

Computing Confidence Values: Does Trust Dynamics Matter?

Joana Urbano, Ana Paula Rocha, and Eugénio Oliveira

Faculdade de Engenharia da Universidade do Porto
Rua Dr. Roberto Frias, 4200-465 Porto, Portugal
{joana.urbano, arocha, eco}@fe.up.pt

Abstract. Computational Trust and Reputation (CTR) systems are platforms capable of collecting trust information about candidate partners and of computing confidence scores for each one of these partners. These systems start to be viewed as vital elements in environments of electronic institutions, as they support fundamental decision making processes, such as the selection of business partners and the automatic and adaptive creation of contractual terms and associated enforcement methodologies. In this article, we propose a model for the aggregation of trust evidences that computes confidence scores taking into account dynamic properties of trust. We compare our model with a traditional statistical model that uses weighted means to compute trust, and show experimental results that show that in certain scenarios the consideration of the trust dynamics allows for a better estimation of confidence scores.

1 Introduction

Computational Trust and Reputation (CTR) systems are systems capable of collecting trust information about candidate partners and of computing confidence scores for each one of these partners. In this document, we envision trust as the confidence that the *trustier* agent has on the capabilities and the willingness of a candidate partner (*trustee*) in fulfilling its assigned tasks, in conformance to a given associated Service Level Agreement (SLA). CTR systems can be centralized, as adequate to electronic institutions and virtual organizations (VO), or decentralized, as adequate to extremely open environments where agents can enter and leave the society at any time.

Although practical examples of CTR systems do already exist (e.g. in e-commerce sites of eBay.com, Amazon.com, and Epinions.com¹), there are still many open questions in this research area. In fact, current work on trust and reputation has diversified in multiple subfields. In the theoretical domain, there is important work on trust and reputation as elements of social intelligence. Conte (2002) addresses the theoretical issues related to reputation and image in artificial societies and social simulation [1], and this cognitive model of reputation was recently extended in order to more thoroughly address the transmission of reputation [2]. In a more practical sense, a great deal of research effort is being put in the representation and aggregation of social

¹ <http://ebay.com>; <http://www.amazon.com>; <http://www.epinions.com>

evaluations into trust and/or reputation scores, which would serve as input to partner selection in electronic business scenarios. These models range from arithmetic means and weighted means ([3] [4] [5]), to Beta ([6]) and Dirichlet distributions ([7]), Bayesian approaches ([8] [9]), and trust learning approaches ([10] [11] [12]). Some of these models are implemented using cognitive based beliefs, desires and intentions (BDI) architectures ([5] [13]). A new trend of investigation in this area is the exploration of the business context to improve the decision making, raising significantly the number and type of information that the evaluator has in order to compute trust. However, few proposals have been made in this specific area ([14]).

Another area of little research work is the consideration of the *dynamics of trust* in the computation of confidence scores. Our hypothesis is that the use of an aggregation engine that encompasses the past experiences of the trustee agent and that accounts for fundamental dynamics of trust could allow for a better estimation of the trustee trustworthiness than probabilistic and statistical approaches that exist in the literature. Due to the relevance of this issue on our work, we dedicate the next section to the presentation of relevant dynamics of trust.

The remaining of this paper is structured as follows. In section 2, we present SinAlpha, a non-statistical aggregation engine that uses an S-shape curve to compute trust scores, taking into account three properties of trust dynamics: the asymmetry, the distinguishability of past evidences and the consideration of distinct maturity phases on the behaviour of target agents. Section 3 presents the experimental phase of our work. It introduces STexVM, a simulated virtual textile marketplace that we have developed in agent technology in order to evaluate the SinAlpha model and to compare it with other strategies. Then it proceeds with the presentation of the results of our experiments and with the analysis of these results. Section 4 presents the concluding remarks and future work.

