

# Making Electronic Contracting Operational and Trustworthy

Joana Urbano, Henrique Lopes Cardoso, Eugénio Oliveira

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

**Abstract.** An Electronic Institution includes a normative environment with rules and norms for agents' interoperability, and is also a service providing platform that assists agents in the task of establishing and conducting normative relationships (contracts). Using this platform, agents representing organizations willing to engage in a collective contractual activity select partners according to different factors, including their capabilities, current business needs and information on past business experiences that may be used as inputs to trust building. In our framework we have designed a tightly coupled connection between electronic contract monitoring and a computational trust model. In this paper, we explain the rationale behind this connection and detail how it is materialized. In particular, we explain how our situation-aware trust model relies on past contractual behavior to dynamically build up a trustworthiness image of each agent that can be helpful for future encounters. Experiments with simplified scenarios show the effectiveness of our approach.

## 1 Introduction


An Electronic Institution (EI) [1] is a software platform including a core infrastructure – a *normative environment* – which embraces the norms that apply to contracts as established among agents. A central role of an EI is to make this environment operational, in the sense that contractual norms will be monitored for compliance, and consequently norm violations will be reacted upon. Furthermore, an EI provides a set of services that assist agents in the task of establishing and conducting normative relationships. Therefore, not only are we interested in monitoring the compliance of agents with the norms they voluntarily adhere to through contracting, but also in providing computational tools that help on automating the creation of such contracts.

The services that we include in an EI are of utmost importance for secure and reliable agent interoperability and cover a broad range of MAS research issues:

- Automatic negotiation [2]: automates partner selection upon a business opportunity and is based on negotiation protocols exhibiting properties such as information privacy, qualitative feedback and adaptation;
- Contract monitoring and enforcement [3]: monitors parties' compliance with

- contractual terms and applies specified sanctions in violations occur;
- Computational trust [4]: aggregates trust information from past contractual behavior, which allows agents to make informed decisions regarding the selection of partners and/or the negotiation of contracts.

When designing an integrated approach that includes these services (Figure 1), we must also consider the interconnections among them. In order to move from the negotiation process to contract monitoring, contract drafting must be taken into account. We may instantiate contract templates with the outcome of negotiation, or include in the negotiation process itself the specification of specific contract clauses. Contracts resulting from successful negotiations may then be validated, registered and digitally signed, before being handed to a normative environment for monitoring and enforcement purposes. Finally, the way agents enact their contracts provides important information for trust building. A repository of trust information may then complete the circle by providing relevant inputs for future negotiations. The integration of all these stages (Figure 1) has been addressed through the development of an EI Platform for B2B Contracting.



**Fig. 1.** Electronic Institution services for B2B contracting.

A specific and most relevant case of joint activity that we have been looking at is the Virtual Enterprise (VE), seen as a temporary organization composed of autonomous enterprises, formed to address a specific business opportunity. One requirement for the successful creation of a VE is a fast set-up phase, which in an open environment may lead to the need to establish contracts with new partners whose past contractual performance is not known. Therefore, the regulation of the operation stage of the VE through an electronic contract is important, as is the aggregation of trust information that can be used to make more informed decisions concerning the selection of future business partners.

An important role of a normative environment in an EI is, thus, to record the way contracts are enacted, making it possible to build up models of agents' trustworthiness that are based on their past performance. This information can then be used as an important source for possible future encounters with other agents. In this paper we will focus on the relationship between contract monitoring and the construction of trust indicators. The main research question we are here dealing with can be summarized as follows: is it possible to benefit from specific information on an

agent's contractual behavior to better tune its trustworthiness and, as a consequence, influence the way partners are selected in future contracts? Our hypothesis is that a better characterization of agents' trustworthiness brings a better ability to select partners and, therefore, results in more successful contracts.

In Section 2 we describe the normative environment and its interfaces to other EI components. Section 3 describes how a trust model can be enhanced using detailed contractual information in order to permit context-aware trust assessment. Section 4 evaluates our situation-aware trust model, and Section 5 concludes the paper.

## 2 A Normative Environment for Contract Monitoring

The role of a normative environment is, besides providing a set of regulations under which agents' collective work is made possible, to check whether agents are willing to follow the norms they commit to (monitoring), and further to employ correction measures as a means of coercing agents to comply (enforcement). We represent in the normative environment's structure the normative relations that correspond to contracts established by real world entities [1]. The "shape" of the environment will therefore evolve and adapt to the actual contractual situations that are established. This contrasts to other approaches of normative environments (e.g. [5]), where the normative relations that can be established are completely predefined.

