Analyzing Control Trust in Normative Multiagent Systems

Joris Hulstijn[1] Yao-Hua Tan[1] Leendert van der Torre[2,3]
1. Vrije Universiteit, Amsterdam, {jhulstijn,ytan}@feweb.vu.nl
2. CWI, Amsterdam, torre@cwi.nl
3. Delft University of Technology

Abstract

It has been argued that transaction trust is composed of party trust and control trust. In this paper we study control trust: trust in an institution that has set up a control mechanism. We present an account of control mechanisms using normative multiagent systems. Control mechanisms consist of constitutive norms which define evidential documents, and regulative norms which define violation conditions and sanctions. The account is illustrated by an analysis of the Letter of Credit trade procedure.

1 Introduction

Trust is an important concept in electronic commerce. When a transaction has to be arranged over a long distance by individuals that do not know each other, and do not share a social or cultural background, lack of trust is likely. We are interested in this notion of *transaction trust*. Suppose individual a, the trustor, is about to engage in a transaction with individual b, the trustee. The success of the transaction for a crucially depends on the behavior of b, but a is not in a position to monitor or control b's behavior. Because a is vulnerable to the actions of b, a's decision to engage in the transaction depends on his level of trust, c.f. (Mayer, Davis, and Schoorman 1995).

In general, transaction trust can be based on a personal relationship, on the social role that the other party is playing in some institution, on personal past experiences, or on a reputation, which is based on past experiences of other trusted parties. Such reasons for trusting the other party are called *party trust*. In the absence of direct reasons for trusting the other party, a trustor may turn to institutions that enforce control mechanisms. Think of an escrow service, or a quality standard maintained by an international standards body. The mere presence of a control mechanism is not enough. The trustor must understand how the mechanism works, and must have trust in the institution that enforces it. This is called *control trust*. The general idea can be summarized as follows (Tan and Thoen 2000; Tan and Thoen 2002).

Transaction Trust = Party Trust + Control Trust

Regarding party trust, a lot of research effort has gone into on-line reputation mechanisms, and methods for learning people's trustworthiness on the basis of past experiences, e.g. (Macy and Skvoretz 1998; Castelfranchi and Tan 2000). Much less attention has been paid to control trust. However, we believe that this factor is as important as the other factors in predicting transaction trust.

In this paper we will have a look at the way an individual understands a control mechanism, and how this affects the decision to trust another party. Thus we have the following two research questions.

1. How can we model control trust?
2. How does control trust affect transaction trust?

The research method is to make a model of the relevant aspects of the individual's decision making process, and use a logic to represent the model. In general such logical models can be used for simulation and analysis, or can be implemented directly in an intelligent trust management system. An example of the latter kind of model is found in declarative policy languages for access control (Blaze, Feigenbaum, and Lacy 1996). Our model is intended to be used for purposes of analysis and simulation.

The particular type of model we use in this paper is that of a normative multiagent system (NMAS). Normative multiagent systems are

> sets of agents (human or artificial) whose interactions can fruitfully be regarded as norm-governed; the norms prescribe how the agents ideally should and should not behave. [...] Importantly, the norms allow for the possibility that actual behavior may at times deviate from the ideal, i.e., that violations of obligations, or of agents rights, may occur (Jones and Carmo 2002).

We will use the version of NMAS developed by Boella et al. (2004; 2005). In particular, we apply the possibility of *recursive modeling* that is afforded by multiagent systems: based on a profile of other agents, an agent can simulate several scenario's, and select an appropriate decision on the basis of those simulations. A profile is nothing but a set of expectations. This reflects the fact that we treat trust as an epistemic notion. So in our account, the amount of trust is based on the information agents have about other agents.

