Extracting Trustworthiness Tendencies Using the Frequency Increase Metric

Joana Urbano, Ana Paula Rocha, and Eugénio Oliveira

LIACC – Laboratory for Artificial Intelligence and Computer Science
Faculdade de Engenharia da Universidade do Porto – DEI
Rua Dr. Roberto Frias, 4200-465, Porto, Portugal
{joana.urbano,arocha,eco}@fe.up.pt

Abstract. Computational trust systems are currently considered enabler tools for the automation and the general acceptance of global electronic business-to-business processes, such as the sourcing and the selection of business partners outside the sphere of relationships of the selector. However, most of the existing trust models use simple statistical techniques to aggregate trust evidences into trustworthiness scores, and do not take context into consideration. In this paper we propose a situation-aware trust model composed of two components: Sinalpha, an aggregator engine that embeds properties of the dynamics of trust; and *CF*, a technique that extracts failure tendencies of agents from the history of their past events, complementing the value derived from Sinalpha with contextual information. We experimentally compared our trust model with and without the *CF* technique. The results obtained allow us to conclude that the consideration of context is of vital importance in order to perform more accurate selection decisions.

Keywords: Situation-aware Trust; Dynamics of Trust; Multi-agent Systems.

1 Introduction

Several technologies are being studied and applied in the general process of computerized supply chain management. Computational trust management is one such technology that will allow extending electronic sourcing to world-wide located, non registered and probably unknown business partners. With this technology, a business entity will be able to search the suppliers offer space and to filter the ones that are fitted to the entity current needs, in a scale of the size of the Internet.

The first generation of Computational Trust and Reputation (CTR) systems addressed the representation and the aggregation of trust evidences about a given agent in evaluation into a trustworthiness score, and most of these proposals are based on some sort of statistical aggregation methods (e.g. [1], [2], [3], [4], [5], [6], [7]). Other works proposed more sophisticated engines that consider the dynamics of trust in the computation of confidence scores, in theoretical and practical terms (e.g. [8], [9], [10], [11]). However, none of the current computational trust approaches is mature enough to be itself trusted by real managers.

J. Filipe and J. Cordeiro (Eds.): ICEIS 2010, LNBIP 73, pp. 208–221, 2011.
© Springer-Verlag Berlin Heidelberg 2011

Trying to cope with this question, trust community is moving towards a second generation of models that explore the situation of the trust assessment in order to improve its credibility. However, few proposals have been made on this specific area (see [12], [13], [14], [15]).

In this paper, we propose a situation-aware technique that allows the extraction of tendencies of agents' behavior[1]. This technique allows, for instance, to detect whether a given supplier has a tendency to fail or to succeed contracts that are *similar* to the current business need (e.g. in terms of good, quantity and delivery time conditions). We performed experiments that show that this technique enhances traditional CTR systems by bringing context into the loop; i.e. it not only concerns if a given supplier is generally trusted good or bad, but if it is trusted good or bad *in the specific contractual situation*. Also, our approach differs from the situation-aware proposals mentioned above in the way that it does not imply the use of hierarchical-based structures (e.g. ontology) and that it is able to detect fine-grain subtle dissimilarities in *related* situations. Moreover, it was designed in order to allow the effective estimation of trustworthiness values even when the trust evidences about the agent in evaluation are scarce.

The situation-aware technique we propose can be used in conjuction with any existing trust aggregation engine. In our experiments, we use Sinalpha, a sigmoid-like aggregator that we have developed ([17]) that distinguishes from traditional trust proposals by embedding properties of the dynamics of trust. In this paper, we review the fundamental characteristics of Sinalpha and present our conclusions about the relevance of the inclusion of such properties in trust aggregation engines.

Although we contextualized the use of our trust system in the sourcing/ procurement part of the supply chain, agent-based trust and reputation systems are of general interest in many other domains (for instance, general business, psychology, social simulation, system resources' management, etc), and apply to all social and business areas of the society where trust is deemed of vital importance.

The remaining of this paper is structured as follows: Section 2 describes our study about the relevance of considering properties of the dynamics of trust in the aggregation engine of CTR systems. Section 3 describes the technique we developed in order to complement traditional CTR engines with situation-aware functionality. Section 4 presents the experiments we run in order to evaluate the proposed situation-aware technique, and Section 5 concludes the paper.