1.1 The Dynamics of Trust

The evolution of trust over time was baptized by Elofson in 1997 [15] as the *dynamics of trust*, and was addressed one year later by Castelfranchi and Falcone [16]. An interesting formalization of the dynamics of trust is presented by Jonker and Treur in 1999 [17], who defend the need for a continuous verification and validation in the trust building process, and define six different types of trust dynamics:

- *Blindly positive*: the agent is unconditionally trusted or after a certain number or sequence of positive trust experiences (i.e. evaluated events) the agent reaches the state of unconditional trust and stays there for good;
- *Blindly negative*: the agent is unconditionally distrusted or after a certain number or sequence of negative trust experiences the agent reaches the state of unconditional distrust and stays there for good;
- *Slow positive, fast negative*: it takes a lot of trust-positive experiences to gain trust and it takes only a few trust-negative experiences to lose trust;
- *Balanced slow*: trust moves in slow dynamics in both positive and negative sense;
- *Balanced fast*: trust moves in fast dynamics in both positive and negative sense;
- *Slow negative, fast positive*: it takes a lot of trust-negative experiences to lose trust and it takes only a few trust-positive experiences to gain trust.

The authors also suggest that the dynamics of trust can be formalized through *trust evolution functions* (mathematical functions that relate sequences of experiences to trust representation) or through *trust update functions* (mathematical functions that relate a current trust representation and a current experience to the next trust representation). They formally define both functions and provide a set of interesting properties that can be associated to each one of the functions. Although this work is based on simple assumptions such as past direct experiences and binary evaluated events, it provides important considerations that shall be taken into account when designing an aggregation engine. Also, the slow positive, fast negative type of trust dynamics responds to the common sense idea that trust shall grow slower and decline faster, as interestingly put in the famous words of the English poet Alexander Pope: ‘*At every word a reputation dies*’. At this respect, Marsh [18] also strongly suggested to penalize deceit behaviour stronger than to award the cooperative ones, as in the real world it is easier to loose, than to gain trust.

Melaye and Demazeau (2005) [19] further explore the dynamics of trust, proposing a Bayesian trust formalism based on Castelfranchi and Falcone’s cognitive model. They use a Kalman filter to address two dimensions of the trust dynamics: the asymmetric increase/decrease of trust and the inherent speed of switching from trust to distrust and *vice versa*, which they name *inertia*; and the *erosion* of trust that happens due to the absence of new observations. In their model, the outcome of an execution is statistically dependent of previous executions, supporting, therefore, the mentioned trust dynamics. The introduction of the erosion dimension is of particular interest, as current trust and reputation systems tend to omit this characteristic, particularly those whose aggregation engine is based on statistical operations. However, the proposed Bayesian presents some drawbacks. In one hand, the model seems not to be scalable in the case of several beliefs and several source beliefs, and the authors assume statistical independence between each one of the belief and source beliefs’ levels. Also, as the authors indicate, the inertia of trust and distrust is fixed *a priori* by a specialist, requiring one instance of the model per context. Finally, the proposed model seems to be *too sensitive* in relation to single occurrences of deceptive behaviour. In fact, in one experiment described in [19], a single negative observation that happens after a high number of previously observed positive experiences makes the trust level to decrease sharply, after which it takes a long sequence of positive observations to getting back to the previously trust value. In our opinion, this strong penalization does not reflect the real world response to one exceptional bad result of a previously trustable partner.

2 The SinAlpha Aggregation Engine

As already mentioned, we are interested in designing and implementing mechanisms that allow for an expressive representation of the dynamics of trust, when aggregating trust evidences. Particularly, we are interested in the *asymmetry* property, that stipulates that trust is hard to gain and easy to lose; in the *maturity phase* of targets property, where the slope of growth can be different in different stages of the partner trustworthiness; and in the *distinguishability* property of past behaviour. The sigmoid curve represented in Figure 1 presents interesting characteristics that seem to fit the

desideratum well. For simplicity, we assume that the available information about a candidate partner is given by a central trust authority (e.g. a CTR service that serves the VO), and that it takes the form of binary values, either representing past successful (1) or violated (0) contracts by the partner.²