We aim at providing an infrastructure in which business entities are represented by software agents and, through them, are able to interact with the computational normative environment using speech acts. When monitoring compliance with norms that apply to specific contracts, the normative environment will be recording a mapping from relevant interactions (which concern business exchanges) that take place. The connection between real-world interactions and the institutional environment is made through illocutions (speech acts) that agents perform with the intent of informing the environment that certain contract-related events have occurred. With an appropriate interface between the normative environment and the statements that agents make, we build an image of relevant real-world transactions that are through this means institutionally recognized. We call this image *institutional reality*, after [6]. The main mechanism that we use in order to certify a real-world event is that of *empowerment* [7]: agents enacting specific roles (e.g. bank, delivery tracker) are seen by the EI as trusted third parties, and are as such certified to obtain specific institutional facts, which are related with business transactions.

We take the stance that it is in the best interest of agents to publicize their abidance to contractual commitments. They do so by interacting with the institution's trusted third parties in order to convince the EI that they are in fact complying.

### 2.1 Modeling and Monitoring Contractual Obligations

When establishing a contract, partners make commitments regarding the business to be enacted, which from the point of view of contract law [8] are expressed as obligations. We model contractual obligations as directed liveline and deadline

obligations to bring about a specific state of affairs [3]:  $Obl_{b,c}(l < f < d)$  represents the obligation of agent  $b$  (the bearer) towards agent  $c$  (the counterparty) to bring about fact  $f$  between liveline  $l$  and deadline  $d$ . A *normative state* records every element that is relevant in contract enactment – institutional reality elements, or *IRE* (Table 1).

**Table 1.** Institutional reality elements used by the computational trust service

Element	Is-a	Specific slots	Description
IRE		contract, when	An IRE pertains to a contract and is obtained at a specific time point
Obligation	IRE	bearer, counterparty fact, liveline, deadline	Prescribed obligation of bearer towards counterparty to bring about a fact between a liveline and a deadline
DeadlineViolation	IRE	obligation	An obligation's deadline was violated
LivelineViolation	IRE	obligation	An obligation's liveline was violated
Fulfillment	IRE	obligation	An obligation was fulfilled
Violation	IRE	Obligation	An obligation was violated

Monitoring rules capture interrelations among these elements, e.g. by saying that if the obliged fact is brought about between the liveline and the deadline then the obligation is fulfilled (this element is added to the normative state). Similarly, if a deadline violation occurs then the obligation might be declared as violated (see [3] for details). Our approach takes advantage of using Jess [9], a forward-chaining rule-based system, which enables a straightforward implementation of monitoring rules.

An electronic representation of a contract includes a set of norms that specify how business is to be enacted. A norm is a rule whose condition analyzes the current normative state and whose conclusion prescribes obligations agents ought to fulfill. Sanctions may be imposed by prescribing obligations upon violation elements.

## 2.2 Interfacing

The normative environment includes a subscription mechanism that enables its use as a tool to alert agents when certain contract-related events occur or are eminent, such as the activation of a contractual obligation or a forthcoming deadline. Agents have to subscribe the normative environment in order to be notified about events related to the contracts in which they participate. This also allows interfacing the contract monitoring service with the Computational Trust service.


Our platform implementation is based on JADE [10], and the subscription mechanism is based on the FIPA-Subscribe interaction protocol [11]. There, an agent sends a subscription message to a service that will be providing notifications of events regarding the kind of information the subscriber is interested in. In our case, the initiator uses a template in order to filter the contracts and events he is interested in.

## 3 Computational Trust System

Whenever an agent (representing a company or individual) is willing to recruit new partners for future joint work, potential risky relationships arises. Electronic

Institutions, because they may track record of all relevant interactions that took place in the past, are well suited to provide trustworthiness information on agents, minimizing the risk in future engagements. Moreover, they are able to follow agents' contract establishment and monitoring, providing an ideal framework for recording the needed information about *how* and *in which context* previous obligations have been dealt with by every agent. This allows us to specify a *situation-aware* trust method that goes beyond traditional, non-contextual trust methods that just estimate a global trustworthiness score for the agents in assessment, lacking a more elaborate and precise information of the agent's adequacy to different particular situations.

We are also concerned with the performance of the method when the trust evidences available about a given target agent are scarce and heterogeneous, and when the activity of the agents under evaluation can span through different situations and contexts. The current implementation of our system that encompasses the proposed method is composed of two different modules, as illustrated in Figure 2.