Control mechanisms are modeled here by a collection of constitutive and regulative rules. Constitutive rules define the institutional concepts. They are modeled here as generic beliefs. Regulative rules define violation conditions and sanctions. Such rules are enforced by an abstract entity, called a normative system. Crucially, trustor agents do not only have a profile of the trustee, but also a profile of the normative system that enforces the control mechanisms. The idea is depicted in Figure 1. Party trust is only based on the expectations of the trustor about the trustee. These expectations can be based on additional background information, such as personal acquaintance or shared cultural values. In the figure, it is represented by the vertical interaction between trustor and trustee. Control trust is more complex. It is based on expectations of the interaction between trustor and trustee, given the interaction of trustee with a normative system.

Obviously, the simulations we run are largely based on assumptions. Often these assumption are derived from specific knowledge of the case. Technically, they are

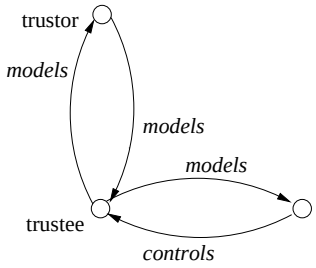

Figure 1: Recursive modeling of trustor, trustee and normative system

stored in a model M. For example, model M_{LC} contains specific assumptions about the Letter of Credit procedure.

The rest of the paper is structured as follows. To illustrate the approach we discuss the Letter of Credit trade procedure in section 2. In section 3 we define normative multiagent systems, as well as constitutive and regulative rules. In section 4 we analyze the Letter of Credit procedure in a normative multiagent system.

2 Letter of Credit

Consider the well known documentary credit procedure of a Letter of Credit (Bons 1997; Lee 2000; Kartseva, Gordijn, and Tan 2004). Suppose we have a potential transaction between a customer and a supplier which is located in a remote and unfamiliar part of the world. Hence there are no common conventions or trade procedures. In such circumstances a lack of trust is warranted. The supplier does not want to ship the goods without first receiving payment, but the customer does not want to pay before the goods have been shipped. To solve this deadlock situation banks introduced the Letter of Credit procedure (LC). A Letter of Credit is an agreement that the bank of the customer, called the issuing bank, will arrange a credit to guarantee payment as soon as the supplier can prove that the goods have been shipped. The supplier can prove this by presenting the shipping documents, such as a Bill of Lading, to his own bank, the corresponding bank. The shipping documents are issued by the carrier of the goods, as soon as the goods have in fact been shipped. The corresponding bank transfers the shipping documents to the issuing bank, which only delivers the shipping documents to the customer after payment. With the shipping documents the customer can then reclaim the goods from the carrier.

A possible way of modeling the procedure is depicted in Figure 2, adapted from (Lee 2000). It indicates the order in which documents must be exchanged, which provides information about the dependencies among actions: which actions will occur only provided some other actions have occurred. For example, the objective of step 4 is to let the supplier know that a credit has been secured, and that he can safely start shipping the goods. Similarly, the objective of the shipping documents, is to provide evidence that the goods have been transferred.