2 Using Trust Dynamics in the Aggregation Engine

2.1 The Sinalpha Model

In [17], we described Sinalpha, a trust aggregation engine based on an S-like curve (Figure 1) that allows for an expressive representation of the following properties of the dynamics of trust:

---

[1] This paper in an extended version of the paper presented at ICEIS 2010, the 12th International Conference on Enterprise Information Systems [16].

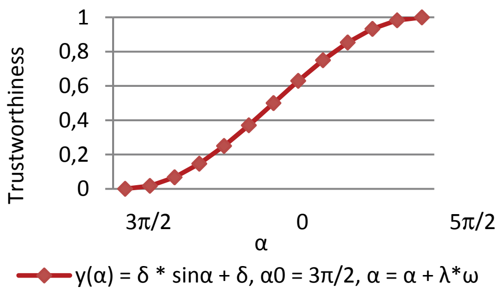

Fig. 1. The S-like curve

- *Asymmetry*: stipulates that trust is hard to gain and easy to lose;
- *Maturity*: 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*: distinguishes between possible different patterns of past behavior.

The choice of the S-like shape was based on the concept of *the hysteresis of trust and betrayal*, from Straker [18]. 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 S 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). It is also worth to mention that the final Sinalpha mathematical formula, depicted at Figure 1 and explained in detail in [17], was based on Lapshin's work [19].

Therefore, the aggregation of the trust evidences of the agent in evaluation using the Sinalpha model presents the following characteristics: the trustworthiness value of the agent grows slowly in the presence of evidences with positive outcomes when the agent is not yet trustable, accelerates when the agent acquires some degree of trustworthiness, and slowly decays when the agent is considered trustable (i.e., in the top right third of the curve). This allows for the definition of three different trust maturity phases, and constitutes the *maturity* property. The decrease movement upon evidences with negative outcomes follows the same logic, although the mathematical formula subjacent to the curve (cf. Figure 1) includes parameter λ, which takes different values dependently on the outcome of the evidence is positive $(\lambda = +1.0)$ or tive $(\lambda = -1.5)$, permitting for slower growths and faster decays. This constitutes the *asymmetry* property. Finally, the mathematics of Sinalpha's formula implies that the aggregation of the same number of positive and negative outcomes that are presented in different ordering results in different trustworthiness values. This constitutes the *distinguishable* property. Considering this last property, we have a somewhat different view than the one presented in [7], where the authors state that the aggregation of evaluations shall not depend on the order in which these evaluations are aggregated. The results we obtained in our experiments and discussed in the next section seem to support our conviction.

One can argue that we could use other S-like curves instead of a sin-based one, such as the Sigmoid curve. However, we intuitively feel that a Sigmoid curve permits a probably too soft penalization of partners that proved to be trustable but that failed the last *n* contracts. This can happens 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 behavior of deceptive provider agents, who tend to build up a trustworthy image using simple contracts and then violate bigger contracts exploring the acquired trustworthiness ([20]).

2.2 Relevance of Sinalpha's Trust Properties

In this section, we review the main results obtained when we experimentally compared Sinalpha to a weighted mean by recency approach ([6]) that represents traditional statistical approaches. A description of the experiments performed is available at [17]. A more exhaustive comparison that we have conducted between trust models based on weighted means and heuristics-based models embedding properties of the dynamics of trust is presented in [21].

The first result we obtained show that Sinalpha is more effective in selecting good partners (i.e. agents that have a higher probability of fulfilling a contract) and in avoiding bad partners (agents that have a higher probability of violating a contract) than the weighted means approach. This is due to the *maturity* property of Sinalpha, where the entire historical path of the agent in evaluation is taken into account in the process of trust construction, meaning that agents have to accumulate several good experiences in the past until they are able to get an average to high trustworthiness score. In opposition, the weighted means approach allows the selection of agents with fewer past evidences, which can promote uninformed decisions.

Another difference between the two approaches is related to the *asymmetry* property of Sinalpha, which allows it to more effectively identifying and acting upon partners that show the much undesired intermittent behavior, by giving more weight to violations' penalty than to fulfillments' rewards.