**Fig. 2.** The current implementation of our trust system

The Aggregator component is responsible for aggregating the available trust evidences of an agent into a trust score. Several trust engines that are defined in the literature can be used ([12]), although we are interested on engines that model the dynamics of trust, as those described in [13] and [14], as they appear to perform better than the traditional statistical approaches. The Contextual Fitness component tunes the outcome of the aggregation step by taking into consideration the specificities of the current business opportunity and the adequacy of the target agent to them.

The idea behind this extension is that if the trust system detects that a target agent has some kind of *handicap* related to the current business necessity, the available overall trustworthiness has to reflect that and further influence the possible selection for future partnerships in similar contexts. At present, agent handicaps are being derived solely from the Fulfilment and Violation IREs (cf. Table 1) received from the normative environment concerning the agent past contractual activities. Our next step is to use other kinds of information related with specific obligations (e.g. liveline and deadline violations) in order to draw a more complete and accurate profile of the agent in that particular situation under assessment.

One good characteristic of this modular approach is that the contextual fitness component can be used together with any conventional trust aggregation engine, being it based on statistical, probabilistic or heuristic models, as it is the case of those reviewed in [12]. Before we further describe the current state of the Contextual Fitness component, we first introduce the notation and the scenario used in the paper.

### 3.1 Scenario and Notation

In this paper, we consider a simulation scenario where, at every round, a given number of agents that want to explore a new business opportunity broadcast a business need specifying a fabric to buy, a quantity and a delivery time. The selection of the best partner to deal with takes into account the estimated trustworthiness of each candidate agent that responds to the specific business need.

We define  $trust_{Ac}(As) \in [0, 1]$  as the trustworthiness value of an agent  $As$ , in the eye of agent  $Ac$ , as computed by a traditional trust aggregator engine, where  $Ac \in C$  is an agent from the set  $C$  of client agents, and  $As \in S$  is an agent from the set  $S$  of supplier agents. We also define adequacy trust  $ad(As, at) \in \{0, 1\}$  as a binary operator for situation-awareness purposes, where  $at \in AT$  describes the *business need*, i.e. an instance of the space  $AT$  of all possible combinations of attribute-value pairs that describe the need (e.g. {fabric='cotton', quantity='900000', delivery time='15'}). For scalability, all numeric values are previously quantified into categories (e.g. *low, medium, high*) using fuzzy logic techniques.

Therefore, the trustworthiness value of agent  $As$  as seen by agent  $Ac$  in the specific context  $at$  is given by the following equation:

$$trust_{Ac}(As, at) = trust_{Ac}(As) * ad_{Ac}(As, at) \quad (1)$$

Finally, a *contractual evidence* represents a transaction at time  $t$  between agents  $Ac$  and  $As$ , for which an outcome  $o \in \{true, false\}$  is generated. It is derived from the *Fulfillment* and *Violation* obligation events produced by the normative environment (Table 1). Therefore, associated to each agent in the Electronic Institution is a history of its past contractual evidences, each represented by the tuple  $\langle Ac, As, at, t, o \rangle$ .

### 3.2 The Contextual Fitness Component

The Contextual Fitness (CF) component is based on an online, incremental and flexible technique of behavior tendencies extraction that we have developed. Current version of CF uses *information gain* ([15]), a well known metric used in machine learning for classification. This metric is based on the *entropy* concept of information theory, and is defined in (2), where  $Gain(S, A)$  is the information gain of attribute  $A$  relative to a collection of samples  $S$ ,  $Values(A)$  is the set of all possible values for attribute  $A$ , and  $S_v$  is the subset of  $S$  for which attribute  $A$  has value  $v$  ([15]).

$$Gain(S, A) \equiv Entropy(S) - \sum_{v \in Values(A)} \frac{|S_v|}{|S|} Entropy(S_v) \quad (2)$$

In our approach, we use this metric to dynamically learn a decision tree from the history of evidences of agent  $As$ , *every time* it is necessary to verify the adequacy of the agent proposal to the current client need. We use all the evidences available, (which might be scarce) about the supplier to build the decision tree. No training or testing phases are performed. After that, the failure tendencies of the agent in evaluation are extracted from the rules pointing to false outcomes. Figure 3 depicts a decision tree that was learnt for a given supplier in a specific experiment we have run.

```

good = cotton
|     dtime = low: false
|     dtime = medium: true
|     dtime = big: null
good = chiffon: null
good = voile: false

```

**Fig. 3. Decision tree generated in our simulations**

For the tree above, our algorithm identified that, at the time of this particular assessment, the agent showed a tendency to fail contracts (outcome=false) that match the tendencies ( $good = cotton, *, dtime = low$ ) and ( $good = voile, *, *$ ). Thus, the trustworthiness value  $trust_{Ac}(As, at)$  of agent  $As$ , as given by Equation 1, would be zero if situation  $at$  matched any of the tendencies derived from the learned decision tree; otherwise, it would be given by the  $trust_{Ac}(As)$  component of Equation 1.