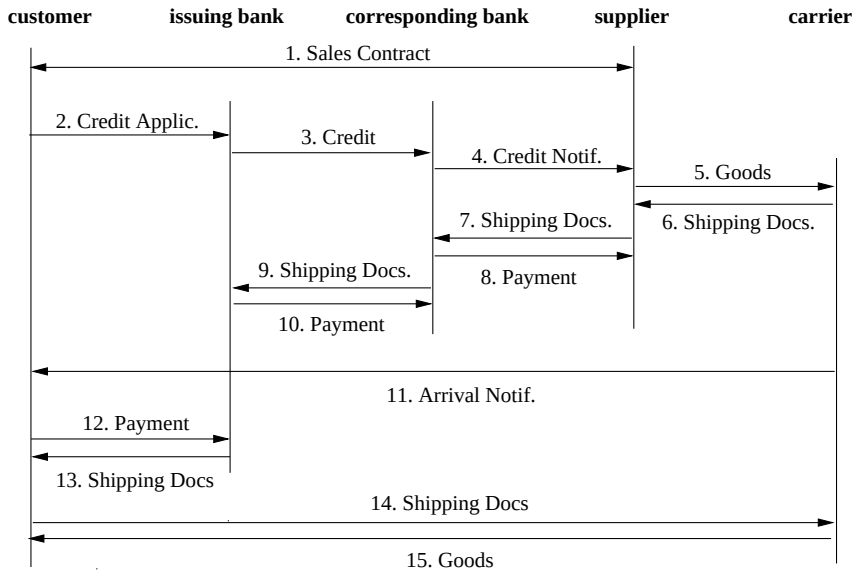

Figure 2: Interaction Diagram of the Letter of Credit Procedure

By means of the Letter of Credit procedure the original lack of trust between customer and supplier is replaced by a professional banking relationship between the issuing bank and the corresponding bank. The banks provide a service, a credit procedure, which is paid for by a fee. The fee must cover the handling costs and the insurance costs against the risk that the credit will not be redeemed. The professional banking relationship is much more trustworthy than the original transaction relationship, for a number of reasons.

First, documentary credit procedures generally conform to the guidelines of the International Chamber of Commerce. Moreover, shipping documents have an *evidentiary effect*, guaranteed by the United Nations (1980). Therefore, in case of a legal conflict, no party can claim not to have known about this evidentiary effect.

Second, trading banks tend to have long term relationships. The success of these relationships is dependent on their reputation. In case a customer has a dispute with a bank over a Letter of Credit procedure, it might start a legal case, possibly damaging the reputation of the bank.

Third, banks are regulated by their respective Central Banks. That means that, at least locally, banks can act as an institution that enforces adherence to the rules.

3 Normative Multiagent Systems

We present the basic idea of normative multiagent systems here. For a more technical exposition, we refer to Boella and van der Torre (2004).

3.1 Beliefs and Goals

Each agent has a profile of the mental attitudes of other agents. A profile contains the beliefs (information) and goals (internal motivations). External motivations, such as social norms or laws, can be represented in the form of obligations. Later we show how obligations are reduced to a combination of beliefs and goals.

Mental attitudes are represented in a logical language, as sets of conditionals or production rules. This expresses the fact that mental attitudes are context dependent (Hansson 1969), and that their application is conditional on certain constraints. So each attitude Bel, $Goal$, is represented by a set of rules of the form $A \rightarrow B$, where both A and B are formulas, composed of facts by means of logical operators $\wedge$ (and), $\vee$ (or), $\neg$ (not) and the constant $\top$ (always true). Here A represents the conditions under which the facts represented by B may be inferred by the agent. Moreover, B may contain special decision variables, also called actions, that will alter that state of the world. For simplicity, both facts and decision variables are represented by boolean variables, being either true or false. The decision making process of an agent is represented by a forward reasoning loop, which runs roughly as follows.

The agent receives input from observation, represented as a set of facts S. Alternatively, the agent may start with a set of initial goals, represented by a set of decision variables S. Now the agent tries to match each rule $A \rightarrow B$ against S. If A is contained in S, and the facts of B do not contradict a fact in S, the rule is applicable. However, there may be several rules applicable to S, from the same and from different mental attitudes, each with a different possible outcome. Using a priority order, the agent selects one rule – this is called conflict resolution – and applies it: the result B is now added to S. This process continues, until a stable set of facts is reached, to which no further additions can be made. Such a stable set, an extension, represents one of the possible outcomes of the decision making. The decision making behavior of an agent crucially depends on the way conflicts among the mental attitudes are resolved. Different priority orders may lead to different extensions, which represent sets of goals and hence lead to different behavior. For example, a selfish agent will prefer goals to social obligations; a social agent will let obligations take priority.

Example 1. Buying and Budget Suppose an agent is a compulsive collector. Whenever it finds a rare toy, it wants to buy it. Because of a tight budget, the agent made up its mind not to buy any more toys. Now suppose the agent finds another rare toy. What will happen?

Belief: rare_toy
Goal 1: rare_toy $\rightarrow$ buy_toy
Goal 2: $\neg$ buy_toy

What happens depends on the relative strength of the agent's urge to buy the car (goal 1) and its resolution not to (goal 2). These relative strengths can be expressed by a priority ordering on rules. In general, belief rules outrank all other goal rules; otherwise the agent would suffer from *wishful thinking*. If the second goal rule would outrank the first goal rule, the agent will refrain from buying. Otherwise, it will buy despite its resolutions.

