Legal Knowledge and Information Systems (Jurix 2008)
E. Francesconi et al. (eds)
IOS Press, 2008, 2003

A Compliance Model of Trust

Guido Governatori [a,1], Subhasis Thakur [b] and Duy Hoang Pham [a,b]
[a] *NICTA, Queensland Research Laboratory, Australia*
[b] *School of ITEE, The University of Queensland*

Abstract. We present a model of past interaction trust model based on compliance of expected behaviours.

1. Introduction

Agents face the challenge of how to decide whether to trust or not other agents. Many models of trust and reputation have been put forward, and the central idea of these models is that they use information from past interaction to deduce the trustworthiness of agents in terms of their competency and reliability. Basically these models get a feedback from the history of interactions of an agent and then, based on some parameters compute how trustworthiness the agent is. A limitation of these models is that, very often the feedback is a subjective feedback from human users of the agent. This means that these models are of limited use for fully autonomous agents, where, a full automation is often required. Additionally agents often operates autonomously in open and unpredictable environments where there could be different reasons for the (partial) failures of an interaction with other agents.

To obviate the issues we have alluded to above we have developed a method to automatically provide feedback for an interaction among agents based on a compliance model. The intuition behind this idea is that interactions among agents do not happen in isolation, but these interactions are governed by rules (norms, social obligations, contracts, policies and so on). The key point is that all these rules describe (and prescribe) the expected behaviour of an agent. The feedback provided is then the degree of compliance of a past interaction with respect to the expected behaviour for that particular interaction.

While autonomous agents are designed for a particular purpose and domain, at design time our model relies on domain experts to provide rules describing the expected and acceptable behaviours for business interactions in the domain. Then the system elaborates the input from domain experts and generates a set of rules (norms) prescribing the behaviour agents have to comply with. At run time the system gathers data about the current interaction (in form of values of environment and state variable, actions performed and so on), and determines what an agent has to do to be compliant with the expected behaviour in the given situation, what scenarios while not compliant are still acceptable and what scenarios are not compliant and not acceptable. This means that the system uses

---

[1] Corresponding author: NICTA, Queensland Research Laboratory, PO Box 6020, St Lucia, 4067, Australia; Email: guido.governatori@nicta.com.au

the case data to identify what rules have been obeyed to and at what level of compliance. The output is a (numeric) score specifying the degree of compliance of an interaction of an agent. The result can then be used as the value of the interaction to compute the trust of the agent based on the history of interactions for the agent.

Notice that the proposed mechanism is not a stand-alone trust management system but it must be used in conjunction with methodologies to give trust metrics based on the history of past interactions. In addition, as it has been recently suggested [6], a trust management system should comprise several components to handle different types of trust information such as interaction trust, witness reputation, role-based trust, certified reputation components etc.

2. Compliance Model of Trust

The assumption underlying our approach is that there is a direct relationship between the trustworthiness of an agent and how compliant the agent was in past interactions. Accordingly we have to provide an account of what compliance is and how to handle it. Compliance can be understood in terms of the normative positions (i.e., obligations, prohibitions, etc.) an agent has to comply with. This means that to tackle this issue one has to adopt a formalism capable to model and reason with such notions.

Many formalisms have been proposed to represent normative notions such as obligations, prohibitions and permissions. In this paper, we adopt FCL (Formal Contract Language) [3] as formalism to model the expected behaviour of an agent in interactions with other agents. FCL is a combination of an efficient non-monotonic formalism (defeasible logic [1]) and a deontic logic of violations [5]. This particular combination allows us to represent exceptions as well as the ability to capture violations, the obligations resulting from the violations, and the compensation for the violations.

2.1. Formal Contract Language

Deontic Logic extends classical logic with the modal operators O, P and F. The interpretations of the formulas OA, PA and FA are, respectively, that A is obligatory, A is permitted and A is forbidden. The modal operators obey the usual mutual relationships

$$OA \equiv \neg P \neg A \qquad \neg O \neg A \equiv PA \qquad O \neg A \equiv FA \qquad \neg PA \equiv FA$$

and are closed under logical equivalence, i.e., if $A \equiv B$ then $OA \equiv OB$, and satisfy the axiom $OA \rightarrow PA$ (i.e., if A is obligatory, then A is permitted) that implies the internal coherency of the obligations in a set of norms, or, in other words, it is possible to execute obligations without doing something that is forbidden. Thus, obligatory actions are the actions expected to be performed by an agent, while prohibitions describe unacceptable outcomes for an agent.

