Social aspects of ReGreT, a reputation model based on social relations

Jordi Sabater, Carles Sierra

IIIA - Institut d'Investigacio en Intel.ligencia Artificial
CSIC - Spanish Scientific Research Council
Bellaterra, Catalonia
{jsabater,sierra}@iiia.csic.es

Abstract

The use of previous direct interactions is probably the best way to calculate a reputation but, unfortunately this information is not always available. This is especially true in large MAS where interaction is scarce. In this paper we present an extension to the ReGreT [11] reputation model that takes advantage of social relations between agents to overcome this problem. We also introduce how to use a reputation system to improve a negotiation model.

Keywords: Multi-agent systems, reputation, negotiation.

[This paper is a short version of [12]]

1 Introduction

The study and modeling of reputation has attracted the interest of scientists from different fields such as: sociology [14], economics [8], psychology [7] and computer science [4, 16]. According to the Oxford English Dictionary, *reputation* is "the common or general estimate of a person with respect to character or other qualities". This estimate is necessarily formed and updated along time with the help of different sources of information. Up to now, the reputation systems have been considering two different sources: (i) the direct interactions and (ii) the information provided by other members of the society about experiences they had in the past [11, 13, 15, 16]. Those systems, however, forget a third source of information that can be very useful. As a direct consequence of the interactions it is possible, even in simple societies, to identify different types of social relations between their members. Sociologists and psychologists have been studying these social networks in human societies for a long time. These studies show that it is possible to say a lot about the behaviour of individuals using the information obtained from the analysis of their social network. In this paper we evolve the **ReGreT** [11] (Reputation for Gregarious Traders) system to incorporate social networks in the reputation model.

In Section 2 we introduce the notion of social network analysis and its application to agent communities. Section 3 introduces the **ReGreT** system and section 4 describe the social aspects of the system in detail. In section 5 we explain how **ReGreT** can be integrated with a negotiation model. Finally sections 6 and 7 present the related work, discussion and future work.

2 Social Network Analysis and agent societies

Social network analysis is the study of social relationships between individuals in a society [14]. Social network analysis emerged as a set of methods for the analysis of social structures, methods that specifically allow an investigation of the relational aspects of these structures. The use of these methods, therefore, depends on the availability of relational rather than attribute data [14].

Relational data can be handled and managed in matrix form or using graphs. A graph structure that shows social relations is called a *sociogram*. A different sociogram is usually built for each social relation under study. Depending on the type of relation we have a directed or non-directed sociogram, with weighted edges or without.

We know that social network analysis can be used to Analise human societies. Could it also be suitable for agent societies? A drawback in social network analysis is that, sometimes, the translation of human social structures into their sociological meaning is difficult due to the complexity of the individuals and their relations. The greater simplicity, in social terms, of multi-agent systems suggests that social network analysis could be applied to autonomous agents with even better results.

We have exposed the pro but, of course, there is a con. Obviously, the more relational data the better the network analysis is. However, these data can be difficult to obtain. Sociologists usually obtain them through public-opinion polls and interviews with the individuals. This procedure is, of course, not possible in agent societies. Moreover, the analysis of human social structures is usually done by a sociologist external to the society. This external position gives the analyst a privileged watchtower to make this kind of analysis. In our case, as we want to use social analysis as part of the reputation system to be included in each agent, each agent has to do this analysis from its own perspective. It is beyond the scope of this paper to propose solutions about the way an agent builds such sociograms. You can see some work on this aspect in [10].

From now on, we will assume that the agent owns a set of sociograms that show the social relations in its environment. This sociograms are not necessarily complete or accurate. We suppose they are built by each agent using the knowledge it has about the environment. Therefore, sociograms are dynamic and agent dependent.

3 The ReGreT system

The **ReGreT** system structure is based on what we call the three *dimensions* of reputation: the *individual dimension*, the *social dimension* and the *ontological dimension*.

If an individual is considering only its direct interaction with the other members of the society to establish reputations we say that the agent is using the *individual dimension*. If the individual also uses the information coming from other members of the society and the social relations, we are talking about the *social dimension*. Finally, we consider that the reputation of an individual is not a single and abstract concept but rather a multi-facet concept. For example, the reputation of being a good carrier summarizes the reputation of having good planes, the reputation of never losing luggage and the reputation of serving good food. The different types of reputation and how they are combined to obtain new types is the base of the third *dimension* of reputation, the *ontological dimension*.