Priority: Belief $>$ Goal 1 $>$ Goal 2
Outcome: $\{$ rare_toy , buy_toy$\}$

The example demonstrates that an agent considers different possible goals. Roughly, there are two kinds of goals. Achievement goals are satisfied once some state of affairs has been realized. For example to buy an object. Maintenance goals on the other hand, are only satisfied for as long as some state of affairs continues to hold, for example the goal to stay within budget.

3.2 Norms and Obligations

Beliefs and goals are mental attitudes of individual agents. But how about social norms? General norms produce obligations for each agent that has accepted the norm. We propose a reduction of obligations to goals of the normative system, where the normative system itself is seen as a separate agent. This may be summarized by the slogan "Your wish is my command". The reduction makes use of a *violation predicate* (Anderson 1958). Violation detection is separated from sanctioning. Sanctioning concerns the decision to sanction an agent in case a violation has been detected, and what sanction is most appropriate. To make sure that normative systems do not act arbitrarily, a number of conditions must be observed (Boella and van der Torre 2004).

Definition 1 (Regulative Norm) Agent a is obliged to bring about x in circumstances C, under sanction s, with respect to a normative system n in a given model M, written as $M \models \mathsf{Obl}_{a,n}(x, s|C)$, if and only if the following conditions hold.

1. If agent n believes that C, then it has as a goal that x should be brought about.

   Goal of n: $C \rightarrow x$

2. If agent n believes that C and $\neg x$ is the case, then it has the goal $\mathsf{Viol}(\neg x, a)$, i.e., to detect $\neg x$ as a violation by agent a.

   Goal of n: $C \wedge \neg x \rightarrow \mathsf{Viol}(\neg x, a)$

3. Agent n has as a goal that there are no violations. This is to prevent arbitrary violation detection.

   Goal of n: $\neg\mathsf{Viol}(\neg x, a)$

4. If agent n believes C and detects $\mathsf{Viol}(a, \neg x)$, then it has the goal to sanction agent a.

Goal of n: $C \wedge \mathsf{Viol}(\neg x, a) \rightarrow s$

5. By itself, agent n does not have a goal to sanction. This is to prevent arbitrary sanctioning.

   Goal of n: $\neg s$

6. Agent a has the goal $\neg s$. Without this condition, sanction s would not deter agent a from violating the obligation.

   Goal of a: $\neg s$

To illustrate the definitions, we will use an example of some aspects of the Letter of Credit procedure. A complete analysis follows in section 4.

Example 2. Evidentiary documents The model M_{LC} is constructed as follows. We use boolean variables 'LC', 'shipping_docs', 'transfer_goods', 'fraud', and 'law_suit' to represent the facts of operating in the context of a Letter of Credit procedure, presenting the shipping documents, having transferred the goods from the custody of the supplier to the carrier, committing fraud, and starting a law suit, respectively. Variable a ranges over the relevant agents: customer (c), issuing bank (ib), corresponding bank (cb), supplier (s) and carrier (ca). We use un to denote the United Nations: the institution that defines the meaning of shipping documents. We use n to denote the normative agent that is able to enforce a norm. Once a letter of Credit procedure has been established, this role is usually taken by the issuing bank. Because the issuing bank has long term trading relationships with many parties involved, it can come up with a credible sanction for each violation. In the example, let us suppose that the sanction will be a lengthy law suit, with possible additional repercussions for reputation. We have an obligation to respect the meaning of the shipping documents, against a penalty of a lengthy law suit.