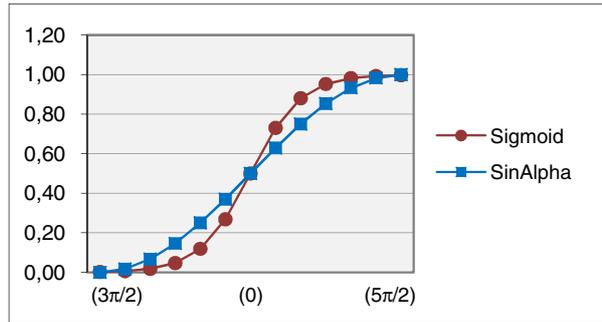


Fig. 1. Two S-shape curves, one exponential (Sigmoid) and one trigonometric (SinAlpha)

The constructing of trust for this partner using the sigmoid curve implies a slow growth upon positive results when the partner is not yet trustable, it accelerates when it is acquiring confidence, and finally slows down when the partner is considered trustable (i.e., in the top right third of the curve). The decrease movement upon negative results follows the same logic. However, we intuitively feel by graphically analysing the curve that it permits a probably too soft penalisation of partners that proved to be trustable but that failed the last n contracts. Therefore, we lightly soften the slope of the sigmoid shape at the top and bottom thirds of the curve, by using instead the trigonometric formula presented in (1) and depicted in Figure 1, with the name of SinAlpha.

$$y(\alpha) = \delta \cdot \sin \alpha + \delta, \quad \alpha_0 = 3\pi/2, \tag{1}$$

$$\alpha = \alpha + \lambda \cdot \omega.$$

In the formula above, δ is a constant value of 0.5, and α ranges from $3\pi/2$ to $5\pi/2$, allowing for aggregated trust scores within the range [0, 1]. The incremental step of α is also shown in (1); ω represents the pace of trust growth (we assume the value of $\pi/2$ in our experiments), and λ is the parameter of the incremental step that allows to differentiate between positive and negative results (in our experiments, λ equals +1 for each positive result to be aggregated, and -1.5 for each violated contract). This way, in each one of the three stages of trust construction, trust grows slower and decreases faster. At this point, we must remind our interest in studying how a curve like the one we propose, which, in a certain way, ‘encompasses’ the historical behaviour

² We use these two assumptions in our experiments, although our proposed aggregation engine might be extended in the future to more complex and diversified representation of trust information. In the same way, the aggregation engine might be used in decentralized systems, to aggregate information from distinct sources of information (e.g. reputation and image).

of the partner under evaluation, is able to catch the dynamics of trust in the presence of certain partners' patterns of behaviour. We are also willing to know how this model can be compared with the common statistical approach that aggregates trust information using weighted means.

3 Experiments

3.1 The STexVM System

In order to run our experiments, we developed the STexVM system. This is a simulated virtual marketplace for trading textile goods that aims to ensure reliable transactions, in a sense that it is able to detect business partners that in some moment start behaving in a defective way. The simulated environment is based on existent online virtual marketplaces where buyers and sellers in the textile and fashion industry can post buying and selling leads (e.g. the Fibre2Fashion marketplace³). It follows the multi-agent paradigm, and is implemented over Jade platform, using the standard behaviours of Jade and FIPA performatives and interaction protocols⁴. The key agents in this environment have the roles either of buyers or suppliers (Figure 2).

At each round, a buyer issues a call for proposal (cfp) stipulating a specific good and associated quantity that needs to be provided, and each candidate partner responds indicating the quantity it is able to provide in the present business opportunity, or refusing the offer. A contract-net like negotiation occurs, and the buyer selects a number $n > 0$ of partners that optimizes the expected utility $E(u)$, using equation (2).

$$E(u) = \arg \max_i \text{ for each } i \sum_j \text{util}_j * \text{trust}_j . \quad (2)$$

In the equation above, i stands for the possible combinations of suppliers' proposals that fit the quantity specified in the current cfp, not exceeding it; j represents the suppliers considered in each of these combinations, and trust_j is the confidence score computed for supplier j at selection time. Finally, util_j is the quantity proposed by each supplier j in the round, normalized by the quantity specified in the cfp, i.e., $\text{quant}_j/\text{Quant}$. In our system, a buyer can accept less quantity than the maximum quantity (Quant) defined in the cfp, but it cannot exceed Quant . Also, a buyer cannot accept partial quantities of the received bids.