A related situation occurs when agents show bursty-like intermittent behavior, i.e. when they present long sequences of positive outcomes and then long sequences of negative outcomes and so on. The results have shown that, in this situation, the weighted means approach can *forgive* past long sequences of negative outcomes if these sequences are followed by a long period of inactivity and a single recent positive outcome. One could argue 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 ([21]). Finally, the forgiveness question does not apply to Sinalpha, due to the action of the *maturity* property, although we realized that it has a somewhat bigger tendency than the weighted means approach to enter a burst of deceptive behavior and also that it can be slower in penalizing good partners immediately after they invert their behavior.

In these experiments, we could not evaluate, however, the potential full benefit of using the Sinalpha round shape at the extremities against simpler curves that do show similar trust dynamics properties (e.g. curves with linear shape). Indeed, in [21] we compared the performance of Sinapha against a simpler curve that uses the same λ and ω parameters (cf. Figure 1) but that lacks the softness round curve at the *creating trust* and *taking advantage* phases. The results of these experiments show similar performance of both curves in the tested scenarios. Therefore, we conclude that we need more complex models of the 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 Our Situation-Aware Trust Technique

3.1 Motivation for Situational Trust

Computational trust estimations help an agent to predict how well a given target agent will execute a task and thus 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 trustworthiness 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 us 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 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 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.

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 no trust 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*. The works in [6] and [15] suggest 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 this situation in this paper, we consider using conceptual clustering of entities' characteristics in the future in order to generate profiles of business entities. Thus, 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 a minimum of the business entities' characteristics is available, which constitutes a reasonable assumption for virtual organizations built upon electronic institutions and even for decentralized approaches where agents are able to present certificate-like information.

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, or even that it is generally trustworthy, but needs to know how well it would adapt to a (even slightly) different business context. For example, the evaluator knows that a given Asiatic supplier is a good seller of cotton zippers but it is afraid that it could fail in providing high quantities of this material in a short period of time, because it does not use to transact with the supplier in these specific conditions. We name this situation the *contextual ignorance*.

In the next sections, we present a situation-aware trust technique that is able to extract tendencies of the behavior of agents in a contextualized way. This technique supports the evaluator agent in addressing the *contextual ignorance* question and can be used with any traditional trust aggregation model, such as the weighted means and the Sinalpha models referred in Section 2 and. Next, we present the scenario and notation used throughout the section.

3.2 Scenario and Notation

Our scenario consists of a social simulation where, at every simulation round, trading client agents attempt to place orders of some type of textile fabrics to the best available supplier agents, taking into account the estimated trustworthiness of the suppliers in the context of the specific orders. At the end of the round, a contract is established between the client agent and the selected supplier agent.

Formally, $ac \in AC$ is a client agent from the set $AC = \{ac_1, ac_2, \dots, ac_n\}$ of all n clients considered in the simulation, and $as \in AS$ is a supplier agent from the set $AS = \{as_1, as_2, \dots, as_m\}$ of all m suppliers considered in the simulation. Also, an order instantiates a business need that comprises the *contractual context* of the future transaction. For the sake of simplicity, we only consider here three contractual terms that are the elements of the set $AT = \{fabric, quantity, deliverytime\}$. Regarding domain values of the contractual terms, the set $FAB = \{cotton, chiffon, voile\}$ contains the possible values for the term $fabric$; the set $QT = \{low, medium, high\}$ contains the possible values for the term $quantity$; and the set $DT = \{short, medium, big\}$ contains the possible values for the term $deliverytime$. Therefore, the contractual context cc is an ordered triple belonging to the 3-ary Cartesian product $CC = FAB \times QT \times DT$. An example of a contractual context associated to an order is $(cotton, high, short)$.

At every simulation round, each one of the n clients broadcasts its current need specifying the corresponding contractual context cc. In response, all suppliers that still have stock on the desired fabric issue a proposal. In this scenario, a *proposal* is simply the identification of the supplier that has proposed. I.e., instead of using selection parameters such as the ones associated with price or payment conditions, the suppliers would be selected uniquely by their estimated trustworthiness.

