

A Trust Aggregation Engine that Uses Contextual Information

Joana Urbano, Ana Paula Rocha, Eugénio Oliveira

Faculdade de Engenharia da Universidade do Porto – DEI, Rua Dr. Roberto Frias, 4200-465
Porto, Portugal

LIACC- Laboratory for Artificial Intelligence and Computer Science at University of Porto
{joana.urbano, arocha, eco}@fe.up.pt

Abstract. Trust estimation of partner agents is considered a fundamental step in the process of selecting partners. In previous work, we proposed SinAlpha, an agent-based aggregation engine that computes the trustworthiness of candidate partners by aggregating their historical contractual evidences, taking into account important properties of the dynamics of trust. In this paper, we further argue on the benefits of the trust dynamics, and we describe the *Contextual Fitness* component, a new element of our computational trust system that brings contextual information into the trust loop, by relating the estimated trustworthiness of partners to the specific current business situation. Experimental results show that such an approach significantly improves partners selection, due to its ability in detecting the flaws of the target population, even when the available historical evidences are scarce.

1 Introduction

Computational Trust and Reputation (CTR) systems are systems capable of collecting trust information about candidate partners and of computing confidence/trust 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 contract established between both parties.

Although practical examples of CTR systems do already exist (e.g. in e-commerce sites of eBay.com, Amazon.com, and Epinions.com¹), several research work on this area is still in progress and has diversified in multiple subfields. Concerning the particular subfield that addresses the representation and the aggregation of social evaluations into trust and/or reputation scores, there exists models that range from arithmetic means and weighted means ([1] [2] [3]), to Beta ([4]) and Dirichlet distributions ([5]), Bayesian approaches ([6] [7]), and trust learning approaches ([8] [9] [10]). Some of these models are implemented using cognitive based beliefs,

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

desires and intentions (BDI) architectures ([3] [11]). A new trend of investigation on 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 the trust. However, few proposals have been made on this specific area ([12]).

Another interesting branch of research work considers the *dynamics of trust* in the computation of confidence scores. The evolution of trust over time was baptized by Elofson in 1997 [13] as the *dynamics of trust*, and was addressed one year later by Castelfranchi and Falcone [14]. An interesting formalization of the dynamics of trust is presented by Jonker and Treur in 1999 [15], who defend the need for a continuous verification and validation in the trust building process. In their work, they define the *slow positive, fast negative* type of trust dynamics that we consider very important, which says that it takes a lot of trust-positive experiences to gain trust and it takes only a few trust-negative experiences to lose trust. At this respect, Marsh [16] 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) further explore the dynamics of trust, proposing a Bayesian trust formalism based on Castelfranchi and Falcone's cognitive model [17]. 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.

Our current work explores two important functionalities that we think shall be present in a trust aggregation engine. First, we consider that an aggregation engine that encompasses the past experiences of the trustee agent and that accounts for fundamental dynamics of trust allows for a better estimation of the trustee trustworthiness than probabilistic and statistical approaches that exist in the literature. In section 2, we describe three of such trust properties, and present the SINALPHA aggregation algorithm that aggregates these properties.

The second functionality that we are currently working on is the ability of the aggregation engine to look at the past trust evidences of a given target agent and to infer, based on the past behavior of the agent, if its profile fits the current specific business need, described in the form of a call for proposals (CFP); we say in this case that we use additional *contextual information* to improve the reliability of the estimated trust values of candidate partners.

The remaining of this paper is structured as follows: section 2 presents the SinAlpha curve, an aggregation algorithm that incorporates properties of the dynamics of trust and that was already implemented and tested by the authors. Section 3 presents the Contextual Fitness module, a component that brings contextual information to the trust loop in an innovative way. Section 4 presents the experiments performed in order to evaluate the Contextual Fitness module. Finally, section 5 concludes the paper and presents future work.

2 The SinAlpha Aggregation Curve

In [18], we described SinAlpha, an S-like, sin-based aggregation curve (see Figure 1) that allows for an expressive representation of the dynamics of trust, particularly, implementing the following properties:

- *Asymmetry* property, that stipulates that trust is hard to gain and easy to lose;
- *Maturity* property, that measures the maturity phase of the partner considering its trustworthiness, where the slope of growth can be different in different stages of the partner trustworthiness;
- *Distinguishably* property, that distinguishes between possible different patterns of past behaviour.