Norms describing the behaviour of an interaction between agents usually specify actions to be taken in case of breaches of some of the norms in them. These can vary from (pecuniary) penalties to the termination of an interaction itself. This type of construction, i.e., obligations in force after some other obligations have been violated, is know in the deontic literature as contrary-to-duty obligations (CTDs) or reparational obligations (because they are activated when normative violations occur and are meant to 'repair' or 'compensate' violations of primary obligations [2]). Thus a CTD is a conditional obligation arising in response to a violation, where a violation is signalled by an unfulfilled obligation. This type of construction identifies situations that are not ideal for the interaction but still acceptable for it. The ability to deal with violations or potential violations and the reparational obligation generated from them is one of the essential requirements for agents where, to the nature of the environment where they are deployed, some failures can occur, but where it does not necessarily mean that the whole interaction has to fail. We represent the norms an agent has to comply as rules, where a rule is an expression

$$r : A_1, \ldots, A_n \Rightarrow B,$$

where r is the (unique) name of the rule, $A_1, \ldots, A_n$ are the *premises* (propositions in the logic), and B is the *conclusion* (also a proposition of the logic). The propositions of the logic are built from a finite set of atomic propositions, and the following operators: $\neg$(negation), O (obligation), P (permission), and $\otimes$ (violation/reparation). Given a rule r we use $A(r)$ to denote the set of premises of the rules, and $C(r)$ for the conclusion. For a set of rules R, we use $R[C]$ do denote the subset of R of rules where the conclusion is C.

If p is an atomic proposition, then $\neg p$ is a proposition. Given a proposition p we use $\sim p$ to denote the complement of p: i.e., if $p = l$, then $\sim p = \neg l$ and if $p = \neg l$, then $\sim p = l$. If p is a proposition, then Op is an *obligation proposition* and Pp is a *permission proposition*; both are called *deontic propositions*. If $p_1, \ldots, p_n$ are obligation propositions and q is a deontic proposition, then $p_1 \otimes \cdots \otimes p_n \otimes q$ is a *reparation chain*. Given a reparation chain C, we use $\pi_i(C)$ to denote the i-th element of the chain; finally $|C|$ returns the length of C, i.e., the number of elements in C.

A simple proposition corresponds to a factual statement. A reparation chain captures obligations and normative positions arising in response to violations of obligations. For example, $B_1 \otimes B_2$ means that the process is obliged to perform B_1; and in case B_1 is not fulfilled (i.e., the obligation is violated), the "secondary" obligation B_2 must be fulfilled. While single obligations and permissions (and their negations) can appear in the premises of a rule, reparation chains can be used only in rule conclusions.

FCL is equipped with a superiority relation over the rule set. The superiority relation ($\prec$) determines the relative strength of two rules, and it is used when rules have potentially conflicting conclusions. For example given the rule $r_1 : A \Rightarrow OB \otimes OC$ and $r_2 : D \Rightarrow O\neg C$. $r_1 \prec r_2$ means that rule r_1 prevails over rule r_2 in situation where both fire and they are in conflict (i.e., rule r_1 fires for the secondary obligation OC).

The aim of this paper is to identify whether a given agent is compliant with a set of rules. Thus we must be able to determine all and only obligations generated by the case data relevant for the agent in a particular interaction with other agents. To this end we use the *normalisation* procedure of FCL (see [4]) that merges rules generating obligations and rules having the negation of the obligation in the premises (or so called contrary-to-duty rules). For example, given two rules $A \Rightarrow OB$ and $C, \neg B \Rightarrow OD$ can be merged into the rule $A, C \Rightarrow OB \otimes OD$. In addition the normalisation process uses a subsumption algorithm to remove redundant rules. At the end of the process the normalisation procedure generates a set of rule with all unique maximal reparation chains. The compliance checkers use the maximal chains to determine whether a task in a process and then a process itself complies with a given set of rules.