Then, after client agent ac selects supplier agent as to provide the good as mentioned in the contractual context cc, both agents establish a contract specifying ac, as and cc. When the transaction is completed, with either the supplier agent succeeding or failing to provide the good in the established conditions, a *contractual evidence* is generated accordingly. In our scenario, an evidence evd is an ordered tuple from the 4-ary Cartesian product that define the set $Evd = AC \times AS \times CC \times O$ of all evidences generated until evaluation time *t*, where $O = \{true, false\}$ is the set of all possible outcomes for the contract (i.e. it takes the value $true$ when the contract is succeeded and $false$ otherwise). In the same way, $Evd_{as_i} \subseteq Evd$ is the subset of all evidences where as_i appears as the supplier counterparty. This means that $Evd_{as_i} = \{(ac, as_i, cc, o) \in Evd: ac \in AC, cc \in CC, o \in O\}$ is the *contractual history* of agent as_i at the evaluation time. If supplier as_i had never transacted before, then $Evd_{as_i} = \emptyset$.

3.3 Our Situation-Aware Trust Model

Figure 2 illustrates the selection algorithm used by our clients at every simulation round in order to select the proposal that best fits their needs.

```
1:  function SELECTION (cc, P) returns a proposal
2:    inputs: cc, contractual context of current order
3:            P: a list with all proposals received by the client
4:
5:    bestProposal ← a proposal chosen randomly from P
6:    higherTrust ← SINALPHA (Evd_bestProposal)
7:    cf ← CONTEXTUAL − FITNESS (cc, Evd_bestProposal)
8:    if cf is 0 then higherTrust ← 0
9:    else higherTrust ← higherTrust * cf
10:   for each remaining p in P do
11:       currentProposal ← p
12:       currentTrust ← SINALPHA (Evd_currentProposal)
13:       cf ← CONTEXTUAL − FITNESS (cc, Evd_currentProposal)
14:       if cf is 0 then currentTrust ← 0
15:       else currentTrust ← currentTrust * cf
16:       if currentTrust is higher than higherTrust then
17:       bestProposal ← currentProposal
18:       higherTrust ← currentTrust
19:    return bestProposal
```

Fig. 2. The SELECTION algorithm

The client picks up randomly a proposal from the set of received proposals and calculates the trustworthiness value of the proponent (lines 5 to 9). Then, for each one of the remaining proposals, it estimates its trustworthiness (lines 10 to 15) and updates the best proposal (lines 16 to 18). The algorithm returns the proposal from the most estimated trusted supplier for the specific context cc (line 19).

As can be observed from Figure 2, our trust model is composed of two components, as described earlier in this document. Indeed, the first component is $sinalpha\ (as) \in [0, 1]$, a function that aggregates all outcomes from the available evidences using the formula illustrated in Figure 1. The second component is our situation-aware technique, $cf\ (as, cc) \in \{0, 1\}$, which is a binary operator that measures the *adequacy* of the proponent supplier to the contractual context of the current need of the client. Therefore, the trustworthiness value of supplier as as estimated by client ac is given in Equation 1.

$$trust_{ac}\ (as, cc) = sinalpha\ (as) * cf\ (as, cc) \tag{1}$$

This is the same as to say that, in a given moment, an agent may be qualified as trustworthy in some situation and as untrustworthy in a (maybe slightly) different situation. Next, we describe our binary operator $cf\ (as, cc)$ in more detail.

3.4 The CF Operator

We developed the CF operator, a technique that allows non-situational aggregation engines (such as Sinalpha or the ones based on weighted means) to make inference about the trustworthiness of an agent based on context. It is an online, incremental technique that does not rely on predefined measures of similarity or ontology-based inference as the ones presented in [12], [13], [14] and [15]. The algorithm of CF is presented in Figure 3.

1: $\boldsymbol{function}\ CONTEXTUAL - FITNESS\ (cc, Evd_{as})\ \boldsymbol{returns}\ a\ value\ in\ \{0, 1\}$
2: $\boldsymbol{inputs}: cc, contractual\ context\ of\ current\ order$
3: $\qquad Evd_{as}: the\ contractual\ history\ of\ agent\ as$
4:
5: $EvdF_{as} \leftarrow all\ evidences\ from\ agent\ as\ that\ have\ a\ false\ outcome$
6: $cc_{false} \leftarrow EXTRACT - FALSE - TENDENCY\ (EvdF_{as})$
7: $\boldsymbol{if}\ there\ is\ a\ match\ between\ cc\ and\ cc_{false}\ \boldsymbol{then\ return}\ 0$
8: $\boldsymbol{else\ return}\ 1$

