Comparing predictions of SPORAS vs. a Fuzzy Reputation System

JAVIER CARBO, JOSE M. MOLINA,
Depto de Informática
Universidad Carlos III de Madrid
Av. Universidad 30, Leganes 28911 Madrid

JORGE DAVILA
Facultad de Informática
Universidad Politécnica de Madrid
MonteGancedo, Boadilla del Monte 28660 Madrid
SPAIN

Abstract: - Autonomous Agents make frequent interactions between people unrelated to each other possible. This is one of the most remarkable and attractive features of Electronic Commerce. Reputation management is therefore an important issue in order to make costly decisions in multi-agent systems. Different models have been proposed to ensure the convergence of the predictions and to avoid the abuse of previously acquired reputation. In this paper we define a fuzzy definition of reputation, and we compare the performance of our approach versus one of the most well known (crisp) reputation mechanism: SPORAS. We finally provide an analysis of the potential benefits of using a fuzzy reputation management according to the experimental results obtained.

Key-Words: - Reputation, Trust, Fuzzy Logic, Recommendations, Multi-Agent Systems, Electronic Commerce

1 Introduction

Internet brings together a broadly extended electronic community offering a wide range of services and products to its members. Most of these services are intended to be paid electronically so commercial interactions play a central role in such electronic context.

All electronic purchases involve six different steps from the customers' point of view [1]: identification need, product profile, merchant selection, automatic negotiation, payment and evaluation. Among them, we are interested in the merchant selection. This stage needs information about the merchant to decide who to trust in, and therefore, who to buy from.

In open systems, agents may not make such decision with complete information due to the variable and large number of agents involved [2]. So cooperation with other agents may contribute to avoid frauds and to make better decisions [3].

Reliable information is acquired and updated along time through the ratings of direct interactions. But the soundness of the recommendations provided by others based on their past experiences is more difficult to guess [4]. And then, decisions made with these recommendations assume some level of uncertainty.

Furthermore, the evaluation of a merchant can not be always built up from objective and commonly accepted criteria [5]. When humans evaluate purchases, they record impressions that reflect how well a merchant suited their expectations and these impressions will often be pretty vague.

Reputation is then a value computed from both sources of information: direct interactions and recommendations. An agent will trust in other one if the reputation of that agent achieves certain decision threshold [6].

In the following section we will outline a brief overview of the solutions proposed to reputation issues. We will explain in detail the algorithm involved in SPORAS since it will be used to test the results of our model of reputation. Next we will present the motivation and considerations hold to define a fuzzy reputation before the explanation of the computations applied over such fuzzy sets. Section 4 will describe two different simulations, where both solutions will show their performance. Finally we will analize the results of the experimentation in section 5.

2 Related work

In Computer Science, reputation has been studied from very different point of views. In some scenarios (closed electronic marketplaces) reputation can be considered a public property. In such context, a central entity computes and updates a global reputation of each agent of the system. Ebay [7], OnSale [8], and Yenta [9], for instance, use a centralized approach.

But in open and dynamic environments (for instance multi-agent systems), the distributed nature of agents suggests a decentralized management of trust (and reputation) to avoid posible bottlenecks [10].

The idea of using recommendations from other agents jointly with direct experiences to build reputation of merchants is predominant, but there are also proposals of decentralized trust based only on direct interactions [11].

Several formalisms represent reputation according to the different point of view of researchers. For instance, bayesian networks, ranges and linguistic categories are proposed in the corresponding [12], [13] and [14] contributions.

Among the different reputation mechanisms, in advance we will focus our attention in the iterative algorithm used in SPORAS [15] to compute reputation.

SPORAS is intended to provide a reputation service in loosely connected communities. It updates reputation recursively. The reputation of certain user at a time i, R_i, is computed from the previous one, R_{i-1} given a purchase rated as W_i. The formulae applied is:

$$R_i = R_{i-1} + \frac{1}{\theta} \bullet \Phi(R_{i-1}) \bullet (W_i - R_{i-1})$$

(1) SPORAS recursive reputation equation

Let θ be the effective number of ratings taken into account in the evaluation (θ>1). The bigger the number of considered ratings, the smaller the change in reputation is.

Furthermore, function Φ is defined in order to slow down the incremental changes for very reputable users. It is computed from:

$$\Phi(R_{i-1}) = 1 - \frac{1}{1 + e^{-(R_{i-1} - D)/\sigma}}$$

(2) SPORAS damping function

Where dominium D is the maximum possible reputation value and σ is chosen so that the resulting Φ would remain above 0.9 when reputations values were below ¾ of D.

