briefly present here the formal model described in [8]. Given a set of experience classes E and a set of trust quantifications T , a mapping for the transition from one trust value t to another trust value $trust(e,t)$ can be defined as:

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

We consider the case in which $E = [-1, 1]$. 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 $T = [-1, 1]$. Parameter $d \in [0.5, 1]$ 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 agents' preferences may change over time). In this trust function after each new experience e , the existing trust value 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 $Time = N^+$ (the set of natural numbers), a trust evolution function $evol$ is inductively defined in [8]. This function is used by an agent when it has to update its trust in another agent at each step from the $Time$ set:

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

The definition of the trust evolution specifies that the initial trust for step 0 is set to a neutral value 0. At each step $i+1$ the trust is updated based on the previous trust (from step i) and the current experience e_i according to the trust function defined above. We use this formal model of subjective trust to represent inter-agent relationships. The following section describes how we integrate it in the agent reasoning mechanism about coalition formation.

3. Coalition Formation Mechanism

We refer to a system of multiple personal assistant agents trading books in an open electronic market. Both the set of customer agents and the set of vendor agents may have size variations over time. We assume that the electronic market provides system matchmaker agents that are responsible for finding suitable partners to interact when a need arises. Agents can form coalitions to improve their individual benefits. The general scheme for coalition formation is shown in Fig. 1. It consists of two important steps. The first step is the interaction between a customer and a vendor. We use the term *interaction* to denote any attempt to execute a transaction between two agents. This step starts with a negotiation for a specific price, it continues with an evaluation of the interaction, and it ends with an update of the trust that each agent has in the other agent. The second step is named Coalition Reasoning Mechanism. It develops in each of the two agents' reasoning mechanism after the interaction is finished. In this step an agent performs a re-evaluation of its status of belonging to the most profitable coalition or not. It also decides what action to take if any is needed. We present each component phase of this general scheme in more detail.

As mentioned above, 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. We use the negotiation mechanism developed by Mudgal [10]. It 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; it can end with an agreement in which case the interaction is *successful* or with a rejection from one of the agents in which case the interaction is *unsuccessful*. The users' preferences play a crucial role in the result of the negotiation (either agreement or rejection). They take into consideration the minimum acceptable price for vendors and the maximum affordable price for customers, the subjective importance of money, the urgency of the current goal of selling or purchasing a certain product, the behavior towards risk, and the time constraints for executing the transaction. These preferences are relevant for establishing a possible compatibility or incompatibility between the agents. For instance, if the vendor has a minimum acceptable price p_1 for a certain product and the customer has a maximum affordable price p_2 for that product, and $p_1 > p_2$, the two agents will never reach an agreement when negotiating on that product. We

consider such cases that lead to rejection on either side to be negative evidence for updating the trust that agents have in one another, since it reveals a certain level of incompatibility between their preferences in ranges of prices.

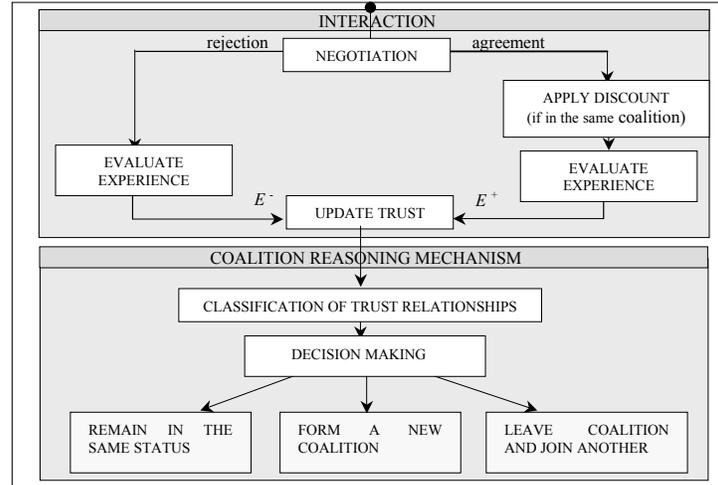


Fig. 1. General scheme for the coalition formation mechanism

When the negotiation terminates with an agreement, the agents agree on a price. If they belong to the same coalition, a certain discount is applied to this price as shown in Fig. 1. After the negotiation is finished, both the customer and the vendor have to evaluate the interaction. We consider a successful interaction as a positive experience because reaching an agreement between two agents is a direct consequence of similar interests in books and compatible preferences of the users. As mentioned above unsuccessful interactions reveal a possible incompatibility between the agents. We consider this as negative evidence for their belief that they will have successful transactions in the future. Positive experiences are assigned values from E^+ , while negative experiences are evaluated in the negative subset of experience classes E^- .