Fig. 3. The CONTEXTUAL-FITNESS algorithm

From the figure above, we verify that the algorithm starts by putting all evidences of agent as that have outcome $false$ in a separate class (line 5). This negative outcome can represent, for example, the past transactions of the supplier agent that triggered relevant contractual sanctions, although the meaning of such outcome can be established by each individual client agent. Then, at line 6, a *behavior tendency* is extracted for this class using the *Frequency Increase* metric [22], presented in the following Equation:

$$\propto = \left(\frac{\#InstAttClass}{\#InstClass}\right)^2 - \left(\frac{\#InstAttTotal}{\#InstTotal}\right)^2. \tag{2}$$

In the equation above, $\#InstAttClass$ is the number of times that a given attribute appears in the class, $\#InstClass$ is the total number of evidences in the class, $\#InstAttTotal$ is the number of times that the attribute appears in all classes, and $\#InstTotal$ is the total number of evidences kept for the agent in evaluation. Therefore, by applying the Frequency Metric to the set $EvdF_{as}$ of false evidences of agent as, the algorithm tests, one by one, which contractual term-value pair can be considered relevant. The parameter α in the equation above is the degree of the required extent of frequency increase, and determines the granularity of tendency extraction. At the end of the procedure, the *most significant contractual characteristics* of the class of false evidences are extracted (line 6). It is worth to mention that, depending on the degree of the required extent of frequency increase and on the evidence set of the agent in evaluation, it is possible that the algorithm does not return any tendency.

Finally, at lines 7 and 8, the extracted false tendency is compared to the contractual context of current order of the client. If there is a match, it means that the supplier has a tendency to fail this type of contracts, and therefore the cf value (and the global trustworthiness value, cf. Equation 1) is zero. Otherwise, there is no evident signal that the supplier is inapt to perform the current transaction, and its final trustworthiness score is given by the Sinalpha output.

Figure 4 illustrates an example of a match between the contractual context of a current order and the failure tendency of the agent in assessment. In there, client ac wants to purchase high quantities (1800000 meters) of chiffon in short delivery times (seven days), as given by cc_{order}. At the same time, supplier as has a failure tendency in delivering in short delivery times, independently of the fabric and the quantity considered, as given by cc_{false}. Therefore, the final trustworthiness value of the supplier for this specific order is zero, as given by Equation 1, which strongly reduces the chance of the supplier being selected by ac to the current business transaction.

$$\boldsymbol{cc_{order}}: chiffon, 1080000, 7$$

$$\boldsymbol{cc_{false}}: *, *, short$$

Fig. 4. An example of a match between the contractual context of an order and the false tendency behavior of an agent

The process just described of extracting tendencies of negative behavior in order to prevent unfitted selection decisions is a dynamic incremental process that shall be repeated at every trustworthiness assessment. The benefits of such an approach are two-folded: first, it allows the extraction of tendencies even when there are few trust evidences about the agent in evaluation; then, it allows capturing the variability of the behavior of the agent at any time, which is a desired requirement for trust models designed to operate in real-world environments.

We draw here a final remark about our trust technique: as described in this section, CF is only dealing with the negative class of the evidences of the agent in evaluation. However, we believe that the use of the positive class and the use of distinct degrees of fitness could allow refining our algorithm, and this constitutes ongoing work.

4 Experiments

In order to evaluate the benefits of the proposed situation-aware trust model, we run a series of experiments using the scenario described in Section 3.2. All experiments were run using the Repast tool (http://repast.sourceforge.net).

4.1 Experimental Testbed and Methodology

The main parameters of the experiments are presented in Table 1.

Table 1. Configuration of the experiments

Parameter	Value
Fabrics	{Chiffon, Cotton, Voile}
Quantities	{Low, Medium, High}
Delivery Time	{Short, Medium, Big}
# buyers	20
# of sellers	50
Types of sellers	Uniform distribution over the types considered in population
Seller stock	Up to 4 contracts per round
# rounds / # runs	60 / 20
CF α threshold	0.25

In these experiments, we wanted to evaluate if the situation-aware technique would improve the ability of the trust system in selecting partners by taking into account the current business context. Therefore, we run the same experiments using, first, just the Sinalpha component (the *SA* model), and then the enhanced version of Sinalpha that uses our proposed situation-aware technique (the *CF* model).