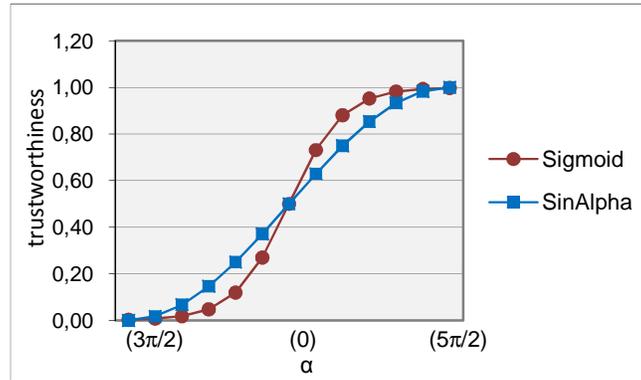


Fig. 1. Two S-shape curves, one exponential (Sigmoid) and one trigonometric (SinAlpha)

For simplicity, we assume that the historical trust evidences of a given candidate partner are its past contractual outcomes, as provided by a trust authority (e.g. a CTR service integrated in an electronic institution environment). However, this is not a hard assumption and other information sources (such as social evaluations) from either central or distributed architectures can be used. In all cases, trust evidences are represented as follows:

$$\langle A_C, A_P, At_1..At_n, t, Res \rangle,$$

where:

- A_C is the client agent, i.e., the agent that received a product or service from the agent being evaluated (the trustier);
- A_P is the agent that provided the good to the client agent (the trustee);
- $At_1..At_n$ are the n attributes established in the contract, each one described as attribute-value pairs (e.g. good=cotton, quantity=360000, deliveryTime=7, for a contract that involves the provision of 360000 meters of cotton in 7 days);
- t is the timestamp of the contract;
- Res is the result of the contract. Currently, it takes the form of binary values, either

representing a successful (1) or a violated (0) contract by the provider partner.²

The constructing of a value of trust for a particular provider agent using the SinAlpha curve implies a slow growth upon positive results when the partner is not yet trustable, an acceleration when it is acquiring confidence, and a slow decay when the partner is considered trustable (i.e., in the top right third of the curve), allowing for the definition of three different trust maturity phases (the Maturity property). The decrease movement upon negative results follows the same logic, although the mathematical formula subjacent to the Sinalpha curve presents a parameter that permits that trust grows slower and decays faster (the Asymmetry property). One can argue that we could use other S-like curves instead of a sin-based one, such as the Sigmoid curve, illustrated in Figure 1. However, we intuitively feel by graphically analysing the Sigmoid curve that it permits a probably too soft penalisation of partners that proved to be trustable but that failed the last n contracts. This can happen accidentally (e.g. due to an unexpected shortage of good or to distribution problems), but it is also described in the literature as a typical behaviour of deceptive provider agents, who tend to build up a trustworthy image using simple contracts and then violate bigger contracts exploring the acquired trustworthiness.

2.1 Evaluation of SinAlpha’s Trust Dynamics Properties

In [18], we provide a detailed description of the SinAlpha curve, as well as an experimental evaluation of its behaviour. In this section, we summarize the main conclusions we obtained when we experimentally compared SinAlpha (the SINALPHA approach) to a weighted mean by recency approach (that we named WMEAN), a common approach seen in literature for trust and reputation aggregating engines (see, for instance, [19]).

In this work, we explored three different scenarios. In the first scenario, we wanted to compare the capacity of both approaches in differentiating between different types of provider agents, namely, the capacity of primarily choosing ‘good’ suppliers that with a high probability do not violate a contract. In such a scenario, we observed that the SINALPHA approach outperforms the WMEAN approach in its capacity of selecting ‘good’ partner agents, in one hand, and in avoiding ‘bad’ partners, in the other hand. One difference between both approaches resides in the fact 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 (the maturity property). In opposition, the WMEAN approach allows the selection of partners with fewer past events. For instance, analyzing the traces of the experiments, we verified that some bad choices of WMEAN happened when the algorithm selected partners with rather few contractual past evidences (e.g. the pattern of the previous evidences to the time

² Once again, we use this assumption in our experiments, although we leave for future work the extension of our proposed aggregation engine to include more complex representations.

of selection were $V-F-F-V-F-F$, where V means a violated contract and F a fulfilled contract).