Although reputations also have a temporal aspect (the reputation value of an agent varies along time), we will omit the reference to time in the notation in order to make it more readable. We will refer to the agent that is calculating a reputation as a (what we call the "source agent") and the agent that is the object of this calculation as b (what we call the "target agent").

The *individual dimension* and the *ontological dimension* were explored in our previous work [11]. In this paper we will focus our attention in the *social dimension*.

4 The social dimension

Although direct interaction is the most reliable source of information, unfortunately it is not always available. Not only because the agent can be a newcomer to a society but also because the society can be very large. Therefore, when the interactions with another agent are scarce it is not possible to assign it a reputation based just on direct experiences. It is in these situations when the *social dimension* of an agent may help by using information coming from other agents. In the **ReGreT** system we use three types of social reputation depending on the information source:

- *Witness Reputation.* Based on the information about the target agent coming from other agents. We note this reputation as: $R_{a \xrightarrow{W} b}(\varphi)$
- *Neighbourhood Reputation.* Uses the social environment of the target agent, that is, the neighbours of the target agent and their relations with it. Noted as: $R_{a \xrightarrow{N} b}(\varphi)$
- *System Reputation.* It is a default reputation value based on the role played by the target agent. Noted as: $R_{a \xrightarrow{S} b}(\varphi)$

We use φ to note the aspect of the agent behaviour considered for that reputation. Each of these reputations requires a different degree of knowledge of the agent society and the target agent. The *System Reputation* is the easiest to calculate. We are assuming that the role an agent is playing is always "visible" information that is available to all the agents in the society. However, the role alone does not convey enough information to compute a reputation on all imaginable aspects. Also, the reliability of this type of reputation tends to be low because it doesn't take into account the peculiarities of the individual and its environment. This is the kind of reputation that an agent can use when it is a newcomer and there is an important lack of interaction with the other agents in the society. The *Witness Reputation* and the *Neighbourhood Reputation*, on the other hand, demand from the agent a moderate to hard knowledge of the social relations in the agent community. We explain below how to calculate all these reputation values.

4.1 Witness reputation

Beliefs about the reputation of others can be shared and communicated by the members of a society. The reputation that an agent builds on another agent based on the beliefs gathered from society members (witnesses) is what we call *witness reputation*. In a world with only homogeneous and trusted agents, this information is as relevant as the direct information. However, in the kind of scenarios we are considering, it may happen that the information be false or that witnesses hide it deliberately.

Besides that, the information that comes from other agents can be correlated (what is called the *correlated evidence problem* [9]). This happens when the opinions of different witnesses are based on the same event(s) or when there is a considerable amount of shared information that tends to unify the witnesses' way of "thinking". In both cases, the trust on the information shouldn't be as high as the number of similar opinions may suggest. Because the event(s) that have generated the opinions for each agent may be hidden, the agent cannot identify directly which agents are correlated. Schillo et. al [13] propose a method based on the analysis of "lying" as a stochastic process to implicitly reconstruct witness observations in order to alleviate this problem. We take a different approach based on the social relations between agents. Analyzing these relations, an agent can obtain useful information to minimize the effects of the correlated evidence problem.

4.1.1 Identifying the witnesses

The first step to calculate a witness reputation is to identify the set of witnesses (**W**). The initial set of potential witnesses might be the set of all agents that have interacted with the target agent in the past. This set, however, can be very big and its members probably suffer from the correlated evidence problem.

We take the stance that grouping agents with frequent interactions among them and considering each one of these groups as a single source of reputation values minimizes the correlated evidence problem. Moreover, assuming that asking for information has a cost, it has no sense to ask the same thing to agents that we expect will give us more or less the same information. Grouping agents and asking for information to the most representative agent within each group reduces the number of queries to be done. A domain dependent sociogram is what **ReGreT** uses to build these groups and to decide who is their most representative agent.