Each supplier that enters the simulated virtual marketplace sells two different types of fabric (e.g. cotton and chiffon). These and their associated quantities (e.g. 180,000 meters) are randomly assigned at creation time. Buyers are characterized by the good and quantity they need to purchase, also randomly picked up at creation time. The remaining agents of the STexVM system are the Agent Simulation Manager, who manages the configuration parameters related with buyers and suppliers; the Agent DF, which registers competences of buyers and suppliers; and the Agent CTR, which gathers information about the performance of suppliers and computes their confidence scores on-demand, when requested by the buyers. Figure 2 illustrates the relation between these agents.

³ <http://fibre2fashion.com/>

⁴ Jade: <http://jade.tilab.com/>; FIPA: <http://www.fipa.org/>

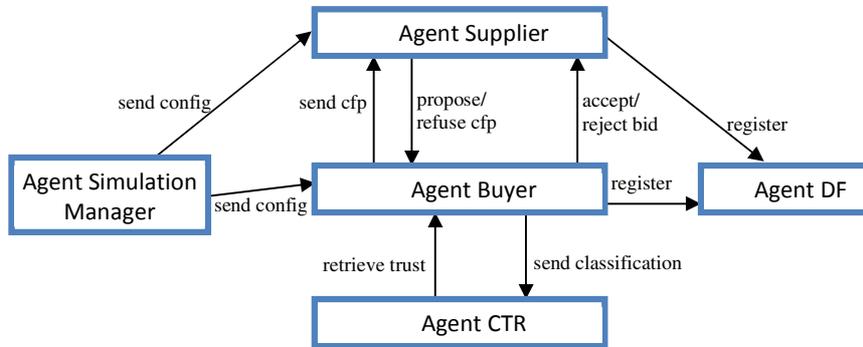


Fig. 2. Interactions between agents in the STexVM system

3.2 Approaches in Evaluation

In the following experiments, a buyer agent can be assigned one of four different approaches for selecting partners. The SINALPHA approach uses equation (1) to estimate a truth score for each candidate partner and weights the resulting score with a recency factor. The ASYM+ approach is a similar, former model that we proposed in [20] that also accounts for the dynamics of trust. The WMEAN approach uses an aggregation engine that computes the mean of the last 10 results weighted by the recency of these results (cf. Huynh, 2006 [21]). As mentioned earlier, there are several CTR models that use weighted means to aggregate social evaluations, therefore the WMEAN approach will allow us to compare SINALPHA with one model that is disseminated in the trust and reputation community. Finally, the QUANT model selects partners by the quantity they are able to provide, and does not take into consideration the trust values of the suppliers.

For all the models, in the first rounds of each experiment the buyers start to explore the space of available candidate partners, by randomly selecting the partners, and after some rounds they progressively increase the exploitation by selecting partners based on the selected model. In the current experiments (Table 2) the exploration phase ends up at round 39 (of 60), for all buyers, meaning that their selection decision relies exclusively on the adopted model after round 40.⁵

3.3 Experimental Methodology

In order to evaluate the approaches described above, we consider that the candidate partners have different behaviours and are divided accordingly into categories “S_A”, “S_B”, and “S_C”. The behaviour of a supplier is related to the results of the contracts it makes, during its lifecycle, with buyer agents. A behaviour is assigned to each supplier at its creation time, following a uniform distribution over the three possible categories. We consider that the capacity of each type of suppliers in fulfilling the contract is modelled by a Markovian process with two states (1 and 0, standing for contract

⁵ In future versions of STexVM, we will allow for the system to keep results between experiments, avoiding the need to bootstrap the system every time an experiment is run.

fulfilment and contract violation, respectively) and transition probabilities P11 (Fulfilment-to-Fulfilment) and P01 (Violation-to-Fulfilment). In these experiments, we consider two distinct populations, A and B, as defined at Table 1.