Another difference between the two approaches is due to the *asymmetry* property of SINALPHA. This seems to be particularly important when identifying and acting upon partners that show intermittent behaviour (e.g. $F-V-V-F-V-F-V-V-F-F-F-F$). This last pattern of behaviour is indeed severely punished by the SINALPHA approach, where violations weight more than fulfillments (therefore penalizing undesirable intermittent patterns), and where the last few positive evidences are not sufficient to ‘push’ the confidence level of the partner to the second third of the SinAlpha curve.

In the second scenario, we intended to study how SINALPHA and WMEAN react in the presence of extreme partners that have a bursty-like behavior (i.e. that switch between sequences of good and deceptive behaviour). By analysis of the traces of the experiments, we realized that both approaches act quite differently as they tend to select different partners in similar conditions. The main point to consider here is that WMEAN, by privileging recency, 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. I.e., WMEAN-like approaches can *forgive too fast* in certain temporal scenarios. One could argue here that this forgiveness issue is solved by increasing the size of the window used (i.e. the number of the last past evidences considered); however, in our experiments we found it hard to select the optimal window size, as it deeply depends on the frequency of the contracts (historical evidences) made in the past. The forgiveness question does not apply to SINALPHA, due to the action of the *maturity* property; however, we realized that SINALPHA has a somewhat bigger tendency to enter a burst of deceptive behaviour and that it can be slower in penalizing good partners immediately after they invert their behavior.

Finally, the last scenario intended to study the *abuse of prior information* scenario defined in [6], where ‘good’ partners definitely invert their behavior after a given number of iterations. The results that we obtained showed that SINALPHA outperforms WMEAN in detecting and penalizing the change of behaviour of originally ‘good’ partners, while WMEAN showed a significantly higher tendency to choose ‘bad’ partners than SINALPHA.

2.2 Remarks about SinAlpha

Taking into account all the experiments performed, we can conclude that the three *properties* of the dynamics of trust embedded in SINALPHA are effective in distinguishing between different types of target agents, therefore in detecting and acting upon undesirable agents’ behaviours. Namely, the asymmetry property penalizes intermittent behavior, the maturity property avoids selection of partners who did not prove to be trustable enough, and the distinguishable past property avoids the

phenomenon of forgiveness described above.³ In these experiments, we could not evaluate, however, the potential full benefits of the SinAlpha shape against simpler curves that do show similar trust dynamics properties (e.g. curves with linear shape).

In fact, as stated previously, SinAlpha considers different growth/decay slopes in different stages of the trustworthiness acquisition of a target agent, and it also presents a sigmoid-like shape. The choice of this shape was based on the concept of *the hysteresis of trust and betrayal*, from [21]. In this work, the author proposes a path in the form of a hysteresis curve where trust and betrayal happens in the balance between the trustworthiness of a self and the trust placed on the self. The SinAlpha curve simplifies the hysteresis approach by using just one curve for both trust and betrayal representation and considering three different growth/decay stages: Creating Trust (first third of the curve), Trust is Given (second third of the curve), and Taking Advantage (last third of the curve).

Performance tests of the SinAlpha representation against a simpler curve were performed. This new simpler curve uses λ and ω parameters from SinAlpha (cf. [18]) to update the trustworthiness value of target agents, but it lacks the softness round curve at Creating Trust and Taking Advantage extremes. The results of these experiments (to appear in a thorough study about SinAlpha) show similar performance of both curves in the tested scenarios. Therefore, we conclude that we need different, much more complex models of target population to further study the impact of the sigmoid-like shape of SinAlpha on its capability of distinguishing between partners. We leave this topic for future work.

3 The Contextual Fitness Model

3.1 The Contextual Fitness Concept

Trust and reputation estimations help the trustier agent to predict how well a given candidate partner will execute a task and to compare between several candidate partners. However, there are some questions that a real-world manager would pose before making a decision that cannot be answered by simply aggregating available trust evidences into trust and reputation values. These questions involve somehow a certain level of *intuition*. We propose to first analyze three scenarios that might occur in real world business and that would help to understand this concept.