Therefore, at every simulation round, each client ac announces a business need in the form of a contractual context cc. Every supplier that still has stock on the specified fabric manifests its intention in providing the material in the conditions specified by cc. Then, client ac selects the best proposal using as a criterion the estimated trustworthiness of each proponent. I.e. it either uses *SA* or *CF* to perform the trustworthiness estimations, depending on the trust approach in evaluation. Finally, an outcome is generated for the transaction between ac and the selected supplier, based on the *type* of this supplier.

Table 2. Characterization of the different types and populations of suppliers

Type	Description	Prob. Success	Pop.A	Pop.B	Pop.C
S_{HQT}	Handicap on *high quantities*	0.05 (handicap), 0.95 (otherwise)	X		X
S_{HDT}	Handicap on *short delivery time*		X		X
S_{HFB}	Handicap on *specific fabric*		X		X
S_{HFBQT}	H. on *fabric* and *high quantity*				X
S_{HFBDT}	H. on *fabric* and *short del. time*				X
S_{HQTDT}	H. on *high quant.* and *short d. time*				X
Good		95%			X
Fair		80%			X
Bad		50%			X
I-S_{HQT}	Same as S_{HQT}, S_{HDT} and S_{HFB}, with a probability of 66.7% of changing handicap at round 30			X	
I-S_{HDT}				X	
I-S_{HFB}				X	

Populations of Suppliers. We used different types of suppliers and populations in the experiments. Each type of supplier reflects its ability in fulfilling or violating a contract, as described in Table 2.

Evaluation Metrics. We used two different metrics to evaluate the trust models. The first one is the *average utility* of clients, which is the ratio of the number of contracts that were fulfilled by the suppliers over the total number of contracts, averaged over all clients and all rounds. The second metric is the *average utility per round*, which measures the same ratio at a round basis. Intuitively, the greater the ability of the trust model in distinguishing between the dynamic behaviors of the proponent suppliers, the better decisions it will make and, consequently, the higher utility it will get.

4.2 Results

Figure 5 shows the results in terms of average utility for populations A, B and C.

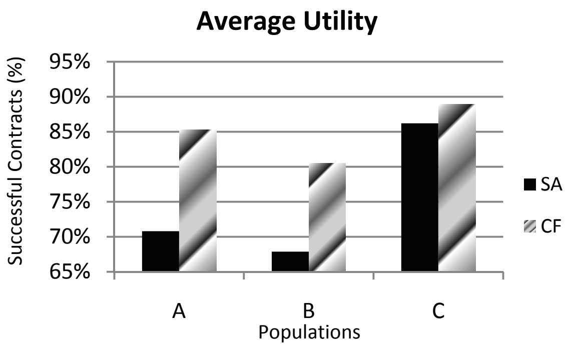

Fig. 5. Average utility for populations A, B and C

In population A, *CF* clearly outperforms *SA*, with an average utility of 85.27% (standard deviation 1.68%) against the average utility of 70.80% (sd. 5.10%) of *SA*. In population B, where agents have two thirds of probability of inverting their behavior at the middle round, *CF* still outperforms SA with 80.54% of average utility (sd. 2.36%) against the 67.87% (4.98%) numbers of *SA*. Finally, at population C, where there is a mix of several different types of suppliers, *SA* raises substantially its performance, achieving 86.22% of average utility (sd. 1.85%), although it still underperforms the *CF* approach that achieves 88.92% of utility (sd. 1.43%).

Another view of the results is given in Figure 6, where the results of the utility are given in a per round basis.

Analyzing the figure above, we verify that in population A (*dotted lines*), the *SA* approach oscillates around an average number of 14 (in 20) successful contracts per round. On the opposite, the *CF* approach shows an ability of *learning* the behavior of the suppliers, continually increasing the utility since the first rounds, where the evidences available about every supplier are still scarce.

Concerning population B (*dashed lines*), we verify that *CF* has a much steeper fall in terms of utility than the *SA* approach at round 30, when the behavior of most of the suppliers is inverted. However, we also verify that the utility of *CF* is always bigger than the utility of *SA* for this population, and that the *CF* starts recovering soon after round 30, showing its ability in dynamically updating the tendencies of behavior in the presence of few new evidences.