Table 1. Transition probabilities. Initial probabilities $P_0 = P_1 = 0.50$ for both populations

	Type "S _A "		Type "S _B "		Type "S _C "	
	P11	P01	P11	P01	P11	P01
Pop A	0.90	1.00	0.80	0.75	0.50	0.50
Pop B	0.90	0.20	0.90	0.20	0.80	0.60

As can be seen from the table above, in population A, suppliers of type S_A have high probability of success and never fail two contracts in a row (once P00 is zero). Types S_B and S_C correspond to progressively worse behaviours. With this population, we want to evaluate the capacity of each strategy in choosing the best partners, i.e., partners of type S_A. Population B, on the other hand, presents a bursty-like pattern of behaviour, where candidate partners of types S_A and S_B generally fulfil several contracts in a row, but when they fail a contract they generally enter in a long burst of violated contracts. Type S_C presents smaller bursts of both positive and negative behaviour. With this population, we pretend to evaluate the performance of each strategy in avoiding long sequences of negative results and in abandoning a good provider when it starts to behave in a deceptive way.

In every experiment, we instantiated 16 suppliers and 8 buyers: two of type SINALPHA, two of type ASYM+, two of type WMEAN and the remaining two of type QUANT. Every buyer was allowed to run 60 rounds, corresponding to the launching of 60 different cfps. Every experience was run 12 times. Finally, the utility gained by each buyer at each negotiation round was recorded, and at the end of the experiments the average utility of a buyer and the corresponding standard deviation were evaluated for each one of the considered approaches. The average utility captures the capacity of the buyer in selecting good partners, and, this way, allows for the evaluation of the performance of each one of the three approaches. Table 2 presents compact data about the experiments.

Table 2. Values and parameters used in the experiments

Fabrics and Quantity	Chiffon, Cotton; 180000
# buyers	2 SINALPHA, 2 ASYM+, 2 WMEAN, 2 QUANT
# of sellers	16
Types of sellers	Chosen upon a uniform distribution over the types {"S _A ", "S _B ", "S _C "}
# issued CFP per buyer, per run	60
# runs per experiment	12
Exploit/Exploration formula	Uniform distribution over $f(x)$, where $f(x) = 100 - \text{round}_i * 7$, $f(x) = 5$, if $(100 - \text{round}_i * 7 < 10)$, $f(x) = 0$ if $\text{round}_i > 40$; $0 < i < 60$
SinAlpha parameters	$\delta = 0.5$; $\omega = \pi/2$; $\lambda = +1$ for successful contracts, and $\lambda = -1.5$ for violations

3.4 Results

We used three different metrics to evaluate the performance of the three approaches in evaluation in the experiments with population A. First, we measured the utility gained by each buyer in the last 20 rounds in each experiment, and averaged the results obtained for each approach over the 12 experiments. The results concerning the average utility showed that the both models that accounts for the described dynamics of trust outperformed the other two models (SINALPHA: 92.1%; ASYM+: 93.3%; WMEAN: 87.1%; QUANT: 79.8%). Then, we counted the number of violated contracts per buyer in its last 10 transactions, and averaged this number per approach and over the 12 runs of the experiment. The results show a clear advantage of SINALPHA in avoiding partners with past bad experience: 11.7% of violated contracts (vc), with standard deviation (sd) of 0.78, when compared to ASYM+ (12.1% vc, 1.38 sd), WMEAN (18.3% vc, 1.37 sd) and QUANT (22.1% vc, 3.20 sd).

Finally, we counted the number of suppliers of types S_A , S_B and S_C that were chosen by each buyer in its last 20 transactions, and averaged this number per approach, over the 12 runs. In this step, we intend to further understand the differences between the proposed SinAlpha curve and the weighted mean approach, and, this way, we only present the results obtained for SINALPHA and WMEAN. The results showed that with the SINALPHA approach the buyers were able to choose the best suppliers (of type S_A) 91% of the times, and suppliers of type S_B the remaining 9%. On the other hand, buyers that used the WMEAN approach were less effective in choosing suppliers of type S_A (they did it 75% of the times) and they even choose bad suppliers of type S_C 5% of the times. Suppliers of type S_B were chosen 20% of the times.