$M_{LC} \models \mathsf{Obl}_{a,n}(\neg\text{fraud}, \text{law\_suit}|\text{LC})$, if and only if:

1. Goal of n: LC$\rightarrow \neg$fraud
2. Goal of n: LC $\wedge$ fraud $\rightarrow \mathsf{Viol}(\text{fraud}, a)$
3. Goal of n: $\neg\mathsf{Viol}(\text{fraud}, a)$
4. Goal of n: LC $\wedge \mathsf{Viol}(\text{fraud}, a) \rightarrow$ law_suit
5. Goal of n: $\neg$ law_suit
6. Goal of a: $\neg$ law_suit

Without the obligation, committing fraud would be desirable. For example, for a carrier it is profitable to issue fake shipping documents. Moreover, everybody knows we are operating in the context of a Letter of Credit procedure.

7. Goal of a: shipping docs$\wedge\neg$transfer_goods
8. Belief of a, n: LC

3.3 Constitutive Norms

A normative system defines when institutional facts hold by *constitutive norms*. In general, there are two kinds of norms. Constitutive norms generate or constitute the institutional context, by means of establishing institutional facts (Searle 1995). For example, the rules of chess constitute the game by defining legal moves and positions. Regulative norms, by contrast, restrict the behavior of agents through obligations and enforcement mechanisms.

A constitutive norm applies only under certain circumstances and is intimately linked to an institution. This institution can be a (legal) person or an abstract entity such as a community of users. Whether a rule applies depends on the jurisdiction of the institution. Thus constitutive rules are of the form "x counts as y under circumstances C in institution i". Several logical representations of counts-as rules have been studied, e.g., (Jones and Sergot 1996). Counts-as rules associate one type of event or fact with another type of event or fact.

In our NMAS model constitutive norms are modeled as belief rules of the institution. Moreover, having the rule is imperative for all agents that fall under the jurisdiction of the institution.

Definition 2 (Constitutive Norm) Event or fact x counts as event or fact y under circumstances C, to all agents a that fall under the jurisdiction of i in a given model M, written $M \models \mathsf{Counts\_as}_{a,i}(x, y|C)$, if and only if:

1. Belief of a, i: $C \wedge x \rightarrow y$
2. If x results from an action, the agent executing the action must play a role in institution i.

There is one relevant constitutive rule in our example.

$M_{LC} \models \mathsf{Counts\_as}_{a,un}((\text{shipping docs} \wedge \neg \text{transfer\_goods}), \text{fraud}|\text{LC})$,
i.e., Belief of un, a: LC $\wedge$ shipping_docs $\wedge \neg$transfer_goods $\rightarrow$fraud

3.4 Recursive Modeling

Now we apply recursive modeling. Suppose agent a operates in the context of a Letter of Credit procedure, has a goal to cheat, and a goal to avoid law suits. So we get an initial extension of the form {LC, fraud, $\neg$ law_suit}. Because the consequences of these potential goals are to a large part controlled by the other party, agent a will try and predict the decisions of agent n. A set of rules to predict another agent's decisions is called a *profile*. Crucially, profiles contain rules that are affected by the actions of the agent itself. That is why this process is called *recursive modeling*. However, for most applications no infinite recursion is needed; recursive models up to three levels of embedding are usually quite sufficient. Here we only need one level of embedding.

First, a compares rule 2 with rule 3. Which of these will get priority depends on a's profile of n. Lets assume that a believes that rule 2 outranks 3. That will produce an extension {LC, fraud, $\mathsf{Viol}(\text{fraud}, a)$ $\neg$ law_suit}. Second, a compares rule 4 and 5. If a expects n not to sanction, for example because of prohibitive costs of lawyers,

a will form the goal to cheat. But if a expects n to act on detected fraud, this would produce an extension { LC, fraud, $\mathsf{Viol}(\text{fraud}, a)$, law_suit, $\neg$ law_suit }, which contains a contradiction. This conflict must be resolved by a's own priority order. In case the penalty outweighs the expected profits, the agent will not form a goal to commit fraud: { LC, $\neg$ law_suit }.

The example shows that an obligation only works when two conditions are met. First, the sanction must outweigh the benefits of a violation.

(Goal of a: $\neg$ law_suit) $>$ (Goal of a: fraud)

Second, the expected relative priorities of being detected and sanctioned, must be sufficiently high.