In the first scenario, an agent may decide to exclude from selection a candidate partner with which it had never entailed business before but that it knows (by other means) that rarely fails a contract, just because the agent intuitively fears that this partnership would not be successful. For example, a high tech company may fear to select a partner from a country of origin without high technology tradition, even

³ Considering this last property, we have a somewhat different view than the one presented in [20], where the authors state that the aggregation of evaluations shall not depend on the order in which these evaluations are aggregated.

though this partner has proved high quality work in the desired task in the recent past. We call this situation the *intuitive fear*. For this scenario, it would be desirable that the selector agent could reason taking into account additional contextual information about the characteristics of the entity represented by the candidate agent. For instance, the presence of key figures such as the annual turnover or the number of employees of the entity would allow the selector agent to better know the entity. Also, the establishment of argumentation between both parties is a real-world procedure that could be automated into the computational decision process. We will address the *intuitive fear* situation in future work.

In the second scenario, the agent may decide to exclude from selection a candidate partner that is currently entering the business, for which there is not trust and/or reputation information yet. This scenario deals with the problem of newcomers, for which there is no information about prior performance, and we name it *absence of knowledge*. Some authors do address the question of newcomers by considering the complementary use of a diversified set of information sources. As an example, in [19] it is suggested that in these cases the use of recommendations and institutional roles could be useful to start considering newcomers in the selection process. Although we do not address in detail the *absence of knowledge* situation in this paper, we propose here a rather different approach that uses conceptual clustering of entities' characteristics in order to generate profiles of business entities. In a second step, the profile of the newcomer is compared with the profiles of business entities for which there is some trust information and an estimation of the newcomer trustworthiness is inferred. This approach implies that the characteristics of the business entities are available. We consider this is a reasonable assumption for centralized virtual market places and virtual organizations built upon electronic institutions, and might also be applied to more decentralized approaches by transmission of this kind of entities' knowledge between communicating agents.

Finally, in the third scenario, the selector entity knows that a candidate partner is well reputed in fulfilling agreements in a given role and context (e.g. selling blue cotton zippers to European countries), but it is afraid that the candidate would not be able to provide high quantities of the material in a short period of time. We name this situation the *contextual ignorance*. In this scenario, the evaluator agent knows that the candidate partner is trustworthy in a particular business scenario, or even that is generally trustworthy, but needs to know how well it would adapt to a different type of business. In this paper, we address this question by presenting a description of our *contextual fitness* module. Inclusion of this component is intended to give extra information to the trustier agent by computing a value of *how well the candidate partner fits in the selector current needs*, as defined in the previously issued call for proposals (CFP). The contextual fitness module is described next.

3.2 The Contextual Fitness Algorithm

The Contextual Fitness (CFit) component is used complementary to the estimation of the trustworthiness of a given target agent by the SinAlpha component (this last process is well documented in [18]). When an agent issues a CFP and receives several proposals from candidate partners, the CFit component checks the adequacy of each proposal to the CFP requirements using the following algorithm:

1. For each candidate partner making a proposal, the agent performs conceptual clustering over its contractual past evidences;
2. For each generated class, a stereotype (i.e. a set of the most discriminate characteristics) is extracted;
3. Stereotypes are compared to the current CFP using a similarity analysis approach, and a contextual fitness value is derived;
4. The values computed by the SinAlpha and the Contextual Fitness are aggregated and a global trustworthiness value is derived for the target agent in evaluation.

3.3 Contextual Fitness Current Implementation

Current implementation of the CFit component is a simplification of the algorithm presented above. Therefore, at step 1 we are currently classifying the contractual evidences of each target agent into two different classes: one with all the evidences related to successful contracts, and the other containing evidences related to violated contracts. Then, a stereotype is extracted for each generated class, following the approach proposed in [22]. The purpose of these two first steps is to detect *tendencies* of success and failure for each particular target agent. Figure 2 illustrates the relation between the content of a hypothetical CFP and a stereotype extracted from a hypothetical agent's contractual history. We shall notice that the values of the contractual attributes are quantified to categories, or quantitative values, using a fuzzy approach prior to step 1.