We repeated the first two procedures described above with population B, and the results are as follows. SINALPHA got an average utility of 79.8%, outperforming ASYM+ (78.8%) and QUANT (62.3%), but underperforming the WMEAN approach, that achieved an average utility in the last 20 rounds of 83.3%. Concerning the average of violated contracts in the last 10 transactions of each buyer, SINALPHA performed a little better than the remaining approaches, getting 22.9% of violated contracts (vc) and standard deviation (sd) of 2.11, against the results of ASYM+ (25% vc, 1.91 sd), WMEAN (24.6% vc, 2.75 sd) and QUANT (38.3% vc, 3.31 sd).

3.5 Interpretation of the Results

Starting with Population A, we verified that the QUANT approach gets the worse results as expected, as it is not able to differentiate between partners of types S_A , S_B , and S_C . We also verified that SINALPHA and ASYM+ tend to perform in a similar way, with SINALPHA slightly outperforming ASYM+. This is because both models use sigmoid like curves and accounts for the same described trust dynamics. Because of space concerns, we relegate a further comparison of these two our models to a future paper. Finally, we observed that the SINALPHA approach outperforms the WMEAN approach, particularly when we attend to the figures related to the violated contracts in the last 10 interactions and to the capacity of each approach in selecting good (S_A) partners. By analyzing the traces of the experiments, we realize that the SINALPHA strategy selects primarily partners of type S_A , while the WMEAN equally selects partners of types S_A and S_B . The difference between both approaches is that in

SINALPHA all the historical path is taken into account in the process of trust construction, and partners have to accumulate several good experiences in the past until they are able to get an average to high trust score. In opposition, by aggregating up to the last N results, the WMEAN approach allows the selection of partners with fewer past events. In reality, we verified that the bad choices of WMEAN on population A were related to the selection of S_B and S_C partners in two distinct situations: i) when they were selected with less than 10 past results (e.g. the pattern of the previous evidences to the time of selection where $V-F-F-V-F-F$, where V means a violated contract and F a fulfilled contract); and ii) when they showed an intermittent pattern followed by a short number of positive evidences (e.g. $F-V-V-F-V-F-V-V-F-F-F-F-F$). This last pattern of behaviour is indeed severely punished by the SINALPHA approach, where violations weight more than fulfilments (therefore penalizing undesirable intermittent patterns), and where the last five positive evidences are not sufficient to ‘push’ the confidence level of the partner to the second third of the SinAlpha curve.

In the experiments with population B, we intended to study the performance of the SINALPHA and the WMEAN strategies in the presence of extreme partners’ behaviour, particularly the cases where good partners, which have been successfully in fulfilling their obligations, suddenly start having systematic deceptive behaviour. By analysis of the traces of the experiments, we realized that both strategies act quite differently as they tend to select different partners in similar conditions. In fact, the WMEAN strategy privileges recency and puts a limit to the historical analyses (in our experiments, it aggregates the last 10 contract results). In one hand, this permits that a candidate partner with few past results (let us say 6 past results, while majority positives), is chosen in detriment of a partner that has been reliable for a long time but that violated the last 2 to 4 contracts. As we are selecting partners using too few results, a considerable risk is associated to the partner’s selection using WMEAN.⁶ On the other hand, in similar conditions, the SINALPHA approach does not select the described partners, as the SinAlpha curve encompasses a “growing path” that partners shall run until they acquire middle to high trust (i.e., until they reach the last two thirds of the curve). This means that in the scenario of population B, the SINALPHA approach has bigger tendency to enter a burst of deceptive behaviour and is somewhat slower in penalizing good partners that inverted their behaviour. However, we detected another problem associated to the WMEAN approach, which helps to explain the results in terms of violated contracts: as this strategy privileges recency, it actually assigns high trust levels to candidate partners that systematically behaved deceptively in the past, had no classification for a long time, and then got one positive classification in the present. Using common sense, we can deduce that this kind of behaviour encompasses a high degree of risk; however, we realized that in the described cases the WMEAN approach chooses these partners in detriment of more stable partners that happened to fail the last couple of contracts.