Besides the normalisation mechanism FCL has a second reasoning mechanism to determine the set of conclusions (obligations and reparation chains) in force for a specific case. As we have already remarked FCL is an extension of defeasible logic with the reparation operator ($\otimes$). Accordingly the reasoning mechanism to derive conclusion is an extension of that for defeasible logic. To derive a conclusion, let us say A, we have to scan the set of rules for a rule having A in its conclusion, and the rule should be applicable; this means that all literals in the antecedent of the rule are either facts or are already provable. Defeasible logic is a skeptical non-monotonic formalism, that is, it does not support conflicting conclusions. Therefore, when we want to prove a conclusion A then we have to ensure that all 'reasons' for $\neg A$ are not usable to prevent the conclusion of A. Thus, for all rules having $\neg A$ in the conclusions we have to check that either the rules do not fire (i.e., at least one of the premises do not hold) or the rule for $\neg A$ is weaker than an applicable for A. In defeasible logic the conclusions of a rule is a single literal and not a reparation chain. Thus the condition that A appears in the conclusion of a rule means in defeasible logic that A is the conclusion of the rule. For FCL have to extend the notion to accommodate reparation chain. The required change is that to prove A, we have to consider all rule with a reparation chain for A, where for all elements before A in the chain, the negation of the element is already provable. Thus to prove A given the rule $P_1,\dots,P_n \Rightarrow C_1 \otimes \cdots \otimes C_m \otimes A \otimes D_1 \otimes \cdots \otimes D_k$, we have that $P_1,\dots,P_n$ must be all provable, and so must be $\neg C_1,\dots,\neg C_m$. For the full details see [3].

2.2. Ideal Semantics

In a way, FCL constraint expressions for a set of rules define a behavioural and state space which can be used to analyse how well different behaviour execution paths (including state constraints) comply with the FCL constraints. Our aim is to use this analysis as a basis for deciding whether the execution of an interaction among agents is compliant with the rules expressed in FCL. The central part of this compliance checking is given by the notions of ideal, sub-ideal, non-ideal and irrelevant situations.

Intuitively an *ideal* situation is a situation where execution paths do not violate FCL expressions, and thus the execution paths (which will then correspond to processes that are related to the norms describing the expected behaviour of the agents) are fully compliant with the norms. A *sub-ideal* situation is situation where there are some violations, but these are repaired, in the CTD sense. Accordingly, processes resulting in sub-ideal situations are still compliant to a contract even if they provide non-optimal performances of the norms. A situation is *non-ideal* if it violates the norms (and the violations are not repaired). In this case a process resulting in a non-ideal situation does not comply with the norms describing the expected behaviour. There are two possible reasons for a process not to comply with the norms: 1) the process executes some tasks which are prohibited by the norms (or equivalently, it executes the opposite of obligatory tasks); 2) the process fails to execute some tasks required by the norms. Finally a situation is *irrelevant* for a set of norms if no rule is applicable in the situation. Irrelevant situations correspond to states of affairs where the set of norms is silent about them.

As discussed in Section 2.1 given a set of rules prescribing the behaviour of an agent we compute the normal form of it, where the normal form contains all conditions that can be derived from the rules and redundant clauses are removed. Thus normal forms are the most appropriate means to determine whether a process (corresponding to the behaviour of an agent) conforms with a set of norms. Accordingly we use the normal form to check the compliance of an agent. We now define conditions under which we are able to determine whether a situation complies with a set of normative specifications or if it represents a violation of some clauses.

First of all we define when a situation (set of literals) is either ideal, sub-ideal, non-ideal or irrelevant with respect to a rule (norm).

DEFINITION 1

- *A situation S is* ideal *with respect to a rule* $\Gamma \Rightarrow A_1 \otimes \cdots \otimes A_n$ *iff* $\Gamma \cup \{A_1\} \subseteq S$.
- *A situation S is* sub-ideal *with respect to a rule* $\Gamma \Rightarrow A_1 \otimes \cdots \otimes A_n$ *iff* $\Gamma \cup \{A_i\} \subseteq S$, *for some* $1 < i \leq n$ *such that* $\forall A_j$, $j < i$, $A_1, \ldots, A_j \notin S$.
- *A situation S is* non-ideal *with respect to a rule* $\Gamma \Rightarrow A_1 \otimes \cdots \otimes A_n$ *iff* $\Gamma \subseteq S$ *and S is neither ideal nor sub-ideal.*
- *A situation S is* irrelevant *with respect to a rule* $\Gamma \Rightarrow A_1 \otimes \cdots \otimes A_n$ *iff it is neither ideal nor sub-ideal nor non-ideal.*