The last phase that takes place in the interaction step is the trust update. When the new experience is evaluated, the trust that the two agents have in each other is updated according to the trust evolution function defined in the previous section. Each agent stores in a vector the representation of its trust relationships with all agents in the system with whom it has ever interacted. The agent's relationships with the other agents are null according to the definition of the trust evolution function. A relationship is represented by the name of the agent to be trusted and a specific value from the set of trust quantifications T .

The update of trust closes the interaction between the vendor and the customer. It also triggers the second step of the coalition formation, namely the Coalition Reasoning Mechanism (shown in Fig. 1). This consists of two parts: first the agent evaluates all its trust relationships with agents with whom it had previous interactions, and second, it makes 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.

The evaluation of trust relationships consists in classifying them and finding the best one. This can be done using different agent strategies: individually oriented and

socially oriented. With the individually oriented strategy – that we denote with *ind* – an agent prefers to be in the same coalition with the agent in whom it has most trust. With the socially oriented strategy the current agent prefers the coalition that it trusts most. Trust in a specific coalition can be calculated as a function of the trust in individual agents from that coalition in which only agents that have a history of experiences with the current agent are significant. We consider the special case of agents out of coalitions as forming coalition 0. We present two possible functions to calculate the relationship established between an agent and a coalition. The first function calculates the relationship as the summative trust in all agents from that coalition. This leads to an agent strategy that we denote by *soc1*. The second function computes the trust of an agent in a coalition as the number of all agents in the coalition with whom the agent has positive relationships (i.e. the number of trustworthy agents). It leads to a different agent strategy that we name *soc2*.

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. We consider as public global knowledge the coalition in which each agent in the system is. Our solution for the decision problem of an agent using the *ind* strategy can be described in pseudocode as a rule-based algorithm. The algorithm ensures that the current agent belongs to the same coalition as its most trusted partner. It finds first this agent. If they are in the same coalition the current agent does not change its status. If it is in a different coalition the agent leaves its coalition. In this latter case if the most trusted agent is in a coalition the current agent joins it, otherwise they form a new coalition.

```

Find  $A_k$  - the most trusted agent by  $A_i$ 
if ( $A_i$  and  $A_k$  in the same coalition)
     $A_i$  does not change its status
elseif ( $A_i$  in a coalition) AND ( $A_k$  is not in  $A_i$ 's coalition) then
     $A_i$  leaves current coalition
if ( $A_k$  out of coalitions) then
     $A_i$  forms a new coalition with  $A_k$ 
elseif ( $A_k$  in a coalition)
     $A_i$  joins  $A_k$ 's coalition

```

The decision-making is the same for both socially oriented strategies. We present a solution for it in the form of a rule-based algorithm that can be used by an agent using the *soc1* or *soc2* strategies after it calculates its trust in coalitions. The agent finds the most trusted coalition first. If it belongs to this coalition no action is needed; if it is in a different coalition it leaves its current coalition. As a result the current agent is out of coalitions. It forms a new coalition with agents that it trusts most from out of coalitions or it joins the most trusted coalition if it exists.

```

Find coalition  $k$  most trusted by  $A_i$ 
if ( $A_i$  in coalition  $k$ ) AND ( $k$  different than coalition 0)
     $A_i$  does nothing
elseif ( $A_i$  in a coalition) AND ( $k$  different than coalition of  $A_i$ )
     $A_i$  leaves current coalition
if ( $k$  is coalition 0)
     $A_i$  forms new coalition
elseif ( $k$  is not coalition 0)
     $A_i$  joins coalition  $k$ 

```

We analyze the three agent strategies presented in this section and compare their effects on the system and on the individual agents in the following section.