Several issues may arise from the use of the information gain criteria in our technique, such as the need to use similar metrics that permit missing attributes. Also, as we pointed out before, the approach may be further enhanced with more specific information about the agent's behaviour, distinguishing beyond contract failure/success and recording more fine grained information on specific contractual clause violations. We address these improvements in future work.

## 4 Experiments

In order to evaluate our trust model, we use the textile scenario mentioned in the previous section. We generate a population in which all suppliers have different handicaps (95% probability to fail) on performing some particular aspect of a business transaction. For instance, some suppliers tend to fail to deliver fabric in short delivery dates, while others might fail to deliver high quantities of any fabric type. Handicaps on more than one contractual attribute are also possible (e.g. a supplier with a handicap in delivering high quantities of material in low delivery times).


We evaluated three different trust models: *SA* (*SinAlpha*), a trust aggregator using dynamics of trust [14]; *CS*, a model that enhances traditional trust models by considering the situation in assessment [16]. It uses domain specific, predefined similarity metrics to predict unanticipated situations (for an overview of similar models, see [4]); and *CF*, our contextual fitness technique described in Section 3.2, used here in conjunction with the *SA* approach. It is a situation-aware trust model, designed to fit well to non parochial open market scenarios, where the number of available trust evidences for a particular partner agent might be scarce.

### 4.1 Evaluation of the Performance of Trust Models

In a first set of experiments, we evaluate the performance of the *CF* technique and compare it with the other two trust models. We use two metrics: the *average utility* of clients at every round, measured by the ratio given by the number of succeed

contracts over the number of all contracts in the round; and the *number of different suppliers* that were selected by all the negotiating clients at every round.

Figure 4 shows the results obtained. We can observe (at the bottom) that both the SA and the CS approaches are relatively conservative (parochial) concerning the selection of partners, with the 20 clients in the experiment choosing in average 9 to 10.5 different suppliers at each round, while CF explores a slightly higher number of suppliers. This fact seems to be related with the utility achieved by each approach, as can be observed from the top plots of the graphic. In fact, the approach that is able to select from a greater number of different suppliers (the CF approach) also gets in average significant better utility (90.46%) than the other two approaches (83.30% for SA, and 85.87% for CS), leading to a number of succeeded contracts very close to the maximum of 19 (i.e. 95% of 20) contracts.



**Fig. 4.** Average utility obtained (*top*) and average number of selected suppliers (*bottom*)

The results obtained show how the combined use of our technique and detailed contractual information is effective in discovering the particular contract enactment handicaps of the agents in assessment, and how it is able to do so, irrespective of the number of trust evidences available for each agent under evaluation.

#### 4.2 Evaluation of CF Behavior in Open Markets

In a second set of experiments, we evaluated whether the CF ability to explore more supplier agents could safely bring higher utility. Two different kinds of CF clients were generated: a *conservative* (parochial) one, which selects from a more restricted set of known partners, based solely on trustworthiness, and a more *explorative* (non-parochial) one, that explores outside this set. We have also introduced the notion of supplier agents' *value*, reflecting characteristics such as their environmental and ethical policies or usual payment method. The value of each supplier – used by explorative agents – is initialized with a random value in set {0.5, 0.6, 0.7, 0.8, 0.9}. This value is only presented to a client after the first transaction between both agents. Before that, clients estimate a value of 1.0 for unknown suppliers, making it attractive to non-parochial clients to explore new partners. Finally, explorative agents select the partners with whom they will trade based on the *utility* expected from the transaction, which is the *product* of the trustworthiness score of agents and their internal *value*.



We use the following metrics: number of successful contracts achieved by all clients at every negotiation round and respective average number of contracts over all rounds; the number of different suppliers selected at every round; and the average utility achieved by the clients at every round and its average score over all rounds.

We verified that, despite similar results concerning the average number of successful contracts per client (conservative: 90.89%; explorative: 90.67%), the latter leads to a significantly higher utility (75.25%) than the former strategy (68.59%). As can be observed in Figure 5, the strategy used by the clients does not alter in a significant way the number of successful contracts achieved by the clients at every round, nor the number of different suppliers chosen per round. However, the big difference on the results obtained by each strategy resides on the utility achieved through them. In fact, after the first rounds of exploration, the non-parochial strategy got systematically higher utility than the parochial strategy.