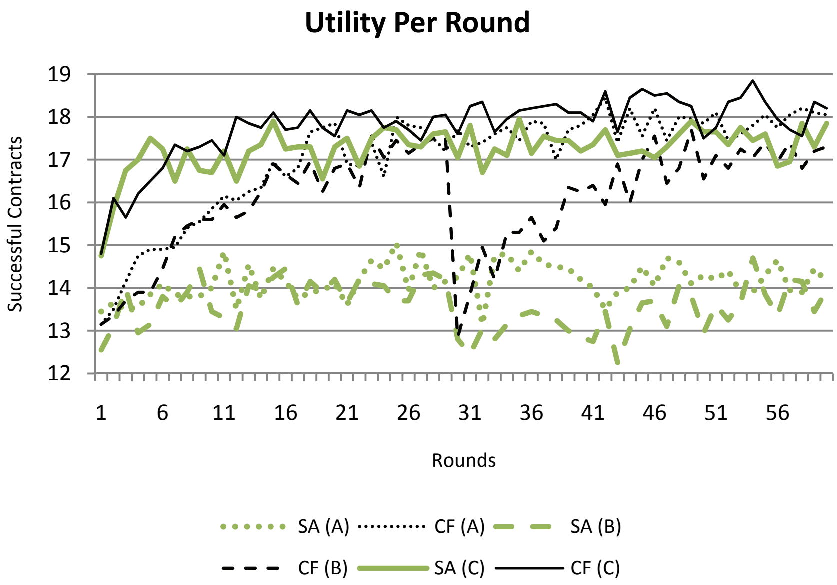

Fig. 6. Utility per round for populations A, B and C

Finally, in population C (*continuous lines*), we verify that both *SA* and *CF* start learning the behavior of the suppliers much earlier than in the previous experiments with populations A and B. In reality, this happens because population C presents a more diversified population, with 4/9 of the population being generally good (with suppliers of types Good, H_{FBQT}, H_{FBDT} and H_{QTDT}). What should be retained from these results is that *CF* outperforms *SA* even in the presence of such good suppliers.

4.3 Discussion

From the results obtained in the experiments, we verify that the *SA* approach, representing the traditional non situation-aware trust models, is effective in differentiating between good and bad suppliers, i.e., suppliers that have a high and a low probability in fulfilling a contract, respectively. However, in real and open business environments, it is expected that the existing population of suppliers does not show this binary behavior, but that instead is composed of suppliers that are more fitted to specific contractual characteristics and that can failed in other types of contracts. In this reality, populations A and B have shown that a situation-less trust model is not effective in differentiating between different suppliers' characteristics, with the clients keeping selecting the same suppliers over and over again, occasionally failing the contracts for which the suppliers have an handicap.

On the contrary, we verify that our *CF* trust model is able to distinguish between suppliers that present different tendencies of failure, and that this ability allows *CF* clients to make better decisions, reflected in the higher utility they achieved in all the experiments with populations A, B and C. Moreover, we observed that *CF* is fast in detecting changes in the behavior of suppliers, as it dynamically updates the failure tendencies extracted when it has new evidences. In the same way, it is able to extract tendencies even when the number of trust evidences about the supplier in evaluation is scarce.

5 Conclusions

In this paper, we presented *CF*, a simple situation-aware technique that extracts failure tendencies from the history of past evidences of an agent, based on the Frequency Increase metric. This technique can be used with any traditional trust system in order to enhance the estimation of trustworthiness scores.

Although other situation-aware approaches are now being proposed in the trust management field, we believe that our proposed *CF* technique presents some benefits over them. First, it can be used with any of the existing traditional trust aggregation engines. Secondly, it is an online process, meaning that it captures the variability in the agents' behavior in a dynamic way. Finally, it does not rely on ontology-based situation representations, and therefore the analysis of the similarity between the situation in assessment and the past evidences of the agent in evaluation does not require specific, domain-based similarity functions. Also, it allows for fine-grain *dissimilarity* detection (e.g. it distinguishes between the *similar though different* situations of providing one container of cotton in 7 or in 14 days) that can be hard to express using pre-defined distance functions.

We evaluated the *CF* technique using Sinalpha, a traditional aggregation engine approach enhanced by the inclusion of properties of the dynamics of trust. We overviewed the benefits of embedding the properties of asymmetry, maturity and distinguishably in trust aggregation engines. However, we also concluded that the study of the benefits of Sinalpha' sinusoidal shape that follows work on the area of Psychology needs proper data/models concerning the behavior of real-world organizations, and we will address the acquisition of such data sets in future work.