4. Evaluation

We have developed a simulation prototype of the proposed coalition formation mechanism in Java. We ran 54 sets of experiments with different configurations of parameters. Each set of experiments consisted of 100 trials over which the results were averaged. In our simulation all agents used the same coalition formation mechanism. Our goal was to evaluate the mechanism at system and agent levels. For the first part we investigated the number of coalitions in the system, the overall dynamics, and how these factors evolve in 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 agent 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 in the previous section (*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. Note that the number of vendor agents was varied to 100 and 1000 while the number of customer agents was varied to 100, 1000, and 10000. Instead of time, we use the number of interactions between agents as parameter accounting for the evolution of different factors over time the time periods in which a certain number of interactions happen can vary significantly.

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

Table 1. Simulation parameters and their values

We considered three different setups for the coalition formation mechanism: a simple one (denoted as *simple*), a setup that accounts for the increased likelihood that agents interact more often with members of their own coalition (denoted as *prob*), and a setup in which agents pay a cost for leaving a coalition (called *costs*). In the *prob* setup the implementation is designed to choose with a 0.6 probability a vendor from the same coalition as the customer agent interested in buying a book and with a 0.4 probability a vendor that is not in the same coalition as the customer. In the other setups the probability of choosing an agent from the same coalition is 0.5. The *costs* setup takes into account the realistic assumption that a customer has to spend time and effort to find better vendors when leaving a coalition, while a vendor that leaves a

coalition might face a decrease in sales by losing its former clients. For vendors the costs are seen as a threshold of trust: given that trust is represented on a scale from -1 to 1, a vendor leaves a coalition only if the trust in the new coalition it wants to join exceeds the trust in the current coalition by 0.2. A customer pays a constant penalty (10 price units) for leaving a coalition to cover expenses for searching for new coalitions.

We plot first the results that show the evolution of the number of coalitions (see Fig. 2). 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 number of coalitions. Graphs a, c, and e (on the left) show configurations with 100 vendors for the *simple*, *prob* and *costs* setup. We used a scale from 1 to 100. Graphs b, d, and f (on the right) show configurations with 1000 vendors for the *simple*, *prob* and *costs* setup. We used a scale from 1 to 1000.

The results in this figure 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. In most cases when costs are involved the decrease in the number of coalitions is merely observable. The meaning of this behavior is that at the beginning agents form coalitions and after a while they start merging (faster when no costs are involved and slower when costs are involved). We focus on analyzing the peak values that reflect the formation of coalitions and the values reached after 1 million interactions that reflect the merging rate of the coalitions.

The peak values are reached between 100 and 10000 interactions (faster for fewer vendors and for fewer customers). They range from 32 to 99 in configurations with 100 vendors and from 90 to 954 in configurations with 1000 vendors. Comparing the setups, we observe that the peak values are higher in the *prob* setup, medium in the *costs* setup, and lower in the *simple* setup (some exceptions apply). We notice that these values are similar for the three agent strategies. As a conclusion we note that the peak values are limited by the number of vendor agents, increase with the number of customer agents, and depend on the setup and, slightly, on the strategy.

As we increase the number of interactions to 1 million, the number of coalitions drops to small values shown in Table 2. These values range from 1 to 99 in configurations with 100 vendors and from 1 to 909 in configurations with 1000 vendors.

Configuration	Simple			Prob			Costs		
	<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	2	5	1	1	11	1	50	5	50
1000C/100V	11	7	1	2	25	1	91	90	91
10000C/100V	35	93	1	95	99	99	99	99	99
100C/1000V	1	7	1	2	20	1	90	7	75
1000C/1000V	1	6	1	16	15	1	500	481	498
10000C/1000V	78	15	1	124	105	1	909	908	909

Table 2. The values of number of coalitions after 1 million interactions

We note an obvious difference in the evolutions of the number of coalitions from the *costs* setup and from the other two setups. While in the *simple* and the *prob* setups the number of coalitions drops in all configurations to small values before 1 million interactions, in the *costs* setup it presents a drop in only 2 out of 18 configurations (100 customers with 100 and 1000 vendors using *soc1*). In the other 16 out of 18 configurations from the *costs* setup the number of coalitions has a very slow drop

until 1 million interactions. The fact that introducing costs for agents leaving a coalition leads to a pronounced decrease in the merging rate of coalitions can be explained by agents becoming less willing to leave a coalition after they join it. We expect them, however, to start merging after more interactions. The *prob* setup also delays the formation of coalitions and their merging speed, but to a much lower degree than the *costs* setup. Comparing the agent strategies we notice that for the *simple* and *prob* setups in all configurations with *soc2* the number of coalitions drops faster to 1; with *ind* it reaches 1 in fewer cases; with *soc1* it never reaches 1. In the *costs* setup the number of coalitions never drops to 1, but to small values.