According to Definition 1, a situation is ideal with respect to a norm if the rule is not violated; sub-ideal when the primary obligation is violated but the rule allows for a reparation, which is satisfied; non-ideal when the primary obligation and all its reparations are violated, and irrelevant when the rule is not applicable. Definition 1 is concerned with the status of a situation with respect to a single rule, while the behaviour of an agent, typically, is described/regulated by many rules, thus we have to extend this definition to cover the case of a set of rules. In particular we will extend it considering all rules in the normal form for the set of rules governing the expected behaviour of an agent.

DEFINITION 2

- *A situation S is* ideal *with respect to an FCL normal form iff there is no rule in the normal form for which S is either sub-ideal or non-ideal or irrelevant.*
- *A situation S is* sub-ideal *with respect to an FCL normal form iff there is a rule for which S is not irrelevant and it is sub-ideal, and there is no norm in the normal form for which S is non-ideal.*
- *A situation S is* non-ideal *with respect to a FCL normal form iff there is no rule in the normal form for which S is not irrelevant and is non-ideal.*
- *A situation S is* irrelevant *with respect to an FCL normal form iff for all rules in the normal form S is irrelevant.*

Definition 2 follows immediately from the interpretation we have provided in Definition 1. On the other hand, the relation between a normal form and the set of rules from which it is obtained seems to be a more delicate matter. A careful analysis of the conditions for constructing an FCL normal form allows us to state the following general criterion:

DEFINITION 3 *A situation S is* ideal *(*sub-ideal, non-ideal, irrelevant*) with respect to a set of FCL rules if S is ideal (sub-ideal, non-ideal, irrelevant) with respect to the normal form of the set of FCL rules.*

It is worth noting that Definition 3 shows the relevance of the distinction between a set of rules and its normal form. This holds in particular for the case of sub-ideal situations. Suppose you have the following set of FCL rules $\Rightarrow OA$ and $\neg A \Rightarrow OB$. The corresponding normal form is $\Rightarrow OA \otimes OB$. While the situation with $\neg A$ and B is sub-ideal with respect to the latter, it would be non-ideal for the former. In the first case, even if $\neg A \Rightarrow OB$ expresses in fact an implicit reparational obligation of the rule $\Rightarrow A$, this is not made explicit. The key point here is that there was no link between the primary and reparation obligations in the original set of rules, but this is made explicit in the normal form. So, there exists a situation which apparently accomplishes a rule and violates the other without satisfying any reparation. This conclusion cannot be accepted because it is in contrast with our intuition according to which the presence of two rules like $\Rightarrow OA$ and $\neg A \Rightarrow OB$ must lead to a unique rule. For this reason, we can evaluate a situation as sub-ideal with respect to a set of FCL rules only if it is sub-ideal with respect to its normal form.

Given a set of rules R and a set of literals S (plain literals and deontic literals), we can use the inference mechanism of defeasible logic to compute the set of conclusions (obligations) in force given the set of literals. These are the obligations an agent has to obey to in the situation described by the set of literals. However, the situation could already be a sub-ideal situation, i.e., some of the obligations prescribed by the rules are already violated. Thus, given a set of literals describing a state-of-affairs one has to compute not only the current obligations, but also what reparation chains are in force given the set.

Consider a scenario where we have the rules $A \Rightarrow OB$ and $\neg B \Rightarrow OC$, and the situation is described by A and $\neg B$. The normal form of the rules is $A \Rightarrow OB \otimes OC$ and $\neg B \Rightarrow OC$. The only obligation in force for this scenario is OC. Since we have a violation of the first rule ($A \Rightarrow OB$ and $\neg B$), then we know that it is not possible to have an ideal situation here. What we have to do is to identify the chains for the ideal situation for the task at hand. To deal with this issue we have to identify the *active* reparation chains.