Acknowledgements. This research is funded by FCT (Fundação para a Ciência e a Tecnologia) project PTDC/EIA-EIA/104420/2008. The first author enjoys a PhD grant with reference SFRH/BD/39070/2007 from FCT.

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. Jøsang, A., Ismail, R.: The Beta Reputation System. In: Proceedings of the 15th Bled Electronic Commerce Conference, Sloven (2002)
3. Zacharia, G., Maes, P.: Trust management through reputation mechanisms. Applied Artificial Intelligence 14(9), 881–908 (2000)
4. 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 (2008)
5. Sabater, J.: Trust and Reputation for Agent Societies. Number 20 in Monografies de l'institut d'investigacio en intelligència artificial. IIIA-CSIC (2003)

6. 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), 119–115 (2006)
7. Sabater, J., Paolucci, M., Conte, R.: Repage: Reputation and image among limited autonomous partners. Journal of Artificial Societies and Social Simulation 9, 3 (2006)
8. Castelfranchi, C., Falcone, R.: Principles of trust for MAS: cognitive anatomy, social importance, and quantification. In: Procs. Int. Conference on Multi-Agent Systems (1998)
9. 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)
10. Marsh, S., Briggs, P.: Examining Trust, Forgiveness and Regret as Computational Concepts. In: Golbeck, J. (ed.) Computing with Social Trust, pp. 9–43. Springer, Heidelberg (2008)
11. 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)
12. Tavakolifard, M.: Situation-aware trust management. In: Proceedings of the Third ACM Conference on Recommender Systems, pp. 413–416 (2009)
13. Neisse, R., Wegdam, M., Sinderen, M., Lenzini, G.: Trust management model and architecture for context-aware service platforms. In: Meersman, R., Dillon, T., Herrero, P. (eds.) OTM 2009. LNCS, vol. 5871, pp. 1803–1820. Springer, Heidelberg (2009)
14. Rehak, M., Gregor, M., Pechoucek, M.: Multidimensional context representations for situational trust. In: IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications, pp. 315–320 (2006)
15. Hermoso, R., Billhardt, H., Ossowski, S.: Dynamic evolution of role taxonomies through multidimensional clustering in multiagent organizations. In: Yang, J.-J., Yokoo, M., Ito, T., Jin, Z., Scerri, P. (eds.) PRIMA 2009. LNCS, vol. 5925, pp. 587–594. Springer, Heidelberg (2009)
16. Urbano, J., Rocha, A.P., Oliveira, E.: Refining the Trustworthiness Assessment of Suppliers through Extraction of Stereotypes. In: Filipe, J., Cordeiro, J. (eds.) ICEIS 2010. LNBIP, vol. 73, pp. 206–219. Springer, Heidelberg (2010)
17. Urbano, J., Rocha, A.P., Oliveira, E.: Computing Confidence Values: Does Trust Dynamics Matter? In: Lopes, L.S., Lau, N., Mariano, P., Rocha, L.M. (eds.) EPIA 2009. LNCS (LNAI), vol. 5816, pp. 520–531. Springer, Heidelberg (2009)
18. Straker, D.: Changing Minds: in Detail. Syque Press (2008)
19. Lapshin, R.V.: Analytical model for the approximation of hysteresis loop and its application to the scanning tunneling microscope. Review of Scientific Instruments 66(9), 4718–4730 (1995)
20. Schlosser, A., Voss, M.: Simulating data dissemination techniques for local reputation systems. In: Procs. of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems, pp. 1173–1174 (2005)
21. Danek, A., Urbano, J., Rocha, A.P., Oliveira, E.: Engaging the Dynamics of Trust in Computational Trust and Reputation Systems. In: Jędrzejowicz, P., Nguyen, N.T., Howlet, R.J., Jain, L.C. (eds.) KES-AMSTA 2010. LNCS, vol. 6070, pp. 22–31. Springer, Heidelberg (2010)
22. Paliouras, G., Karkaletsis, V., Papatheodorou, C., Pyropoulos, C.D.: Exploiting Learning Techniques for the Acquisition of User Stereotypes and Communities. In: Procs. of UM 1999 (1999)