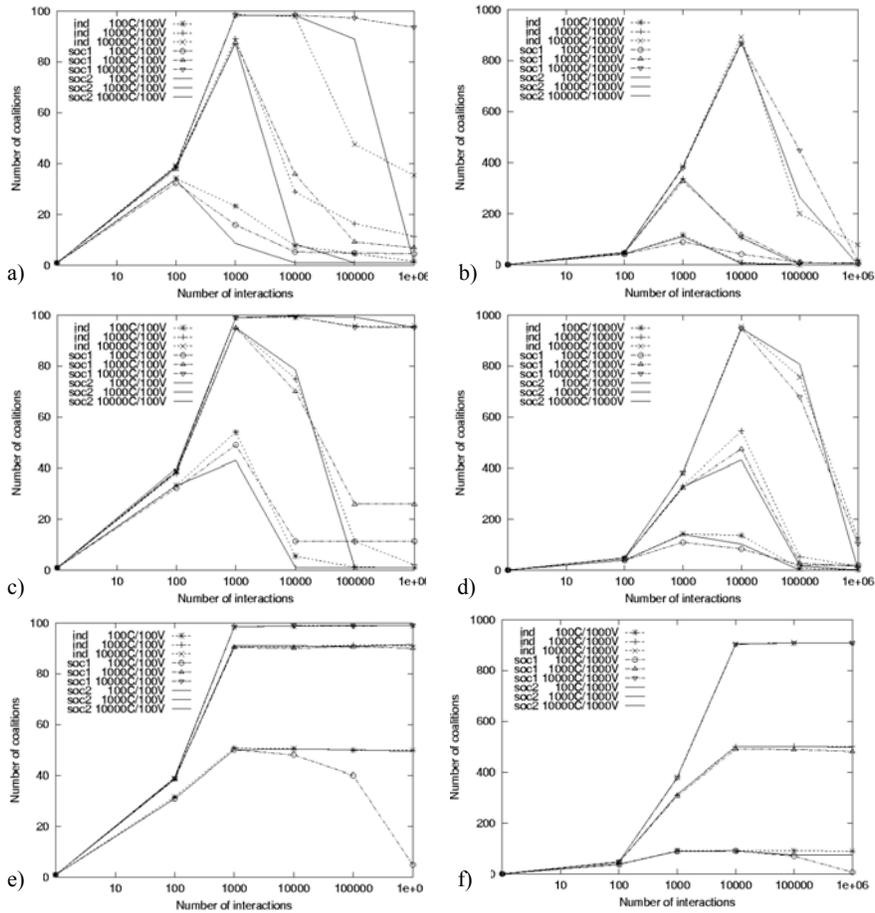


Fig. 2. Number of coalitions for: a) 100 vendors *simple*; b) 1000 vendors *simple*; c) 100 vendors *prob*; d) 1000 vendors *prob*; e) 100 vendors *costs*; b) 1000 vendors *costs*.

To conclude our analysis of the number of coalitions we note that it has a predictable and controllable evolution over time. The number of coalitions is limited by the number of vendor agents in the system; this led us to using different scales for

configurations with 100 and 1000 vendors. Another interesting observation is that the curves for configurations with the same proportion between the number of vendors and the number of customers are the same (e.g. 100 customers with 100 vendors and 1000 customers with 1000 vendors as well as 100 customers with 1000 vendors and 1000 customers with 10000 vendors). In these cases the proportion between the existing number of coalitions and their upper limit is the same in all sample points. A useful consequence can be drawn from this: when dealing with large numbers of vendors and of customers we can easily simulate the experiments for smaller numbers with the same proportion between the two numbers and scale up the results. 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. When agents join the coalition with the highest number of trusted agents (*soc2* strategy) and no costs are involved the number of coalitions tends to drop faster to 1.