There are many heuristics that can be used to find groups and to select the best agent to ask. In the **ReGreT** system we use a heuristic based on the work by Hage and Harary [6]. Taking the subset of the selected sociogram over the agents that had had interactions with the target agent as the initial graph, the heuristic that **ReGreT** follows is:

1. To identify the *components* of the graph. A *component* is defined as a maximally connected subgraph.

2. To find the set of *cut-points* (CP) for each component. A *cut-point* is a node whose removal would increase the number of components by dividing the sub-graph into two or more separate sub-graphs among which there are no connections. A cut-point can be seen from a sociological point of view as indicating some kind of *local centrality*. Cut-points are pivotal points of articulation between the agents that make up a component [14].

3. For each component that does not have cut-points, to choose as a representative for that component the node with the larger degree. If there is more than one node with the maximum degree, choose the one that represents the agent more trustworthy. The degree can be regarded also as a measure of *local centrality* [14]. We refer to this set of nodes as LCP.

4. The set of selected nodes is the union between the set of *cut-points* and the set of LCP. That is, $\mathbf{W} = CP \cup LCP$.

Figure 1 shows an example of the application of the heuristic. At this point, the agent has to ask

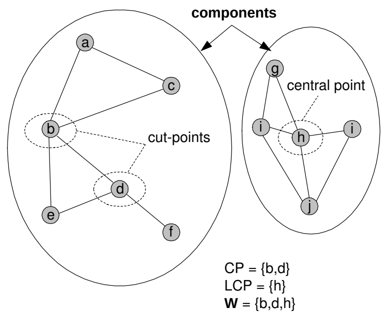

Figure 1: Witness selection within **ReGreT**.

for information to all the agents in the so calculated set of witnesses $\mathbf{W}$.

Once the information is gathered we obtain

$$\{\langle R_{w_i \to b}(\varphi), RL_{w_i \to b}(\varphi)\rangle \mid w_i \in \mathbf{W}' \subseteq \mathbf{W}\}$$

where $\mathbf{W}'$ is the subset of agents that answered a's query and $\langle R_{w_i \to b}(\varphi), RL_{w_i \to b}(\varphi)\rangle$ is the information received from the witness, that is, the reputation that witness w_i assigns to agent b together with its reliability. The reliability value is a measure of the reliance that the agent has on the reputation value (see [12] for more details).

The next step is to aggregate these values to obtain a single value for the *Witness Reputation*. However, as we said before, it is possible that this information be false so the agent has to be careful to give the right degree of importance and reliability to each piece of information. The degree of importance relies on the trust that each witness has. The system uses two different methods to calculate this trust: *social trust* and *outcome trust reputation*.

We define $socialTrust(a, w_i, b)$ as the trust degree that agent a assigns to w_i when w_i is giving information about b and taking into account the social relations among a, w_i and b.

Social ReGreT uses fuzzy rules [17] to determine how a social structure provides a reliability degree on the information coming from a given agent in that structure. The antecedent of each rule is the type and degree of a social relation and the consequent is the reliability of the information from the point of view of that social relation. For example, in a scenario where there is a type of relation based on cooperation, a possible rule could be:

IF $\text{coop}(w_i, b)$ is high
THEN $\text{socialTrust}(a, w_i, b)$ is very_bad

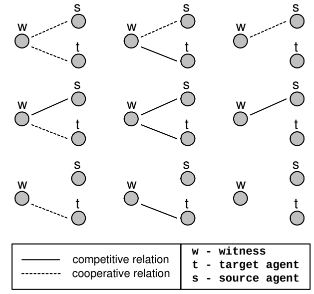

Figure 2: Relevant social structures in the example.

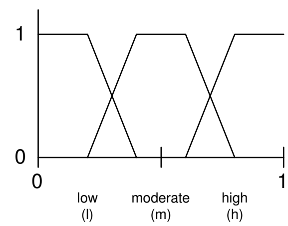

Figure 3: Intensity of a social relation.

that is, if the level of cooperation between w_i and b is high then the trust from the point of view of a on the information coming from w_i related to b is very bad. The heuristic behind this rule is that a cooperative relation implies some degree of complicity between the agents that share this relation so the information coming from one about the other is probably biased.

Which relations are relevant to calculate the reliability of an agent depend on the meaning that each relation has in the specific agent society. For instance, figure 2 shows the social structures that should be taken into account for an agent in a society where the cooperation and the competition are the relevant relations to measure the reliability of another agent.