Profile of a:
(Goal of n: LC $\wedge$ fraud $\rightarrow \mathsf{Viol}(\text{fraud}, a)$) $>$ (Goal of n: $\neg\mathsf{Viol}(\text{fraud}, a)$)
(Goal of n: LC $\wedge \mathsf{Viol}(\text{fraud}, a) \rightarrow$ law_suit) $>$ (Goal of n: $\neg$ law_suit)

Violation detection and sanctioning can be delegated to specific agents, such as a police force. In thew Letter of Credit case, violation detection is to a large extend the responsibility of the victims. For example, when a carrier has not transferred the goods to the customer, the customer can take the shipping documents, and file a complaint, i.e., notify the normative system of the violation.

4 Letter of Credit in NMAS

In this section we remodel the Letter of Credit procedure in normative multiagent systems, to indicate how control trust plays a role in transaction trust.

We will model two sets of circumstances: one in which no control mechanism is in place, and one in which a Letter of Credit procedure has been established.

4.1 Without Letter of Credit

We have the following scenario. Supplier s must decide to trust customer c in a transaction that consists of two actions: payment from c to s, and shipment from s to c, where price and details of delivery have been previously agreed. The sales contract would consist of two obligations:

$\mathsf{Obl}_{c,n}$(payment, in_debt | shipping)
$\mathsf{Obl}_{s,n}$(shipping, in_debt | payment)

However, initially there is no credible normative system n that could enforce this contract. Parties are left to enforce their own contract, i.e. use external motivations like threats. But although parties can easily detect violations, they do not have the power to make a credible threat.

Without documentary credits, there are still reasons for an agent to honor the contract. First, an agent could have the internal motivation that being in debt is morally wrong, and therefore undesirable. If this disposition were known, for example through

shared cultural values, the agent can be trusted. Second, when the trade relationship is supposed to last for a longer period, a credible threat would be to end the trade relationship and ruin possible future transactions. In both cases adding such information would beg the question.

So, in the absence of any further information, in the profile of the supplier, the customer has no incentive to pay after delivery. Similarly, the customer will expect the supplier not to deliver the goods after payment.

Profile of s:
(Goal of c: $\neg$ payment) $>$ (Goal of c: $\neg$ in_debt)
Profile of c:
(Goal of s: $\neg$ shipping) $>$ (Goal of s: $\neg$ in_debt)

4.2 With Letter of Credit

What we need is a mutual obligation in which non-compliance of one, irrevocably leads to failure for the other. Such a direct exchange of goods for money is only possible, when parties are physically present, and in possession of the valuables.

$\mathsf{Obl}_{c,n}(\text{payment}, \neg\text{shipping}|\text{shipping})$
$\mathsf{Obl}_{s,n}(\text{shipping}, \neg\text{payment}|\text{payment})$.

Because a simultaneous exchange is not possible in international trade, the mutual obligation must depend on evidentiary documents.

$\mathsf{Obl}_{c,n}(\text{payment}, \neg\text{shipping}|\textit{evidence of shipping})$
$\mathsf{Obl}_{s,n}(\text{shipping}, \neg\text{payment}|\textit{evidence of credit})$.

Such documents can be created through a letter of Credit procedure. The rules of the procedure (figure 2) are constituted by the international trading conventions, encoded by the United Nations. The issuing bank acts as a normative system n to enforce adherence to these procedures. The issuing bank is a good candidate for enforcing the rules. First, its sanctions are credible. It is in a position to damage the other parties, for example by starting a law suit. Starting a lawsuit would not have been possible for the supplier for example, because it would involve a foreign legal system. Second, violation detection is relatively easy. Because of the evidentiary documents, typically the victim of fraud will be able to prove that a violation occurred.

So we must show that because of the dependencies specified in the constitutive rules of the Letter of Credit procedure, and the regulative enforcement by the issuing bank, the link between 'evidence of shipping' and actual shipping, and the link between 'evidence of credit' and actual payment, is strong enough to let both parties engage in the transaction (Figure 2).