Our second evaluation factor is the system dynamics defined as the sum of coalitions visited by each agent. In Fig. 3 we display similar plots: a, c, and e (on the left) for configurations with 100 vendors using the *simple*, *prob* and *costs* setup; b, d, and f (on the right) for configurations with 1000 vendors using the *simple*, *prob* and *costs* setup. The system dynamics is represented in thousands on the Y-axis (from 0 to 200) as a function of the number of interactions shown on a logarithmic scale from 1 to 1,000,000. Generally, the dynamics is insignificant for small values of the number of interactions (from 1 to 1000). It shows a slow increase when the interactions grow to 10,000 and then to 100,000. When they grow further (to 1 million) the dynamics presents an exponential increase in several cases or stabilization in the other cases. The exponential increase is observable in configurations using the *ind* and the *soc1* strategies with 1000 and 10000 customers in the *simple* and *prob* setups. In configurations using the *soc2* strategy the dynamics stabilizes between 100,000 and 1 million interactions in the *simple* and the *costs* setups for all configurations. In the remaining cases the system dynamics has a linear increase.

For easier comparison, we show in Table 3 the highest dynamics reached after 1 million interactions. They range from 0 to 183 (for 100 vendors) and from 2 to 172 (for 1000 vendors).

Configuration	<i>Simple</i>			<i>Prob</i>			<i>Costs</i>		
	<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	10	44	0	1	76	0	3	25	0
1000C/100V	23	52	3	12	133	4	2	5	1
10000C/100V	128	82	30	182	183	45	10	10	10
100C/1000V	3	52	3	11	116	4	2	13	2
1000C/1000V	8	35	8	13	87	8	2	4	4
10000C/1000V	128	110	35	81	172	45	11	11	11

Table 3. The values of system dynamics after 1 million interactions

Comparing the values of the system dynamics for different setups we observe that they are higher in the *prob* setup, medium in the *simple* setup, and lower in the *costs* setup for all configurations. In the *costs* setup the dynamics is drastically reduced compared to the other two setups. The *prob* setup brings an increase in the dynamics in 15 out of 18 cases compared to the *simple* setup. Therefore, introducing costs stops agents from moving from one coalition to another, while considering higher probabilities for agents to interact within the same coalition increases their dynamics.

Comparing the agent strategies, we observe that *soc2* results in the lowest dynamics, followed by *ind* and, lastly, by *soc1* (some exceptions apply). We also note that the system dynamics has similar values for configurations with 100 and 1000 vendors, but it has higher values when the number of customers is increased. This means that customer agents tend to move more from one coalition to another, since they account for the overall dynamics more than the vendor agents.