DEFINITION 4 *A reparation chain C is* active *given a set of literals S, if*

1. $\exists r \in R[C] : \forall a_r \in A(r), a_r \in S$ *and*
2. $\forall s \in R[D]$ *such that* $\pi_1(C) \in D$, *either*
   1. $\exists a_s \in A(s) : \sim a_s \notin S$, *or*
   2. $\exists i\ \pi_i(D) = \sim \pi_1(C)$ *and* $\exists k, k < i, \sim \pi_k(D) \notin S$, *or*
   3. $\exists t \in R[E]$: $\pi_j(E) = \pi_1(C)$, $\forall a_t \in A(t), a_t \in S$, $\forall m, m < j, \sim \pi_m(E) \in S$ *and* $t > s$.

Let us examine the following example. Consider the rules

$$r_1 : A_1 \Rightarrow OB \otimes OC, \qquad r_2 : A_2 \Rightarrow O\neg B \otimes OD, \qquad r_3 : A_3 \Rightarrow OE \otimes O\neg B.$$

The situation S is described by A_1 and A_3. In this scenario, the active chains are $OB \otimes OC$ and $OE \otimes O\neg B$. The chain $OB \otimes OC$ is active since r_2 cannot be used to activate the chain $O\neg B \otimes D$. For r_3 and the resulting chain $OE \otimes O\neg B$, we do not have the violation of the primary obligation OE of the rule (i.e., $\neg E$ is not one of the literals in S), so the resulting obligation $O\neg B$ is not entailed by rule r_3.

2.3. Degree of Compliance

Compliance distance is measured to indicate the degree of match between a the set of rules describing the expected behaviour of an agent for a particular type of interaction and the actual behaviour in an instance of an interaction of the given type. The degree of compliance measures how distant is the execution of a particular instance of an interaction by an agent from the ideal execution of the same.

Given a set of rule R defining the intended behaviour of an agent and a set of literals F (corresponding to the case data describing the actual behaviour of an agent), we determine C, the set of the reparation chains relevant w.r.t the case data.

```
Input: $F$: set of literals; $C$: set of reparation chains
Output: $d \in [0,1]$: degree of compliance
  for each $c \in C$
    let $i \in \mathbb{N} : \pi_i(c) \in F$ and $\forall j < i : \pi_j(c) \notin F$
      $d(c) = \frac{1+|c|-i}{|c|}$
  $d = \frac{\sum_{c \in C} d(c)}{|C|}$
```

The idea behind the above algorithm, is that the degree of compliance is the ratio between the number of fulfilled obligations and the total number obligation in force for a transaction. First of all, given a scenario for an agent, we consider only the reparation chains that are relevant to the particular behaviour of the agent. Then for every relevant chain we identify whether the situation (identified by the case data) is ideal, sub-ideal or non-ideal.

Every ideal situation is compliant, the agent did what it was expected to do, and only one obligation (the primary obligation, the most preferred outcome) was in force for that reparation chain. Thus the execution get a score of 1 for every chain for which the situation is ideal. For chains, for which the situation is sub-ideal, the score the agent get is between 0 and 1. Consider the chains

$$OA \otimes OB \otimes OC, \qquad OD \otimes OE \otimes OF, \qquad OG \otimes OH$$

and we have A, B, $\neg D$, E and $\neg H$. For the first rule we have $d = 1$, the behaviour of the agent complies with the rule. For the second rule $d = 2/3$, the primary obligation is violated, but the compensation is fulfilled. Lastly for the third rule we do not have G, so the obligation is not fulfilled. In addition, the violation has not been compensated since we have the violation of the contrary-to-duty resulting from the violation of the primary obligation; accordingly the degree of compliance for the last rule is 0. Based on these values we have that the degree of compliance for the interaction is $(1 + .66 + 0)/3 \approx .56$

3. Related Work

The model we have proposed captures the degree of compliance of a single interaction between two agents. As such it does not provide trust metrics based on the history of interactions between agents. Therefore our approach must be supplemented by additional mechanism to 'aggregate' previous interactions to provide a trust value. In this section we quickly review popular trust/reputation mechanisms and we show how to integrate our method to the aim of removing subjectivity from these models.