In other words, we must show that, in the profile of the supplier,

1. the corresponding bank has an incentive to pay, upon receiving the shipping documents (step 7-8), and consequently
2. the carrier will have an incentive to deliver the shipping documents, in return for the goods (step 4-6).

Similarly, we must show that, in the profile of the customer,

1. the carrier has an incentive to transport the goods and deliver them to the customer, in return for the shipping documents (step 14-15), and hence
2. the corresponding bank and the issuing bank have incentives to transfer the shipping documents from the supplier to the customer, in return for their respective payments (steps 9-10,12-13)

For lack of space, we can only sketch two of these requirements. The other steps are analogous.

Step 7-8 According to the supplier, why should the corresponding bank pay? First, not paying would count as violation for the corresponding bank at this stage of the Letter of Credit procedure. It would risk a law suit. Second, it does not have much reason to withhold payment, since earlier it received credit from the issuing bank (step 3). And third, it needs the shipping documents, in order to get payment (step 9-10). Because the supplier will not hand over the shipping documents, there is a direct motivation for complying. Only the first suggestion is worked out in detail. It is a special case of example 2.

Profile of s:
$\mathsf{Obl}_{cb,ib}$(payment,law_suit) |LC$\wedge$ shipping_docs)

As we analyzed, this requires the following conditions:

Profile of s:
(Goal of cb: $\neg$ transfer_goods) $<$ (Goal of cb: $\neg$law_suit))
(Goal of ib: LC $\wedge$ shipping_docs $\wedge\neg$ payment $\rightarrow \mathsf{Viol}(\text{fraud}, cb)) \ >$ (Goal of ib: $\neg\mathsf{Viol}(\text{fraud}, cb)$)
(Goal of ib: LC $\wedge\ \mathsf{Viol}(\text{fraud}, cb) \rightarrow$ law_suit) $>$ (Goal of ib: $\neg$ law_suit)

Step 14-15 Again, this is a special case of example 2. Note that the carrier must already be in possession of the goods; otherwise the arrival notification (step 11) must have been false. When the goods are not handed over, the customer has a right to redress: he could notify the issuing bank (ib) that the carrier violated the Letter of Credit procedure. In that case, the carrier risks a law suit.

Profile of c:
$\mathsf{Obl}_{ca,ib}$(transfer_goods,law_suit|LC $\wedge$ shipping_docs$\wedge$arrival_notif)

As indicated, this will only work in case the following conditions are met.

Profile of c:
(Goal of ca: $\neg$ transfer_goods) $<$ (Goal of ca: $\neg$ law_suit)
(Goal of ib: LC $\wedge$ shipping_docs $\wedge$ arrival_notif $\wedge\neg$ goods $\rightarrow \mathsf{Viol}(\text{fraud}, ca)) \ >$
(Goal of ib: $\neg\mathsf{Viol}(\text{fraud}, ca)$)
(Goal of ib: LC $\wedge\ \mathsf{Viol}(\text{fraud}, ca) \rightarrow$ law_suit) $>$ (Goal of ib: $\neg$ law_suit)

Note that customer c is in a position to assess these conditions, at least for ib. Typically, such expectations are based on aspects of party trust (personal experience, role, reputation).

5 Conclusions

In this paper we studied transaction trust and its components party trust and control trust. Where party trust is based on expectations of parties about each other, control trust also involves expectations of parties about a normative system, and the way the other party will react to that normative system.

The research questions can now be answered as follows.

1. Control trust is generated by the existence of a control mechanism, and an institution that is perceived to be actively enforcing it. control mechanisms may contain both evidentiary documents and obligations. Evidentiary documents can be accounted for by constitutive norms. Obligations can be reduced to a system of beliefs and goals of the normative system and the subject, according to the principle "Your wish is my command'.

   Our case study of the Letter of Credit procedure, showed that it makes sense to model control mechanisms as a combination of constitutive norms, that define institutional facts, and regulative norms, that enforce compliance to obligations through violation detection and sanctioning.