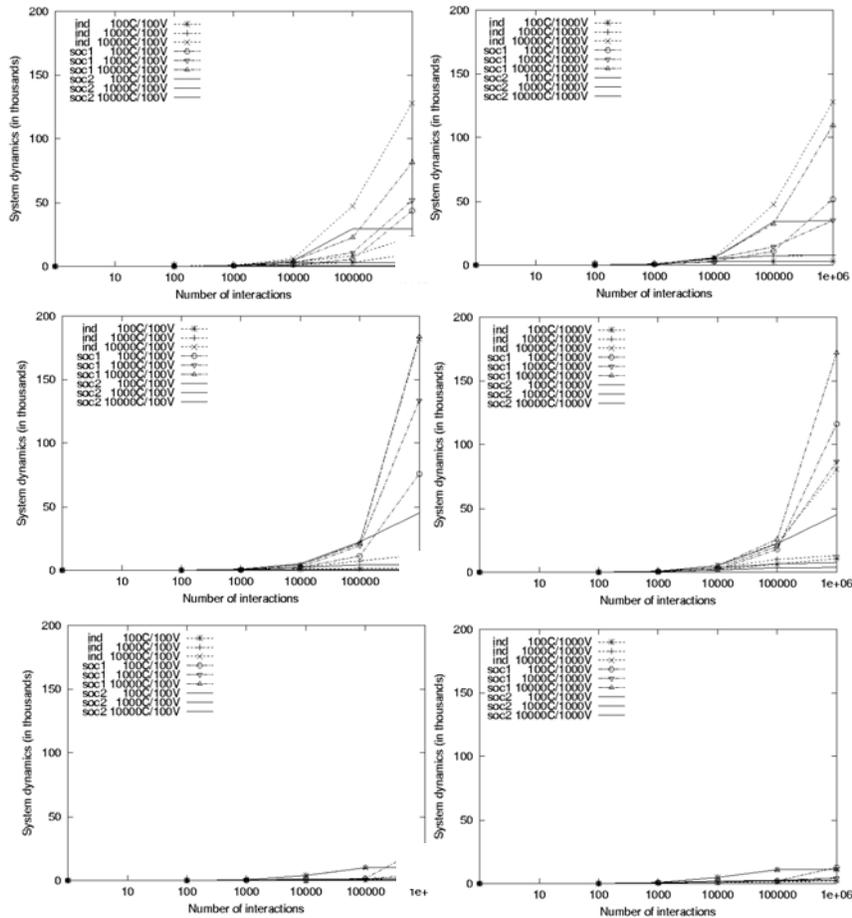


Fig. 3. System dynamics for: a) 100 vendors *simple*; b) 1000 vendors *simple*; c) 100 vendors *prob*; d) 1000 vendors *prob*; e) 100 vendors *costs*; f) 1000 vendors *costs*.

To conclude the analysis of the system dynamics we observe that the increase in the system dynamics is related to the merge of coalitions: they both start between 1000 and 10000 interactions and last the same period of time. The evolutions of the number of coalitions and of the systems dynamics are similar for all cases. Small differences occur as a consequence of delays that appear when taking into account diverse realistic conditions, such as increased probabilities or costs. Overall, our

system analysis shows that in all conditions the system is predictable (in number of existing coalitions) and that it reaches a stable state (in the overall dynamics) after a certain number of interactions that depends on the characteristics of the environment.

The individual gain - evaluated for the agent level of our analysis – calculates the average of the sum of all discounts that each customer receives. We plot it in Fig. 4 as a function of the number of interactions shown on a logarithmic scale. It ranges from 0 to 15 thousands. The general trend of the curves is to grow very slowly from 1 to 10,000 interactions, after which they raise linearly until 100,000 interactions and exponentially until 1 million interactions. The configurations with 100 customers have higher values than those with 1000 or 10000 customers due to the fact that after a number of interactions the overall discounts are similar, but divided among less customer agents. We observe that the values of system dynamics are higher for the *prob* setup, medium for the *simple* setup, and lower for the *costs* setup.