Table 1 shows a possible set of fuzzy rules for this case. Figure 3 contains the fuzzy set values -defined over the intensity label on the arcs in the sociogram- for the variables $\text{coop}(w_i, a)$, $\text{coop}(w_i, b)$, $\text{comp}(w_i, a)$, and $\text{comp}(w_i, b)$, and figure 4 the fuzzy set values for the variable $\text{socialTrust}(a, w_i, b)$. Although currently the kind of influence of each social structure is static and based in human common sense, we plan to improve the system with a rule learning mechanism.

A second way to calculate the degree of trust of an agent is using the *outcome trust reputation* of the

IF	coop(w, a) is l	THEN	socialTrust(a, w_i, b) is g
IF	coop(w, a) is m	THEN	socialTrust(a, w_i, b) is vg
IF	coop(w, a) is h	THEN	socialTrust(a, w_i, b) is vg
IF	comp(w, a) is l	THEN	socialTrust(a, w_i, b) is sb
IF	comp(w, a) is m	THEN	socialTrust(a, w_i, b) is b
IF	comp(w, a) is h	THEN	socialTrust(a, w_i, b) is vb
IF	coop(w, b) is l	THEN	socialTrust(a, w_i, b) is sb
IF	coop(w, b) is m	THEN	socialTrust(a, w_i, b) is b
IF	coop(w, b) is h	THEN	socialTrust(a, w_i, b) is vb
IF	comp(w, b) is l	THEN	socialTrust(a, w_i, b) is sb
IF	comp(w, b) is m	THEN	socialTrust(a, w_i, b) is b
IF	comp(w, b) is h	THEN	socialTrust(a, w_i, b) is vb
IF	no_rel(w, b)	AND THEN	no_rel(w, a) socialTrust(a, w_i, b) is g

Table 1: Witness reputation fuzzy rules.

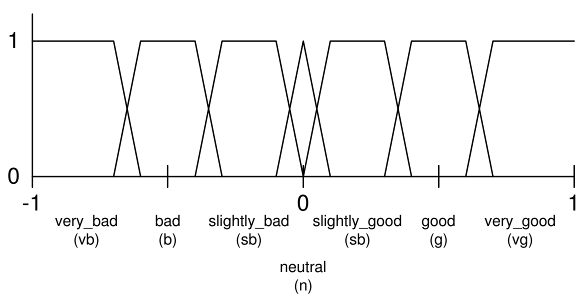

Figure 4: Fuzzy sets for reputation.

'trust' that other agent deserves. The *outcome trust reputation* is calculated using the *individual dimension* and takes into account the direct experiences with the other agent (outcomes). In this case the outcomes used to calculate this type of reputation are build from the information that the other agent sent previously and how accurate was that information. We note the *outcome trust reputation* as $R_{a \overset{O}{\to} w_i}(trust)$ and its reliability as $RL_{a \overset{O}{\to} w_i}(trust)$. For more details about how to calculate an outcome reputation see [12].

The trust values calculated using an outcome trust reputation are more useful than those based on social relations (socialTrust) because an outcome trust reputation is based on the individual experiences and thus takes into account its particularities while the analysis of social structures rely on global expected behaviours. However, in those situations where there is not enough information to calculate a reliable outcome trust reputation, the analysis of social relations can be a good solution. Usually, social relations are easier to obtain than a set of outcomes (necessary to calculate an outcome trust reputation) especially if the agent has just arrived to a new scenario. Given that, we define a's trust degree for an agent w_i when it is giving information about b as:

$$trust(a, w_i, b) = RL_{a \overset{O}{\to} w_i}(trust) \cdot R_{a \overset{O}{\to} w_i}(trust) + (1 - RL_{a \overset{O}{\to} w_i}(trust)) \cdot socialTrust(a, w_i, b)$$

That is, the agent uses the trust reputation based on direct interactions if it is reliable, if not, it uses the social trust.