3.1. Centralised reputation mechanism

We start by discussing centralised trust models, and we show how the trust model proposed in the previous section can be integrated in the proposed trust/reputation methodologies.

3.1.1. eBay

In the eBay model [8], buyers place bids to buy from a seller. The buyer is obliged to pay the agreed amount and the seller is responsible for shipment of the good in agreed time frame and also for accurateness of the advertised good. The buyer and the seller can report to a central authority about their interaction as ‘not delivered in time’ or ‘check bounced’ etc. The central authority calculates these feedbacks and combines with the history of the buyer and the seller and calculates the reputation. The reputation is then available for future buyers or sellers.

In this model, subjectivity lies in the feedback that users can send to the central authority. The available feedbacks are [Positive, Negative, Neutral, None].

We can use our model of compliance in the eBay model as follows: The central authority should have access to the ‘contract’ between buyer and seller, determining the expected behaviour for the transaction. For example, the contract is as follows:

1. $\mathit{Paid} \Rightarrow O\mathit{DeliveryInTime} \otimes O\mathit{CashBack}$
2. $\mathit{Paid} \Rightarrow O\mathit{MatchAdvertisedProperties} \otimes O\mathit{CashBack}$
3. $\mathit{WillingToPay} \Rightarrow O\mathit{PayByCheque} \otimes O\mathit{MoneyOrder}$

The first two rules show the seller’s obligation in the interaction. According to first rule, if the buyer has paid then she must get the good within the agreed delivery time; otherwise the seller should return the money. The second rule suggests that the seller should deliver the exact product that she has advertised, otherwise she should return the payment from buyer. The third rule says that the buyer who has won a bid and willing to pay can pay by cheque and in case that the cheque bounces then the payment is done by money order.

In case the seller delivers a product which does not match with the advertisement, she returns the buyer’s payment. The buyer reports a compliance value 0.5. Similarly if the delivery was not in time and the seller has refunded the price, the buyer reports compliance value 0.5. All of the reports from the buyer can be aggregated to generate the overall report. Also, the seller can generate a report about the buyer.

SPORAS [12] extends [8]: in particular, its trust calculation algorithm includes the following:

- New user has a minimum reputation.
- Reputation of a user does not fall below the reputation of a new user.
- Trust value of user is updated according to feedback from other agents.
- Most recent rating has more weight.

Our compliance model can be incorporated in SPORAS as in [8].

4. Decentralised trust management systems

In this section, we quickly outline decentralized models of trust and we show how our compliance based trust model can be used in conjunction with decentralised trust models to minimise the degree of subjectivity.

4.1. Reputation management

In [7] a decentralised reputation management system has been proposed. In this model agents have incentives to reveal their interaction results, there is a group of broker agents who collect these interaction reports and generates the reputation measure for each agent, where the reputation r_i of an agent a_i is calculated as

$$r_i = \frac{\sum_{j=1}^{k} report_j}{N}$$

where $report_j$ is generated by the other agents about agent a_i. Each report $report_j$ can take value either 0 or 1. So 'subjectivity' is introduced in this model as interaction can be mapped to two predefined values (0 for noncompliance and 1 for compliance). In this model, agents' interactions are modelled using the Prisoner Dilemma where players use pure strategies. Clearly, we can use our trust model to generate the report and to provide more accurate report. However, this requires agents to apply mixed strategies.

4.2. Regret

In the Regret trust model [9], agents keep the record of their interactions as they maintain rating of agents in a database. The database is queried to calculate the trust. Each rating is associated with a weight, which depends on the recency of the rating.

In this model $outcome(O)$ is defined as dialogue between the agents which reflects either negotiation to establish a contract between them or a course of actions taken by agents with some initial contracts. The reputation/trust of an agent is computed based on the 'impression' the agent made on other agents in previous interaction. The impression is modelled as tuple $Impression = (a, b, O, \chi, t, W)$ where a, b are agents (a is judging b's impression), O is outcome as described before, χ is the variable that is 'topic' of interaction between a and b, $W \in [-1, 1]$ is rating of b by a.

The Regret model is silent on how the actual values of W are computed. Our compliance model can be integrated with Regret as follows: The variable outcome O corresponds to the agents actual behaviour and variable χ corresponds to the expected behaviour on an agent. Then we can apply the techniques discussed in 2.1 and 2.3 to derive the exact value of W.