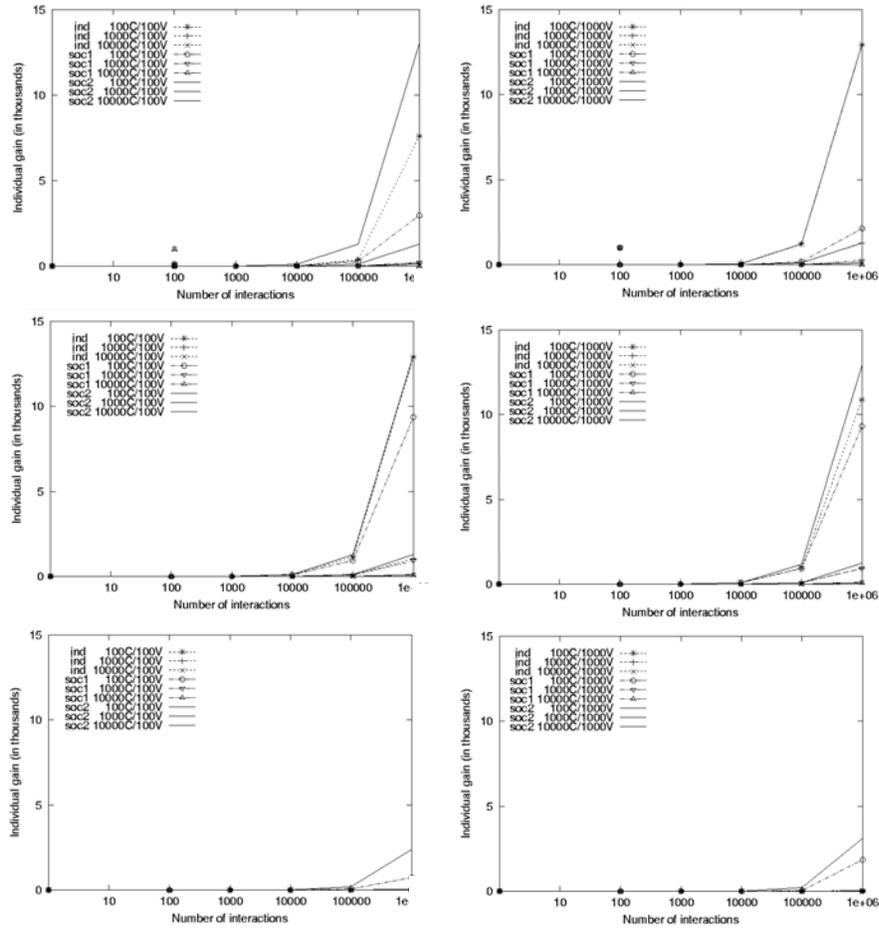


Fig. 4. Individual gain for: a) 100 vendors *simple*; b) 1000 vendors *simple*; c) 100 vendors *prob*; d) 1000 vendors *prob*; e) 100 vendors *costs*; f) 1000 vendors *costs*.

There is a small increase in the individual gain of customers in the *prob* setup compared to the *simple* setup. This is due to the higher probability of agents to interact with members of the same coalition that increases the customer's chances to get discounts. In the *costs* setup the individual gain has reduced values compared to the *simple* setup. This can be explained by the fact that the costs that agents pay when leaving a coalition are subtracted from their total benefits.

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 when the number of coalitions stabilizes to 1 the system dynamics stabilizes and the individual gain increases exponentially. This is observable in most cases that use the *soc2* strategy and some with *ind* in the *simple* and the *prob* setups. It is a direct consequence of agents finding compatible partners faster and stabilizing in certain coalitions.
- In the remaining cases from the *simple* and the *prob* setups (all with *soc1* and several with *ind*) the number of coalitions is still dropping (slowly) to its lower limit and the dynamics is increasing (linearly or exponentially), while the individual gain is increasing slower. This means that it is harder for the agents to find compatible partners and most profitable coalitions. We expect that in these cases the system will also reach equilibrium in the number of coalitions and in the overall dynamics shortly after 1 million interactions.
- When costs are involved (for all agent strategies) the number of coalitions drops very slowly from its peak values and the overall dynamics seems to stabilize after insignificant increases, but the individual gain is low. This can be explained by the fact that agents are more reluctant to leave their coalitions and join different ones. It drastically delays both the drop of the number of coalitions and the dynamics (that hardly increases). We expect that the number of coalitions will drop and the dynamics will increase and then stabilize, but much slower.

Overall, our results show that the proposed coalition formation mechanism is beneficial for the customer agents and for the system. It ensures exponential 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 equilibrated state for the *soc2* strategy. Although the system dynamics increases exponentially 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. Costs for leaving a coalition reduce the dynamics and the gain drastically, but also delay the drop in the number of coalitions. Increasing the probability that agents have interactions inside their own coalitions lead to an increase in the dynamics as well as in the gain.

5. Conclusion

In this paper, we proposed and evaluated a coalition formation mechanism that takes into consideration trust relationships between agents. 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 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 considering setups with agents with different coalition strategies, 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

1. 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, pp. 130-144.
2. O. Shehory, K. Sycara *Multi-agent Coordination through Coalition Formation* Proceedings of International Workshop on Agent Theories, Architectures, and Languages (ATAL97), Providence, RI, pp 135-146, 1997.
3. K. Lermann, O. Shehory *Coalition Formation for Large Scale Electronic Markets* Proceedings of the Fourth International Conference on Multiagent Systems ICMAS'2000, Boston, July 2000, pp. 216-222.
4. M Tsvetov, K. Sycara *Customer Coalitions in the Electronic Marketplace* Proceedings of the Fourth International Conference on Autonomous Agents, Barcelona, 2000, 263-264.
5. D. P. Reed *That Sneaky Exponential—Beyond Metcalfe's Law to the Power of Community Building* Context magazine Spring 1999.
6. Customer Relationships Management, available on line at <http://www.crm.com/>
7. J. Vassileva *Goal-Based Autonomous Social Agents Supporting Adaptation and Teaching in a Distributed Environment* Proceedings of ITS'98, San Antonio, Texas. LNCS No1452, Springer Verlag: Berlin pp.564-573, 1998.
8. 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, pp. 81-94.
9. 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, pp. 65-76.
10. C. Mudgal, J. Vassileva *Bilateral Negotiation with Incomplete and Uncertain Information: A Decision-Theoretic Approach Using a Model of the Opponent* Proceedings of the 4th International Workshop on Cooperative Information Agents (CIA IV), Boston, July 2000, pp. 107-118.