3 AFRAS

AFRAS is a multiagent system devoted to manage reputation using fuzzy logic. AFRAS stands for A Fuzzy Reputation Agent System. A preliminary version of the system was presented in [16]. The adopted architecture of the agents in AFRAS is based on the BDI paradigm and it was described in [17]. Previous works also tackled some security and privacy issues involved in the overall design (like protecting recomendations [18] and fair initial matchmaking [19]).

We take the stance that reputation is a single (not multi-facet) value due to the generic and abstract view of trust concept. Our system is a distributed approach, where each agent has its own opinion about the rest of the agents in the system. Every one of them may anytime act as seller, buyer or recommendator. Reputation is then built as a result of all the actions hold, irrespective of the role that the agent is currently playing. With this approach we pursue the aim of reflecting the way humans trust in others through a ‘word of mouth’.

3.1 Motivation of a fuzzy approach to reputation

We can represent reputation with fuzzy sets due to the vague, uncertain and incomplete nature of the information and opinions used to define reputation. First, when this concept is applied to merchants, it involves how much the services provided suit the user’s expectations. These expectations are user-defined, and therefore they will often be pretty vague. Furthermore, due to the untraceability of every agent in open and dynamic systems, the information that any agent perceives is incomplete. Finally, as we face competitive scenarios, benevolence should not be assumed, and therefore such information will be considered as uncertain.

We represent a fuzzy value with the four squares that define a trapezium. When these values are associated to the reputation of a merchant, the fuzzy set may be interpreted in such a way that its regularity and its reliability are implicit in them. Therefore, our system allow us to distinguish, through the shape of the trapezium, between a merchant very reliable but irregular and other one more regular although a bit less reliable.

The design of the agents in the system is driven to model human-like atributes. In this way, the human atributes of susceptibility, sociability and shyness are represented through trust thresholds used to make buying/answering/asking decisions. If these limits were crisp, they would be difficult to concrete from human interaction, and users won't feel comfortable with such numerical values when his agent would explain its behaviour to him. As frequent changes of mind should be avoid, representing the vague and abstract concepts of being reputed 'enough to be asked/answered/bought' as fuzzy thresholds makes sense.

Together with these three human attitudes associated to the agent's behaviour, we count on another human-like atribute: remembrance. By such concept we mean the ability to impress the memory of the user and to lead him towards a more cautious behaviour. Any deceptions will afect all future purchases, not only the ones from the disappointing merchant. It also implies a certain level of reactivity to failure results. In one way, the smaller the remembrance or memory, more weight will be assigned to the satisfaction generated by the last merchant's behaviour. In the opposite way sensitivity will be reduced when no deceptions were found.

3.2 Computing a fuzzy reputation

When a purchase is carried out, the given rating will be used to update the reputation of the corresponding merchant. Supposed i previously performed purchases. Let R_{i-1} be the reputation of the merchant until that instant. And let S_i be the rating (satisfaction provided by the merchant) in the purchase which took place at i. The new reputation of the merchant is then dependent on these fuzzy values, S_i and R_{i-1}. We combine them using different weights. These weights are defined from the importance attached to the story of purchases over the last experience. We call remembrance of the agent to this concept.

Formally, the new reputation of a merchant after the ith purchase, can be computed from the following formulae:

$$R_i = \frac{R_{i-1} \bullet \Phi + S_i \bullet (2-\Phi)}{2}$$

(3) First approach of a fuzzy reputation equation

If we suppose that remmembrance Φ was equal to 0, the last rating would count as much as the previous reputation. If remembrance was Φ=1, then reputation would remain constant. In other intermediate cases, both factors would be weighted with values between 0 and 1.

Remembrance should also be updated after each purchase according to the success of the last prediction. So we compare the similarity between the reputation prevision and the purchase rating to update remembrance. Simililarity Δ is computed through the support of the matching between the fuzzy S_i and R_{i-1}. Therefore, the formulae applied to update remembrance is the next one:

$$\Phi' = \frac{\Phi + \Delta}{2}$$

(4) Updating remembrance

The outcome of this equation follows the next simple principles:

- When the prediction fitted well the rating, $\Delta \approx 1$ and therefore remembrance (the importance given to the story of purchases over the last one) will be increased in 1/2 - Φ/2.
- If $\Delta \approx 0$ then remembrance will be halved.

These properties avoid Φ to be below 0 and above 1.