Stereotype: Agent X, null, null, low, false
CFP: chiffon, 1080000, 7

Fig. 2. Possible instances of a CFP (fabric, quantity, delivery time) and a stereotype for target Agent X ($\langle X, At_{fabric}, At_{quant}, At_{dtime}, Res \rangle$)

In the example of Figure 2, we can observe that Agent X has a tendency to violate contracts that stipulate low delivery times, no matter the good and quantity to deliver.

At step 3, the stereotypes extracted for each target are compared to the current CFP and a contextual fitness value is calculated. In the current implementation, the similarity analysis is quite simple, and a non-correlated comparison for each one of the attributes is performed. As a result, a binary value is derived: a zero value for a full match with a *false* stereotype, and a one value for all the remaining cases. In the example above, a delivery time of 7 days is quantified to the *low* category; therefore a match is detected between the stereotype of agent X and current CFP. Consequently, the proposal of the target agent in evaluation is valued zero and is put in a *does_not_fit_the_CFP* set. Finally, at step 4, the evaluating agent chooses the winning proposal from the ones that are not in the *does_not_fit_the_CFP* set, taking into account their trustworthiness value as computed by the SinAlpha component. If there are no proposals in this condition, a proposal is randomly taken from the *does_not_fit_the_CFP* set.

4 Experiments

In order to evaluate the benefits of the Contextual Fitness component (CFit), we run a series of experiments where the standalone SinAlpha component was compared to the global solution constituted by both SinAlpha and CFit components.

4.1 Experimental Testbed and Methodology

All experiments were done using the Repast tool [23]. In the experiments, we simulated a virtual textile marketplace, where client agents post buying leads (in the form of call for proposals) discriminating a fabric to buy together with the correspondent quantity and delivery time, and supplier agents propose, in response to these leads, if they still have the described quantity of the fabric. Table 1 presents the configuration options for the experiments.

Table 1. Configuration of the experiments

Parameter	Value
Fabrics	{Chiffon, Cotton, Voile}
Quantities	{Low, Medium, High}
Delivery Time	{Low, Medium, Big}
# buyers	20
# of sellers	50
Types of sellers	Chosen upon a uniform distribution over the types {"S _{HQT} ", "S _{HDT} ", "S _{HFB} "}
# rounds	100
# runs per experiment	10
SinAlpha parameters	As described in [18]. No bootstrapping is done in current experiments

With these experiments, we wanted to evaluate the capability of CFit in selecting partners taking into account the current business needs, and to understand the benefits that would arise from such an approach. Therefore, we run the same experiment using, first, just the SinAlpha component, and then the global solution of SinAlpha plus CFit. We used the utility criterion to compare both approaches: in each round, the utility of a client agent is 1 if the contract done in this round is successful and 0 if the contract is violated. Therefore, the best approach is the one that gets the higher average utility over all clients, in all rounds; i.e. the one that is more efficient in selecting the best partners for every CFP attributes at any time.

Table 2 presents the types of suppliers used in the experiments. Supplier type estimates its behavior in a probabilistic way. For the experiments described in this section, we defined three different types of suppliers, each one showing some kind of handicap in fulfilling a contract. For example, a S_{HQT} supplier would have a handicap in providing any fabric if the quantity to provide is high.

Table 2. Different types of Suppliers

Supplier Type	Description
S_{HQT}	Probabilistically succeeds 90% of the established contracts, except the ones that involve the delivery of <i>high quantities</i> , which probabilistic fails 75% of the time
S_{HDT}	Probabilistically succeeds 90% of the established contracts, except the ones where the <i>delivery time is low</i> , which probabilistic fails 75% of the time
S_{HFB}	Probabilistically succeeds 90% of the established contracts, except the ones that involves the delivery of <i>a given fabric</i> , which probabilistic fails 75% of the time

4.2 Results

Table 3 shows the results obtained for the experiments described above. As can be observed, the addition of CFit to the SinAlpha aggregation engine improves overall utility in almost 5% for the tested population, which can be considered a relevant result as SinAlpha already outperforms the WMean and Random approaches, as described in section 2.

Table 3. Average utility for each approach in evaluation

SinAlpha Utility	(SinAlpha + CFit) Utility
74,76%	79,51%

Also, when we analyzed the traces of the experiments, we observed that the learning curve for the global solution is consistently more evident than the learning curve we get when we use SinAlpha by itself. In Figures 3 and 4, we plot the relative number of unsuccessful contracts per type of supplier without and with CFit, respectively. No matter what the first choices are concerning the initial selection of partners, the global is successful in stabilizing on selecting all three types of suppliers with low values of violated contracts. On contrary, this behavior is not achieved with SinAlpha when used standalone.