Although not mentioned in section 3.3, we run a different type of experiment using population A and the Repast agent simulator⁷. In this experiment, we run the SINALPHA and the WMEAN approaches separately; i.e., we maintained the number of

⁶ A confidence value could be used along with the computed trust score. However, these experiments aim to evaluate the aggregating processes by themselves.

⁷ <http://repast.sourceforge.net/>

suppliers (16), but all the 8 buyers were either using SINALPHA or WMEAN. Each run took 100 rounds. In the first 40 rounds, the selection of suppliers was done randomly, and in the last 60 rounds the selection was done taking into account the approach used by the buyers. At round 70, for both approaches, all suppliers of type S_A abruptly changed their behaviour and assumed the characteristics of type S_C , and the remaining suppliers kept their initial behaviour. With this experiment, we intended to study the *abuse of prior information* scenario defined in [8]. The results that we obtained showed a similar capacity of SINALPHA and WMEAN in detecting and penalizing the change of behaviour of suppliers that were originally of type S_A . However, the most interesting result of this experiment was the capacity of SINALPHA buyers in adapting to the situation by massively choosing suppliers of type S_B after round 70. Concerning WMEAN buyers, although they also increased the number of selected S_B partners after round 70, they show an undesirable side effect of also considerably increasing the selection of partners that were originally of type S_C . This behaviour shall be analysed with further detail in future work.

4 Concluding Remarks and Future Work

The work presented in this paper started by empirically searching a mathematical function that would allow to aggregate evaluations on a given partner that encompasses the evolutionary performance of the target. This search was driven by economical common sense. For example, and simply putting, a business player would certainly distinguish between the following patterns of behaviour: *good-good-good-good-bad-bad-bad-bad*, *good-bad-good-bad-good-bad-good-bad* and *bad-bad-good-bad-good-good-bad*. We have put thorough attention in the sequences of possible results, always taking into consideration common sense about business notions. For instance, a partner that ever succeed with its obligations and achieved a high degree of trustworthiness should not be severely punished if he accidentally is not successful in the last actual transaction; a partner that succeed the first two obligations but does not entered any other transaction after that cannot be considered highly reputed, as there is not enough information on the past he had carried out until be considered reputable. Also, a partner that achieved a given trust level should not maintain this level if he starts behaving in an intermittent way. Following this reasoning process, we came across the SinAlpha function.

The experiences we run on two different populations allow us to conclude that the SinAlpha function actually gets better results than a weighted mean by recency approach, because it takes into account the dynamics of trust. Also, the results obtained seem to show that there is still margin to improve the performance of the proposed mechanism by adapting/learning the values of parameters λ and ω to the perceived patterns of the current populations. In fact, the next phase of our work would be dedicated to this topic, and to the inclusion of the erosion property of trust in our approach. Different type of behaviour patterns shall be defined, as well as experimental procedures (e.g. the use of the metric *converge speed* defined in [8]).

In the same way, we will continue progressively improving our approach, and several research challenges would certainly be presented. Namely, we propose as future

work to identify and to categorize patterns of behaviour as new target evidences appear, through the usage of clustering techniques; and to consider multi-attribute evaluations (e.g. price, delivery time, and quality).

Acknowledgements. The first author enjoys a PhD grant with reference SFRH/BD/39070/2007 from the Portuguese Fundação para a Ciência e a Tecnologia.