The initial value of remembrance associated to any agent joining the system, should be 0, but it will increase when the number of relative success becomes relevant. The actual value of remembrance depends mainly on the variability of the all the merchants in the system and on the number of interactions that carried out by the agent until that instant.

The result of the applying equation 3 over R_{i-1} and S_i can be seen graphically as a movement of the four squares of the trapecium that defines the fuzzy reputation R_i towards the last rating of the merchant. However when the outcome from that merchant does not overlap the expected fuzzy set of its reputation, the reliability of the resulting reputation should be decremented. We apply the desired decrement modifying the shape of the trapezium obtained from the first approach of a fuzzy reputation formulae (equation 3).

We use again the similarity level between them to modify the square of the resulting reputation R_i that is nearest to the rating of the merchant S_i. The smaller the similarity between R_{i-1} and S_i, the less steep the side of the trapezium is.

Let R^0 be the square of the reputation nearest to the rating of the merchant, and S^0 the corresponding square of such rating. The computation of the new R^{0}' can be formalized as:

$$R'^0 = R^0 \bullet \Delta + S^0 \bullet (1-\Delta)$$

(5) Modifying the nearest square of the trapezium

We can see the effect of equation (5) in the next figure:

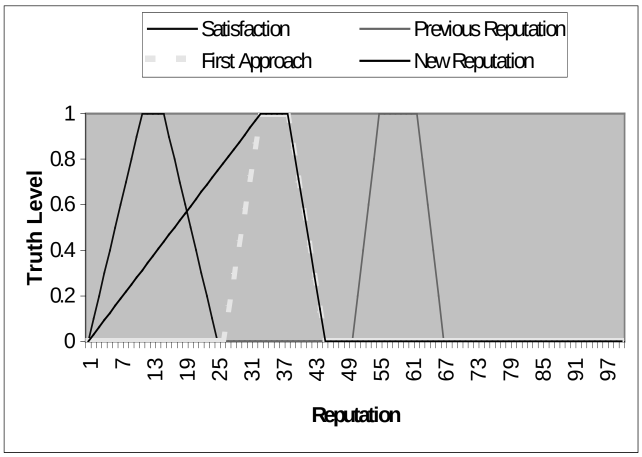

Fig. 1 Reliability Modification

4 Experimental Results

In order to compare Sporas with the fuzzy approach of AFRAS, we have applied both algorithms to the simulations proposed by Zacharia and Maes [15]. In them the convergence of the predictions and the abuse of prior performance were evaluated. No recommendations were involved in these simulations. They studied the level of success obtained from the point of view of an agent continously buying to all the merchants. The merchant chosen in each iteration to interact with, was selected randomly among all the merchants.

Our experiments have tested the results of one agent randomly buying to 10 merchants in 120 purchases. The same order of merchants selected along the 120 iterations and the same response from those merchants was applied to AFRAS and SPORAS in order to obtain a consistent comparison. Due to this reason, we have also adapted the range of reputation and ratings values of SPORAS to be from 0 to 100.

We have represented the reliability and the standard deviation of SPORAS as the width of the sides of the trapezium in AFRAS.

4.1 Convergence of the predictions

In order to evaluate the convergence of the reputation estimations, the simulation has satisfied the next properties:

- Each of them has assigned a prefixed behavior (uniformly distributed). And all of them use such prefixed behavior along the 120 purchases.
- The satisfaction provided at any time is drawn from a normal distribution. The mean of that distribution is equal to the prefixed behavior. The standard deviation is assumed to be 3.33 (over 100).
- Initially the reputation of all merchants is 10 (over 100) with a reliability of 1 (over 100).

The error produced in each purchase has been quantified in AFRAS through the difference between the corresponding means of maxima of the fuzzy sets R_{i-1} and S_i.

Next, we will show the average error committed in the predictions computed by both alternatives in order to compare the speed of convergence.