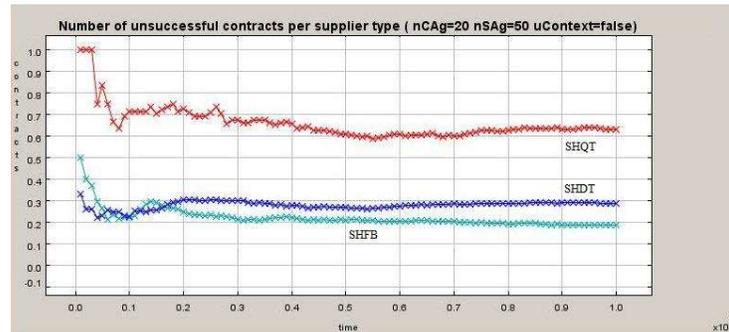


Fig. 3. Relative number of unsuccessful contracts per supplier type for SinAlpha

Finally, we observed that although simple, the algorithm for CFit is indeed able to extract correct stereotypes for each target agent with a few number of past contractual evidences, sometimes as low as five contractual evidences, for the population used. This is an important issue in certain industries, such as the textile industry, where direct or even indirect evaluations of a given target agent might be scarce.

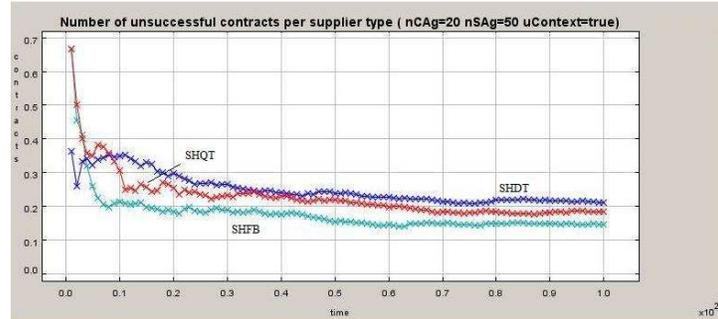


Fig. 4. Relative number of unsuccessful contracts per supplier type for Sinalpha plus CFit approach

4.3 Interpretation of the Results

The results obtained show that an approach that is not able to correctly relate the current business needs to the historical behavior of a supplier is also not able to find out the best characteristics of each supplier. For instances, looking at Figure 3, we observe that suppliers with an handicap on quantity have suffered from a cold start, most probably because they were initially selected to provide high quantities of material. As SinAlpha by itself is not able to capture the handicaps – and as suppliers of this type would tend to succeed on all other contracts they are engaged to, therefore maintaining some level of trustworthiness – the algorithm will continue to select suppliers with quantity handicap to provide high quantities of material. On the other hand, we have shown that approaches that are contextual fitness-aware are able to understand the handicaps of the population (potential partners) and selectively choosing the best suppliers taking into account the current needs.

5 Conclusions and Future Work

In this paper, we introduced a computation trust building model that aggregates historical contractual evidences of business agents in order to compute their trustworthiness. This model leads to a system composed of two different modules: the SinAlpha component is an aggregation curve that incorporates fundamental properties of the dynamics of trust. This component was already implemented and tested in previous work and shows to respond well to specific populations of agents. Then, we presented the Contextual Fitness module, a component that is able to tune the

trustworthiness estimations of any traditional aggregation engine, by relating the past behavior of candidate partners to the specificities of the current business needs. Even though the paper describes a simplified version of the Contextual Fitness algorithm, experimental results show that the inclusion of such an approach improves the process of selection of partners in a relevant way.

In the future, we intend to develop further the global trust system in several ways. In one hand, work will be done on studying the benefits of a sinusoidal like shape in our aggregation engine, following [21] work on the area of Psychology. For this, we have already started a team to work on the acquisition of data/models concerning the behavior of real-world organizations. On the other hand, there is work, already in progress, to extend information representation mechanisms, in order to allow for more complex and appropriate data representations. Finally, we are currently working on the full implementation of the Contextual Fitness module.