Now we have all the necessary tools to calculate the *witness reputation* and its reliability considering that the information coming from the witnesses can be false. The formulae we propose to calculate these values are:

$$R_{a \overset{W}{\to} b}(\varphi) = \sum_{w_i \in \mathbf{W}'} \omega^{w_i b} \cdot R_{w_i \to b}(\varphi)$$

$$RL_{a \overset{W}{\to} b}(\varphi) = \sum_{w_i \in \mathbf{W}'} \omega^{w_i b} \cdot \min(trust(a, w_i, b), RL_{w_i \to b}(\varphi))$$

where $\omega^{w_i b} = \frac{trust(a, w_i, b)}{\sum_{w_j \in \mathbf{W}'} trust(a, w_j, b)}$

These formulae require some explanations. To calculate the *witness reputation* the agent uses the normalized trust of the witness to weight each opinion in the final value. For the reliability, we want that in the final reliability value, the contribution of each individual reliability be in the same proportion that its related reputation. Therefore, the agent uses the same weights in the reliability formula that in the reputation formula. Finally to calculate the reliability of an individual reputation, the agent uses the minimum between the trust of the witness that sent the reputation and the reliability that the witness himself gives to that reputation. We use this method to model the idea that if the witness is a trusty agent, then we can use his own measure of reliability for the reputation, if not, we cannot rely on his information and we have to calculate our own measure of trustworthiness for that reputation based on the outcomes and the social relations of that witness.

4.2 Neighbourhood reputation

The reputations of the individuals that are in the neighbourhood of the target agent and their relation with him are the elements used to calculate what we call the *Neighbourhood Reputation*. Neighbourhood in a MAS is not related with the physical location of the agents but with the links created through interaction. The main idea is that the behaviour of these neighbours and the kind of relation

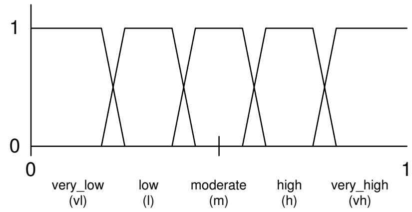

Figure 5: Fuzzy sets for reliability.

they have with the target agent can give some clues about its possible behaviour. The set of neighbours of an agent b is noted as $\mathbf{N}_b = \{n_1, n_2, \cdots, n_n\}$.

To calculate a *Neighbourhood Reputation* we use fuzzy rules as well. These rules, that are domain dependent, relate the outcome reputation of a target's neighbour and the social relation they have, with a reputation of the target agent.

The application of these rules generates a set of *individual neighbourhood reputations* noted as $R_{a \xrightarrow{n_i} b}(\varphi)$. For instance:

IF $R_{a \to n_i}$(swindler) is X AND coop$(b, n_i) \geqslant$ low
THEN $R_{a \xrightarrow{n_i} b}(swindler)$ is X
IF $RL_{a \to n_i}$(swindler) is X' AND coop(b, n_i) is Y'
THEN $RL_{a \xrightarrow{n_i} b}(swindler)$ is T(X', Y')

With these rules we are saying that if the neighbour of the target agent is a swindler and there is a relation of cooperation between the target and this neighbour, then the target is also (assumed to be) a swindler. The fuzzy sets for the reliability were shown in figure 5.

Finally table 2 shows a possible set of values for function T.

X' Y'	**l**	**m**	**h**
vl	vl	vl	vl
l	vl	vl	l
m	vl	l	m
h	l	m	h
vh	m	h	vh

Table 2: Function T used in reliability rules.

The general formulae we use to calculate a *neighbourhood reputation* and its reliability are:

$$R_{a \xrightarrow{N} b}(\varphi) = \sum_{n_i \in N_b} \omega^{n_i b} \cdot R_{a \xrightarrow{n_i} b}(\varphi)$$

$$RL_{a \xrightarrow{N} b}(\varphi) = \sum_{n_i \in N_b} \omega^{n_i b} \cdot RL_{a \xrightarrow{n_i} b}(\varphi)$$

where $\omega^{n_i b} = \frac{RL_{a \xrightarrow{n_i} b}(\varphi)}{\sum_{n_j \in N_b} RL_{a \xrightarrow{n_j} b}(\varphi)}$

In this case we are using the reliability of each *individual neighbourhood reputation* to weight the contribution to the final result, both for the reputation and for the reliability.

4.3 System reputation

The idea behind *System reputations* is to use the common knowledge about *institutional structures* and the role that the agent is playing for that *institutional structure* as a mechanism to assign default reputations to the agents. An *institutional structure* is a social structure the members of which have one or several *observable* features that unambiguously identify them as members of that social structure. The fact that there are *observable* features to identify its members is what differentiates an institutional structure from other social structures. Examples of institutional structures in human societies are the police, a company or a family. We assume that the role that an agent is playing and the institutional structure it belongs to is something "visible" and trustworthy for each agent.