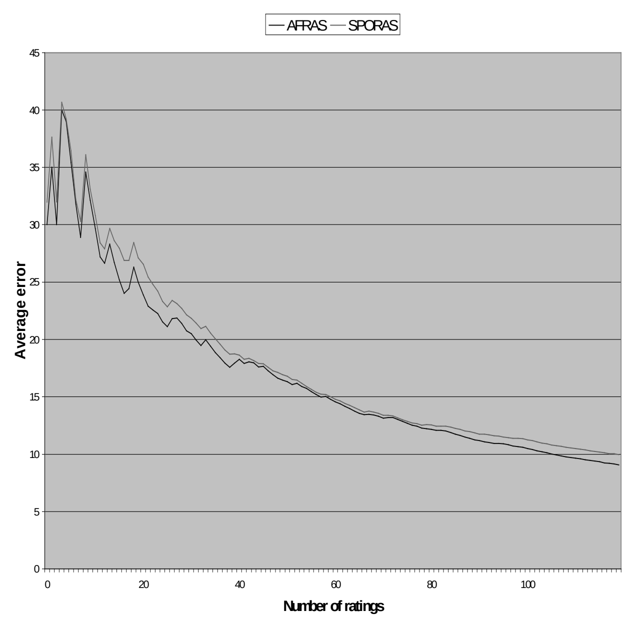

Fig. 2 Convergence of SPORAS & AFRAS

In figure 2 we can see how both algorithms, SPORAS and AFRAS, converge to the real behavior of each merchant since the average error committed along the 120 iterations decreases continously. The curves of figure 2 show how AFRAS fits slightly better than SPORAS but we can not label such difference as significant.

4.2 Abuse of prior performance

The second kind of experiments hold has evaluated the possible advantage that anyagent could obtain from a previous right behaviour. So this simulation consists of two stages. First, one merchant behaves reliably until the merchant would reach a high reputation. Second, the merchan begins abusing of its prior reputation to commit fraud.

We will compare the evolution of reputation in AFRAS and in SPORAS to observe which system avoids better the abuse of previously acquired reputation.

In the example to be shown in this paper, the fraudulent merchant sells with a rating of 80 over 100 along the first third of the interactions (40 over 120). After the 40th purchase with any merchant, the malitious merchant behaves as a bad seller (24 over 100) until the end of the simulation.

The rest of the merchants behave constantly in the same way as in the first simulation (uniformly distributed with a standard deviation of 3.33). They have also assigned the same initial reputation values (10 over 100 with a reliability of 1).

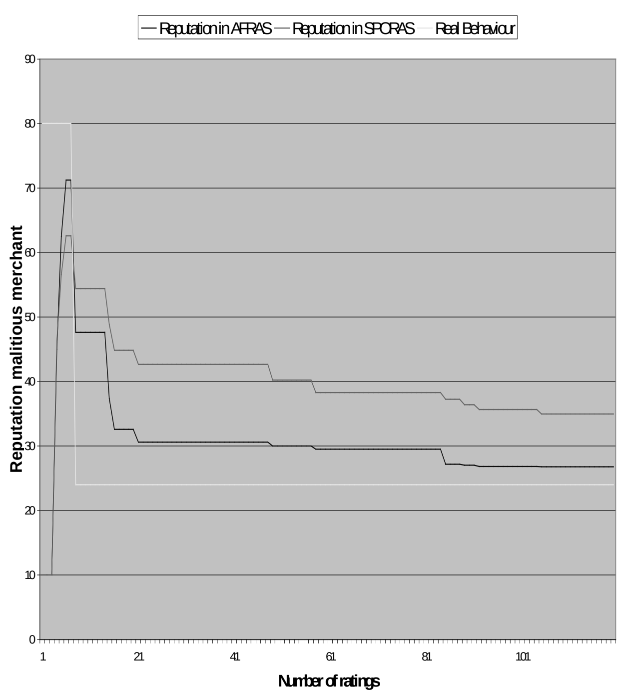

Fig. 3 Fraud detection in AFRAS & SPORAS

Figure 3 shows how fraud is detected first in AFRAS and less abuse can be committed with our fuzzy approach than in SPORAS. The improvement in the predictions of AFRAS vs. SPORAS is now relevant along the last two-thirds iterations of the simulation.

The fact of mixing purchases from different merchants adds a higher value to the results from AFRAS because this system updates reputation with a remembrance value dependent of the overall variability of merchants. The constant behaviour of the rest of the merchants is introducing a distorting factor in the predictions about the malitious one.

In spite of the small range of experiments showed in the graphics, other initial conditions (with different number of merchants, iteration, initial reputation ratings, response of merchants) do not change significantly the results of the comparison included in this paper.

5 Conclusions and Future Work

We can explain the good adaptation of AFRAS against SPORAS in the detection of the sudden fall of the merchant behavior studying the details of each formulae.

SPORAS introduces a damping factor used to soften the changes when agents have a high reputation. It also reduces the contribution of new ratings as the number of effective ratings grows up.

This second factor is tackled partially in REGRET [20] giving more (fixed) weight to last iterations over previous ones. We will also compare our fuzzy approach with REGRET in the very next future.