Acknowledgements. The first author holds a PhD grant with reference SFRH/BD/39070/2007 from the Portuguese Fundação para a Ciência e a Tecnologia. The authors would like to thank Agnieszka Danek for her help in the evaluation phase of SinAlpha.

References

1. 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)
2. Sabater, J.: Trust and Reputation for Agent Societies. Number 20 in *Monografies de l'institut d'investigacio en inteligencia artificial. IIIA-CSIC*, (2003)
3. Carbo, J., Molina, J., Davila, J.: A BDI Agent Architecture for Reasoning about Reputation. In *IEEE International Conference on Systems, Man, and Cybernetics* (2001)
4. Jøsang, A., Ismail, R.: “The Beta Reputation System”, in *Proceedings of the 15th Bled Electronic Commerce Conference*, Sloven (2002)
5. 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)
6. Zacharia, G., Maes, P.: Trust management through reputation mechanisms. *Applied Artificial Intelligence*, 14(9), 881–908 (2000)
7. Haller, J.: A Bayesian Reputation System for Virtual Organizations. *Dagstuhl Seminar Proceedings 06461, Negotiation and Market Engineering* (2006)
8. 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 - Volume 3*, Estoril, Portugal, May 12 – 16 (2008)
9. Erete, I., Ferguson, E., Sen, S.: Learning task-specific trust decisions. In *Procs. 7th Int. J. Conf. on Autonomous Agents and Multiagent Systems - Vol 3*, Portugal (2008)
10. Smith, M., desJardins, M.: Learning to trust in the Competence and Commitment of Agents. In *Autonomous Agent Multi-Agent Systems*. 18:36-82 (2009)
11. Sabater, J., Paolucci, M.: On Representation and Aggregation of Social Evaluations. In *Computational Trust and Reputation Models, Int. J. Approx. Reasoning* (2007)
12. Rettinger, A., Nickles, M., Tresp, V.: A statistical relational model for trust learning. In *Procs. 7th Int. J. Conf. on Autonomous Agents and Multiagent Systems – Vol. 2* (2008)
13. Elofson, G. Developing Trust with Intelligent Agents: An Exploratory Study, In *Proceedings of the First International Workshop on Trust*, pp. 125-139 (1998)

14. Falcone, R., Castelfranchi, C.: Principles of trust for MAS: cognitive anatomy, social importance, and quantification, Procs. Int. Conference on Multi-Agent Systems (1998)
15. Jonker, C.M., Treur, J.: Formal Analysis of Models for the Dynamics of Trust Based on Experiences. In Procs. of the 9th European Workshop on Modelling Autonomous Agents in A Multi-Agent World: Multiagent System Engineering. F. J. Garijo and M. Boman, Eds. Lecture Notes In Computer Science, vol. 1647. Springer-Verlag, London, 221-231 (1999)
16. Marsh, S.P.: Formalising Trust as a Computational Concept, PhD Thesis, University of Stirling (1994)
17. Melaye, D., Demazeau, Y.: Bayesian Dynamic Trust Model. CEEMAS 2005: 480-489 (2005)
18. Urbano J, Rocha AP, Oliveira E (2009). Computing Confidence Values: Do Trust Dynamics Matter? The 14th Portuguese Conference on Artificial Intelligence, EPIA'2009, Aveiro, Portugal, October 2009
19. 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, Vol. 13, N. 2, September 2006, pp. 119–15 (2006)
20. Sabater, J., Paolucci, M., Conte, R. RePAGE: Reputation and image among limited autonomous partners. Journal of artificial societies and social simulation, Vol. 9, 3 (2006)
21. Straker, D. Changing Minds: in Detail. Syque Press (2008)
22. Paliouras, G., Karkaletisis, V., Papatheodorou, C. , Spyropoulos, C.D. Exploiting learning techniques for the acquisition of user stereotypes and communities. In UM'99: Proceedings of the seventh international conference on User modeling, pages 169-178, USA (1999)
23. North, M.J., N.T. Collier, and J.R. Vos, Experiences Creating Three Implementations of the Repast Agent Modeling Toolkit, ACM Transactions on Modeling and Computer Simulation, Vol. 16, Issue 1, pp. 1-25, ACM, New York, New York, USA (January 2006).