Each time an agent performs an action we consider that it is playing a single role within the institutional structure. An agent can play the role of buyer and seller but when it is selling a product only the role of seller is relevant. Although we can think up some situations where an agent can play two or more different roles at a time, we consider that there is always a predominant role so the others can be disregarded.

In **ReGreT** the reputations associated to each role within an institutional structure are domain dependent and part of the initial knowledge of the agent. The value for these reputations can be different depending on which institutional structure the agent belongs to. This models the idea that groups (in our case institutional structures) influence the point of view of their members [7]. An institutional structure does not always has to associate a reputation value to each contract issue.

System reputations are calculated using a table for each institutional structure where the rows are the possible roles, and the columns the value for the reputation.

System reputations are noted as $R_{a \xrightarrow{S} b}(\varphi)$ and its reliability as $RL_{a \xrightarrow{S} b}(\varphi)$.

5 Integrating ReGreT with a negotiation model

In this section we show how **ReGreT** can be used together with the negotiation model presented in [5] to improve the success of the agent in its negotiations. The purpose of the negotiation is to reach an agreement about the provision of a service. This service is reflected in the contract, and the issues of this contract are what the agents negotiate. It is beyond the scope of this article to explain how the negotiation model works. If you want a complete description refer to [5]. We will focus our attention on the connection point between the negotiation model and **ReGreT**.

In the negotiation model, when an agent a receives an offer from agent b at time t, it evaluates the offer using an utility function. If the utility value for that offer is greater than the utility value that the counter offer that agent a is ready to send at time $t+1$ has then, agent a accepts. Otherwise, the counter offer is submitted.

The problem arises when the fulfillment of the contract is not as expected. If the fulfillment of the contract is worse than expected, probably the agent will not cover its needs. If the fulfillment is better than expected then it means that the agent may have lost time in negotiation. For instance, the agent could have accepted a not so good offer knowing that at the end it would be good enough. Also, arriving to an agreement sooner minimizes the risk of partners withdrawal.

The general formula used in the negotiation model to calculate the utility that has a contract for an agent a is:

$$V^a(X^c) = \sum_{i \in I} \omega_i^a V_i^a(x_i^c)$$

where X^c is the contract, I the issues under negotiation, $x_i^c \in [min_i, max_i]$ a value for issue i, ω_i^a the importance of issue i for agent a and $V_i^a : [min_i, max_i] \rightarrow [0,1]$ the scoring function for issue i. We propose to use reputation to "modulate" the shape of the score function. The idea is that depending on the reputation, the scoring function will have a more preservative or a more confident behaviour. Therefore, a very good reputation meaning that the partner is splendid in the fulfillment of contracts implies that the agent can accept a deal that does not satisfy completely its requirements knowing that, in the end, it will be OK. On the other hand, a bad reputation implies that the agent has to be harder in its negotiation.

To achieve this, we introduce the notion of an utility function that takes into account reputation. A possible utility function that considers reputation could be:

$$VR^a(X) = \max(0, \min(1, V^a(X) + k \cdot RL_{a \to b}(\varphi) \cdot R_{a \to b}(\varphi)))$$

where φ is the reputation type selected by the agent.

A good reputation with a high reliability increases the value that the agent gives to the offer. The constant k models the relevance the agent gives to reputation in the negotiation process.

6 Related work

The idea of using the opinion of other agents to build a reputation is not new. The work of Michael Schillo, Petra Funk and Michael Rovatsos[13] and the work of Bin Yu and Munindar P. Singh [15] are good examples of this. In both cases they use a trust-net for weighting the other agents' opinions. Our structure to calculate the witness reputation can be considered also a trust-net. In our case, however, besides the previous experiences with the witnesses we also consider the information about the agents' social relations.

The model described in [15] merges information that comes from agents that have good reputation. In [13] the same agents that can provide you with information are also competing with you. Although agents are assumed to never lie, they can hide information or bias it to favour their goals. We go one step further and consider that the agents can also lie.