AFRAS gives to agents more sensitivity to last experiences when they had recent deceptions. This weight depends on the overall variability of the behavior of all the merchants. Previous works of AFRAS studied the adaptative response of the predictions according to such variability [21].

If we had used a specific remembrance value for each merchant, the behavior of the systems would be probably even better. But we tried to represent with this factor a general cautious attitude of humans when they face a deception from any source. By this way, we also reduce the computational space required to update reputation computations.

Other classic recommender systems as Ebay or Onsale could be used for further comparisons, but SPORAS computations seemed to be better founded and improved the adaptative behavior of the predictions [15]. We will focus on the proposal from Sing and Yu [13] soon to complete the benchmark of our fuzzy proposal.

Every one of these reputation systems is designed to make predictions using recommendations. The experiments showed in this paper do not include them. Without this cooperative attitude of the agents, these systems are just simple machine-learning algorithms. We will also observe in future works how the use of recommendation fastens the adaptation of the predictions and how it reduces the possible abuse of prior performance. When recommendations will be involved in such decisions, we will study how much distortion includes the collusion of a malitious agent sending false recommendations about a fraudulent merchant.

The role played by thresholds will be also evaluated according to the results, and to the compational time required.

References:

[1] Terpsidis I., Moukas A., Pergioudakis B., Doukidis G., Maes P. The potential of electronic commerce in reingeneering cosumer-retail relationships through intelligent agents. *Advances in Information Technologies: the business challenge*, IOS Press 1997.

[2] Rosenchein, J. Genesereth, M. Deals among rational agents., Proceedings of the 1985 *International Joint Conference on Artificial Intelligence*, pp. 91-99.

[3] Rassmuson, L. Rassmuson A. Jansson S. Proceedings of the 19th *National Information Systems Security Conference*. 1996

[4] Karlins M., Abelson H.I. *Persuasion, how opinion and attitudes are changed*. Crosby Lockwood & Son, 1970.

[5] Josang A. The right type of trust for distributed systems. Proceedings of *New Security Paradigms* Workshop, 1996.

[6] Castelfranchi C., Falcone R. Principles of trust for MAS. Proceedings of the *3rd Int. Conference on Multi-Agent Systems*, pp 72-79, 1998.

[7] Ebay. http://www.ebay.com

[8] OnSale. http://www.onsale.com

[9] Foner L., Yenta: A multi-agent referal-based matchmaking system. Proceedings of the *1st Int. Conf. on Autonomous Agents* pp. 301-307. 1997

[10] Blaze M., Feigenbaum J., Lacy J. Decentralized trust management. Proceedings of *IEEE Conference on Security and Privacy*, 1996.

[11] Marsh S. Formalising trust as a computational concept. Ph.D. Thesis, Univ. Stirling 1994.

[12] Banerjee B., Biswas A., Mundhe M., Debnath S., Sen S. Using bayesian networks to model agent relationships. *Applied Artificial Intelligence* 14 pp 867-879. 2000

[13] Yu B., Singh M.P. A social mechanism of reputation management in electronic communities. Proceedings of *Cooperative Information Agents* Conference, pp 154-165. 2000.

[14] Abdul-Rahman A. Hailes S. Supporting trust in virtual communities. Proceedings of the *IEEE International Conference on System Sciences*, 2000

[15] Zacharia G., Maes P. Trust management through reputation mechanisms. *Applied Artificial Intelligence* 14 pp 881-907 2000

[16] Carbó J., Molina J.M., Dávila J. A fuzzy model of reputation in Multi-Agent Systems *Advances in Fuzzy Systems and Evolutionary Computation* AI series, pp 22-27 Ed. Nikos Mastorakis, World Scientific and Engineering Society Press, 2001.

[17] Carbó J., Molina J.M., Dávila J. 1. A BDI Agent Architecture for Reasoning about Reputation, *IEEE Conference on Systems, Man and Cybernetics*. 2001.

[18] Carbó J., Molina J.M., Dávila J. Secure and fair introduction of newcommers in a fuzzy reputation system. IASTED Int. Conf. on Artificial Intelligence and Applications. 2002

[19] Sierra J.M., Carbó J. Fast and Secure Recommendation Protocols among members of a community of agents. *Electronic Notes in Theorical Computer Science*, vol 63, to be published in 2002

[20] Sabater J., Sierra C. REGRET: a reputation model for gregarious societies. *5th International Conference on Autonomous Agents*, pp 194-195, 2001

[21] Carbó J., Molina J.M., Dávila J. Agents Selecting Merchants with Good/Bad Memory. *UK Workshop on Computational Intelligence*, 2001