References

1. Conte, R., Paolucci, M.: *Reputation in Artificial Societies: Social Beliefs for Social Order*. Kluwer Academic Publishers, Dordrecht (2002)
2. Salvatore, A., Pinyol, I., Paolucci, M., Sabater-Mir, J.: Grounding Reputation Experiments. A Replication of a Simple Market with Image Exchange. In: *Proceedings of the Third International Model-to-Model Workshop*, France, pp. 32–45 (2007)
3. Ramchurn, S., Sierra, C., Godo, L., Jennings, N.R.: Devising a trust model for multi-agent interactions using confidence and reputation. *Int. J. Applied Artificial Intelligence* 18, 833–852 (2004)
4. Sabater, J.: *Trust and Reputation for Agent Societies*. Number 20 in *Monografies de l'institut d'investigacio en intelligençia artificial*. IIIA-CSIC (2003)
5. Carbo, J., Molina, J., Davila, J.: A BDI Agent Architecture for Reasoning about Reputation. In: *IEEE International Conference on Systems, Man, and Cybernetics* (2001)
6. Jøsang, A., Ismail, R.: The Beta Reputation System. In: *Proceedings of the 15th Bled Electronic Commerce Conference*, Sloven (2002)
7. Reece, S., Rogers, A., Roberts, S., Jennings, N.R.: A Multi-Dimensional Trust Model for Heterogeneous Contract Observations. In: *22nd AAAI Conf. on Artificial Intelligence* (2007)
8. Zacharia, G., Maes, P.: Trust management through reputation mechanisms. *Applied Artificial Intelligence* 14(9), 881–908 (2000)
9. Haller, J.: A Bayesian Reputation System for Virtual Organizations. In: *Dagstuhl Seminar Proceedings 06461, Negotiation and Market Engineering* (2006)
10. Hang, C., Wang, Y., Singh, M.P.: An adaptive probabilistic trust model and its evaluation. In: *Proceedings of the 7th international Joint Conference on Autonomous Agents and Multiagent Systems*, Estoril, Portugal, May 12 – 16, 2008, vol. 3 (2008)
11. Erete, I., Ferguson, E., Sen, S.: Learning task-specific trust decisions. In: *Procs. 7th Int. J. Conf. on Autonomous Agents and Multiagent Systems*, Portugal, vol. 3 (2008)
12. Smith, M., desJardins, M.: Learning to trust in the Competence and Commitment of Agents. *Autonomous Agent Multi-Agent Systems* 18, 36–82 (2009)
13. Sabater, J., Paolucci, M.: On Representation and Aggregation of Social Evaluations. *Computational Trust and Reputation Models*, *Int. J. Approx. Reasoning* (2007)
14. Rettinger, A., Nickles, M., Tresp, V.: A statistical relational model for trust learning. *Procs. 7th Int. J. Conf. on Autonomous Agents and Multiagent Systems* 2 (2008)
15. Elofson, G.: Developing Trust with Intelligent Agents: An Exploratory Study. In: *Proceedings of the First International Workshop on Trust*, pp. 125–139 (1998)
16. Falcone, R., Castelfranchi, C.: Principles of trust for MAS: cognitive anatomy, social importance, and quantification. In: *Procs. Int. Conference on Multi-Agent Systems* (1998)
17. Jonker, C.M., Treur, J.: Formal Analysis of Models for the Dynamics of Trust Based on Experiences. In: Garijo, F.J., Boman, M. (eds.) *MAAMAW 1999*. LNCS, vol. 1647, pp. 221–231. Springer, Heidelberg (1999)

18. Marsh, S.P.: Formalising Trust as a Computational Concept, PhD Thesis, University of Stirling (1994)
19. Melaye, D., Demazeau, Y.: Bayesian Dynamic Trust Model. In: Pěchouček, M., Petta, P., Varga, L.Z. (eds.) CEEMAS 2005. LNCS (LNAI), vol. 3690, pp. 480–489. Springer, Heidelberg (2005)
20. Urbano, J., Rocha, A.P., Oliveira, E.: Trust Evaluation for Reliable Electronic Transactions between Business Partners. In: Fischer, K., Muller, J.P., Odell, J. (eds.) Proc. The AAMAS 2009 Workshop on Agent-based Technologies and applications for enterprise interOPERability (ATOP), Budapest, pp. 85–96 (2009)
21. Huynh, T.D., Jennings, N.R., Shadbolt, N.R.: An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems* 13(2), 119–154 (2006)