Fig. 5. Comparison between parochial and non-parochial client strategies

An important conclusion taken from these experiments is that with our situation-aware trust model, that uses contractual trust evidences, a client agent may feel safe to explore other potentially trustable partners outside its previous group of acquaintances, which in turn can bring him increased benefits (e.g. better prices).

## 5 Conclusions

Doing business electronically should benefit from the development of tools that enable the automation of e-contracting tasks. In this paper we have introduced a way to take advantage of two such tools, made available as services in an Electronic Institution platform. A normative environment providing a contract monitoring facility is used as a source of contract enactment events that feeds a computational situation-aware trust engine.

We experimentally evaluated the benefits derived from combining our situation-aware trust model with the contract events received from the normative environment. We conclude that this combination allows business agents to seek business partners outside their breeding trading acquaintances in a safe environment, conferring to these

agents higher levels of utility than the ones obtained with other trust approaches found in the literature. As future work, we intend to further explore the information provided by the normative environment in order to increase even more the capabilities of the trust system and the effectiveness of the partners' selection process.

**Acknowledgments.** This research is funded by FCT (Fundação para a Ciência e a Tecnologia) project PTDC/EIA-EIA/104420/2008. The first two authors enjoy FCT grants SFRH/BD/39070/2007 and SFRH/BD/29773/2006, respectively.

## References

1. H. Lopes Cardoso and E. Oliveira. Electronic Institutions for B2B: Dynamic Normative Environments. *Artificial Intelligence and Law*, 16(1):107-128, 2008.
2. E. Oliveira, A.P. Rocha. Agents Advanced Features for Negotiation in Electronic Commerce and Virtual Organisations Formation Process. In Dignum & Sierra (eds), *Agent Mediated Electronic Commerce: The European AgentLink Perspective*, pp. 78-97. Springer, 2000.
3. H. Lopes Cardoso, E. Oliveira. Monitoring Directed Obligations with Flexible Deadlines: a Rule-based Approach. In M. Baldoni, J. Bentahar, J. Lloyd, and M. B. Van Riemsdijk, eds., *Declarative Agent Languages and Technologies VII, LNAI*, pp. 51-67. Springer, 2010.
4. J. Urbano, A. P. Rocha, E. Oliveira, Refining the Trustworthiness Assessment of Suppliers through Extraction of Stereotypes. In 12th International Conference on Enterprise Information Systems, Madeira – Portugal, 2010
5. J. L. Arcos, M. Esteva, P. Noriega, J. A. Rodríguez-Aguilar, and C. Sierra. Environment engineering for multiagent systems. *Engineering Applications of Artificial Intelligence*, 18(2):191-204, 2005.
6. J. R. Searle. *The Construction of Social Reality*. Free Press, New York, 1995.
7. A. Jones, M. Sergot. A Formal Characterisation of Institutionalised Power. *Logic Journal of the IGPL*, 4(3):427-443, 1996.
8. P. S. Atiyah. *An Introduction to the Law of Contract*. Clarendon Law Series. Clarendon Press, Oxford, 4th edition edition, 1989.
9. E. Friedman-Hill. *Jess in Action*. Manning Publications Co., 2003.
10. F. Bellifemine, A. Poggi, G. Rimassa. Jade: A FIPA-compliant agent framework. In *Practical Applications of Intelligent Agents and Multi-Agents*, 97-108, 1999.
11. FIPA, FIPA Subscribe Interaction Protocol Specification, available at <http://www.fipa.org/specs/fipa00035/index.html>
12. J. Sabater, M. Paolucci, On Representation and Aggregation of Social Evaluations. *Computational Trust and Reputation Models, Int. J. Approx. Reasoning* (2007).
13. C. M. Jonker, J. 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 Multi-Agent World: Multiagent System Engineering*. (1999).
14. J. Urbano, A. P. Rocha, E. Oliveira, Computing Confidence Values: Does Trust Dynamics Matter?, in L. Seabra Lopes, N. Lau, P. Mariano & L. M. Rocha (eds.), *Progress in Artificial Intelligence (EPIA 2009)*, Springer, ISBN 978-3-642-04685-8, 520-531 (2009).
15. J. R. Quinlan, "Induction of Decision Trees", *Machine Learning*, 1(1), 1986, 81-106.
16. M. Rehak, M. Gregor, M. Pechoucek, Multidimensional context representations for situational trust. *IEEE Workshop on Distributed Intelligent Systems: Collective Intelligence and Its Applications*, 315-320 (2006).