In electronic marketplaces, the reputation that a user has is the result of aggregating all the experiences of the other users that interacted with him/her in the past. Amazon Auctions [1], eBay [2] and OnSale Exchange [3], for instance, are online auction houses where users buy and sell goods. Each time a new transaction is finished, the buyer rates the seller. These ratings are used to build the reputation of a seller. Sporas [16] is an evolved version of this kind of reputation models. Sporas introduces the notion of reliability of the reputation and is more robust to changes in the behaviour of a user than reputation systems like Amazon Auctions, based on the average of all the ratings given to the user. In all these systems each user has a global reputation shared by all the observers instead of having a reputation biased by each observer. Histos [16], also oriented to electronic commerce, is a more personalized reputation system where reputation depends on who makes the query, and how that person rated other users in the on-line community.

A model that uses also social networks to calculate a measure of reputation is [10]. The main difference with **ReGreT** is that this model uses only the location of each agent in the community social network to calculate a measure of reputation. This work propose a method to calculate what, in terms of **ReGreT** terminology, is called the *neighbourhood reputation.*

7 Discussion and future work

In this paper we have presented how social network analysis can be used in a reputation system that takes into account the social dimension of reputation. The use of the social network analysis techniques as part of a reputation system opens a new field for experimentation. Our first objective is to validate the system in a realistic e-commerce environment where social relations are an important factor. To be able to exploit all the capabilities of the **ReGreT** system we need environments more sophisticated than the actual e-markets like Amazon Auctions or eBay. We are working in several tools that allow the specification and implementation of these kind of e-markets.

Once you introduce the social dimension in reputation systems and the agents start to take into account social relations, it becomes more and more important to consider not only which is the reputation of the other agents, but what can an agent do to get and maintain a good reputation. Using the **ReGreT** system, we want to study reputations from this new perspective.

Acknowledgements

This work has been supported by the European project SLIE, IST-1999-10948, and the Spanish MCYT project e-INSTITUTOR, MCYT 2000-1414.

References

[1] *Amazon Auctions.* http://auctions.amazon.com.

[2] *eBay.* http://www.eBay.com.

[3] *OnSale.* http://www.onsale.com.

[4] C. Castelfranchi and R. Falcone. Principles of trust for mas: Cognitive anatomy, social importance, and quantification. In *Proceedings of the 3th International Conference on Multi-Agent Systems*, pages 72—79, 1998.

[5] P. Faratin, C. Sierra, and N.R. Jennings. Negotiation decision functions for autonomous agents. *Robotics and Autonomous Systems*, (24):159—182, 1997.

[6] P. Hage and F. Harary. *Structural Models in Anthropology.* Cambridge University Press, 1983.

[7] Marvin Karlins and Herber Abelson. *Persuasion, how opinion and attitudes are changed.* Crosby Lockwood & Son, 1970.

[8] R. Marimon, J.P. Nicolini, and P. Teles. Competition and reputation. In *Proceedings of the World Conference Econometric Society, Seattle*, 2000.

[9] J. Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference.* Morgan Kaufmann, 1988.

[10] J.M. Pujol and R. Sangesa. Reputation measures based on social networks metrics for multi agent systems. In *Proceedins of the 4t Congres Catala en IA, Barcelona, Catalunya*, pages 205—213, 2001.

[11] J. Sabater and C. Sierra. Regret: A reputation model for gregarious societies. In *Proceedings of the Fourth Workshop on Deception, Fraud and Trust in Agent Societies, Montreal, Canada*, pages 61—69, 2001.

[12] J. Sabater and C. Sierra. Social regret, a reputation model based on social relations. In *Proceedings of the First International Join Conference on Autonomous Agents and MultiAgent Systems, Bologna, Italy*, 2002.

[13] Michael Schillo, Petra Funk, and Michael Rovatsos. Using trust for detecting deceitful agents in artificial societites. In *Applied Artificial Intelligence, Special Issue on Trust, Deception and Fraud in Agent Societies*, 2000.

[14] J. Scott. *Social Network Analysis.* SAGE Publications, 2000.

[15] Bin Yu and Munindar P. Singh. A social mechanism of reputation management in electronic communities. In *Cooperative Information Agents, CIA-2000, Boston, MA, USA*, pages 154—165, 2000.

[16] Giorgios Zacharia. Collaborative reputation mechanisms for online communities. Master's thesis, Massachusetts Institute of Technology, September 1999.

[17] L. A. Zadeh. Fuzzy sets. *Information and Control*, (8):338—353, 1965.