4.3. Referral systems

In referral systems [10,11], agents keep track of agents they know and their expertise. An agent can query its known agents about information. And, if these agents could not answer the query then they will refer to some other agents who may answer.

The idea of these models is that when agent a_i evaluates reputation of agent a_j, a_i will use results of its direct interaction with a_j and other agents' recommendation about a_j. The direct interactions with a_j are represented as $S_{j1}, S_{j2}, \ldots S_{jH}$ (last H interactions) and $S_{jk} \in [0, 1],\ k \leq H$. This model does not discuss about how S_{jk} is mapped into $[0, 1]$. We can use our model of compliance to derive the exact mapping from S_{jk} to $[0, 1]$.

4.4. FIRE

FIRE [6] allows combination of different information sources as direct interaction, witness report, third party references etc in order to generate reputation of an agent. A direct interaction between two agents a and b, a's rating about agent b is represented as $r = (a, b, c, i, v)$, where i is the interaction between agents, v is the rating which takes a value in terms of c, c can have values as [quality, honesty] and v has value in the range

$[-1, +1]$. [6] does not provide exact mechanism how v is calculated and predefines values in $[-1, +1]$ as subjective measure as -1 for negative, 0 for neutral etc. Similarly the model provides the same structure for roles instead of agents.

In this model the trust reputation of an agent is based on a vector of values where each value represents a parameter of interest for agent interactions. Our model can be applied to define a set of rules for each parameter, therefore, determining the degree of compliance v for each of such parameter.

5. Conclusion

In this paper we proposed a qualitative method to remove the subjectivity in current trust models. As far as we are aware of this is the first work that proposes a formal approach to trust based on compliance. The qualitative evaluation is achieved by distinguishing the desired behaviours with the actual behaviours of an agent w.r.t. a set of norms governing the interaction between agents. Also our mechanism can be easily incorporated into existing trust management mechanisms.

References

[1] Grigoris Antoniou, David Billington, Guido Governatori, and Michael J. Maher. Representation results for defeasible logic. *ACM Transactions on Computational Logic*, 2(2):255–287, 2001.

[2] José Carmo and Andrew J.I. Jones. Deontic logic and contrary to duties. In D.M. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic, 2nd Edition*, volume 8, pages 265–343. Kluwer, Dordrecht, 2002.

[3] Guido Governatori. Representing business contracts in RuleML. *International Journal of Cooperative Information Systems*, 14:181–216, 2005.

[4] Guido Governatori, Zoran Milosevic, and Shazia Sadiq. Compliance checking between business processes and business contracts. In *10th International Enterprise Distributed Object Computing Conference (EDOC 2006)*, pages 221–232. IEEE Computing Society, 2006.

[5] Guido Governatori and Antonino Rotolo. Logic of violations: A Gentzen system for reasoning with contrary-to-duty obligations. *Australasian Journal of Logic*, 4:193–215, 2006.

[6] Trung Dong Huynh, Nicholas R. Jennings, and Nigel R. Shadbolt. An integrated trust and reputation model for open multi-agent systems. *Autonomous Agents and Multi-Agent Systems*, 13(2):119–154, 2006.

[7] Radu Jurca and Boi Faltings. An incentive compatible reputation mechanism. In *AAMAS*, pages 1026–1027, 2003.

[8] P. Resnick and R. Zeckhauser. Trust among strangers in internet transactions: Empirical analysis of ebay's reputation system. In M. R. Baye, editor, *The economics of the internet and e-commerce*, volume 11, 2002.

[9] Jordi Sabater and Carles Sierra. REGRET: A reputation model for gregarious societies. In *Fourth workshop on deception fraud and trust in agent societies*, pages 61–70, 2001.

[10] Bin Yu and Munindar P. Singh. An evidential model of distributed reputation management. In *AAMAS*, pages 294–301, 2002.

[11] Bin Yu and Munindar P. Singh. Searching social networks. In *AAMAS*, pages 65–72, 2003.

[12] Giorgos Zacharia and Pattie Maes. Trust management through reputation mechanisms. *Applied Artificial Intelligence*, 14(9):881–907, 2000.