2. Control trust will affect transaction trust by the fact that, according the trustor's profile, the trustee will let its decisions depend on the way a normative system is enforcing a control mechanism. That means that, according the trustor's profile, (1) the normative system must be willing and able to detect and sanction possible violations, and (2) the trustee must prefer to avoid a sanction over a violation.

   In our case study, we contrasted a situation without a Letter of Credit, in which the trustor would be most likely not to trust the trustee, with a situation in a which a Letter of Credit procedure is established. In that case, the trustor is indeed likely to trust the trustee, provided conditions (1) and (2) are met.

Trust aspects of international trade procedures can be generalized to electronic commerce. Like international trade, electronic commerce often takes place between parties that are separated in time and space, do not know each other, and do not share a cultural background. Based on our initial analysis, we suggest the following steps for designing a control mechanism:

1. Standards for evidentiary documents (constitutive norms) must be known and respected globally.
2. Use of a standard must be enforced locally through regulative norms, by a normative system that has the institutional power to apply credible sanctions.
3. Start with a mutual obligation, in which the sanction for both parties is simply non-compliance of the other party.
4. Set up a causal chain such that evidence of compliance can replace compliance itself, as an incentive for the other party to engage in a transaction.

In further research will be about mutual obligations during a simultaneous exchange. These can be analyzed as a mutual threat or promise: I will comply only if you comply.

References

Anderson, A. (1958). A reduction of deontic logic to alethic modal logic. *Mind*, 67:100–103.

Blaze, M., J. Feigenbaum, and J. Lacy. (1996). Decentralized trust management. In *IEEE Symposium on Security and Privacy*. IEEE, pages 164–173.

Boella, G. and L. van der Torre. (2004). Regulative and constitutive norms in normative multiagent systems. In *Procs. of 9th International Conference on the Principles of Knowledge Representation and Reasoning*, pages 255–265.

Boella, G., J. Hulstijn, and L. van der Torre. (2005). Virtual organizations as normative multiagent systems. In *Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05)*, volume 7, page 192c.

Bons, R. W. H. (1997). *Designing Trustworthy Trade Procedures for open Electronic Commerce*. Ph.D. thesis, University of Rotterdam.

Castelfranchi, C. and Y. Tan, editors. (2000). *Trust and Deception in Virtual Societies*. Kluwer.

Hansson, B. (1969). An analysis of some deontic logics. *Noûs*, 3:373–398.

Jones, A. and J. Carmo. (2002). Deontic logic and contrary-to-duties. In D. Gabbay, editor, *Handbook of Philosophical Logic*. Kluwer, pages 203–279.

Jones, A. and M. Sergot. (1996). A formal characterisation of institutionalised power. *Journal of the Interest Group in Pure and Applied Logic*, 3:427–443.

Kartseva, V., J. Gordijn, and Y.-H. Tan. (2004). Analysing preventative and detective control mechanisms in international trade using value modelling. In *Proceedings of the Sixth International Conference on Electronic Commerce (ICEC'04)*, pages 51–18. ACM Press.

Lee, R. M. (2000). Documentary Petri Nets: A modeling representation ffor electronic trade procedures. In W. van der Aalst, editor, *Business Proces Management*, LNCS 1806. Springer Verlag, pages 359–375.

Macy, M. W. and J. Skvoretz. (1998). The evolution of trust and cooperation between strangers: A computational model. *American Sociological Review*, 63:638–660.

Mayer, R., J. Davis, and F. Schoorman. (1995). An integrative model of organizational trust. *Academy of Management Review*, 20(3):709–734.

Searle, J. (1995). *The Construction of Social Reality*. The Free Press, New York.

Tan, Y.-H. and W. Thoen. (2000). An outline of a trust model for electronic commerce. *Applied Artificial Intelligence*, 14(8):849–862.

Tan, Y.-H. and W. Thoen. (2002). Formal aspects of a generic model of trust for electronic commerce. *Decision Support Systems*, 33(3):233 – 246.

United Nations. (1980). United nations convention on international multimodal transport of goods. Technical report, UNCITRAL (Geneva, 24 May 1980).