

Abstract Norms and Electronic Institutions

F. Dignum

Institute for Information and Computing Sciences
Utrecht University
dignum@cs.uu.nl

Abstract. Electronic institutions regulate the interactions between parties that conduct transactions. The rules (or norms) that the parties should comply to within the institution are described in statutes and regulations of the institution. Unfortunately these descriptions are often very abstract and vague. Therefore it is practically impossible to check whether the protocols that are followed during the interaction actually fulfill all these norms or not. In this paper we discuss this problem and describe how the connections between abstract norms and the concrete institutional protocols can be made.

1. Introduction

The study and modeling of norms has attracted the interest of scientists from different disciplines such as sociology, economics, philosophy, and computer science.

According to sociology, a norm is a rule or standard of behavior shared by members of a social group (Encyclopedia Britannica). According to philosophy, a norm is an authoritative rule or standard by which something is judged and on that basis approved or disapproved (Columbia Encyclopedia). Examples of norms include standards of right and wrong, beauty and ugliness, and truth and falsehood. According to economics, a norm (from *norma*, Latin for carpenter's level) is a model of what should exist or be followed, or an average of what currently does exist in some context, such as an average salary among members of a large group.

Norms can be either informal, when there is no formal regulation of the norm and no formal sanction for violating the norm, or they can be formal when they are incorporated in laws or regulations from the institutions that regulate the behavior of the people within the society. In this paper we will concentrate on the formal norms and in particular on formal norms in the context of electronic institutions. See e.g. [3,4] for more information on the treatment of informal norms.

Electronic institutions, such as auctions and market places are the electronic counterparts of institutions that are established in our societies. They are established to regulate the interactions between parties that are performing some (business) transaction. One of the main roles of institutions is to inspire trust into the parties that perform the transaction. (see [6] for more details on the roles of institutions). One way to inspire trust in the parties is by incorporating a number of regulations (norms) in the institution that indicate the type of behavior to which each of the parties in the trans-

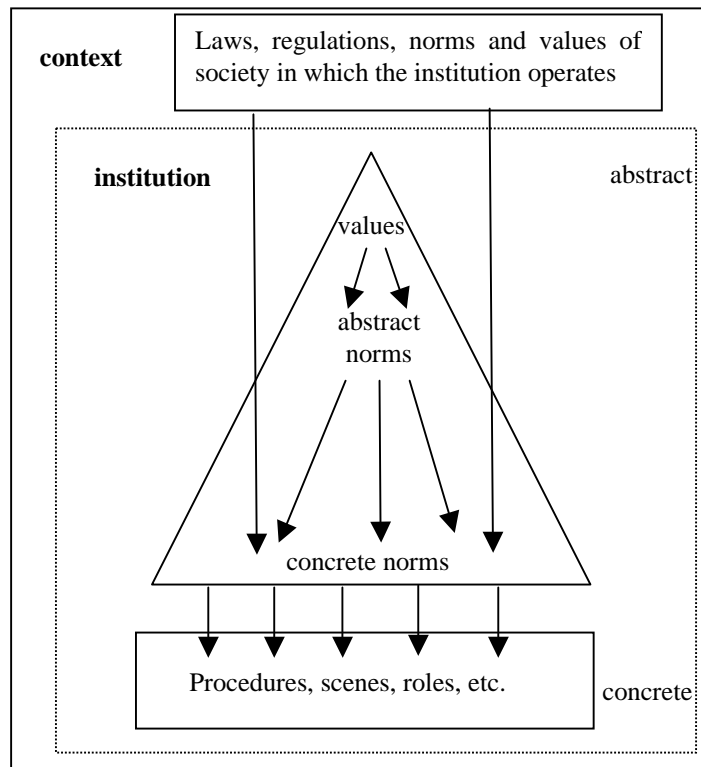
action should adhere within that institution. The main question of this paper is how the norms specified in the regulations can be incorporated in the structure of an institution such that the agents operating within the institution will operate according to these norms or can be punished when they are violating the norms.

Many attempts have been made to formally specify norms in e.g. deontic logic ([13,16]). Although it is possible to capture the norms in this way and even give them a certain kind of semantics and reason about the consequences of the norms etc. this kind of formalization does not yet indicate how the norm should be interpreted within a certain institution. For instance, we can formalize a norm like “it is forbidden to discriminate on the basis of age” in deontic logic as “ $F(\text{discriminate}(x,y,\text{age}))$ ” (stating that it is forbidden to discriminate between x and y on the basis of age). However, the semantics of this formula will get down to something like that the action “ $\text{discriminate}(x,y,\text{age})$ ” should not occur. However, it is very unlikely that the agents operating within the institution will explicitly have such an action available. The action actually states something far more abstract.

We claim that the level on which the norms are specified is more abstract and/or general than the level on which the processes and structure of the institution are specified. Therefore we need to “translate” the norms specifically to a level where their impact on the institution can be described directly. This translation is dependent on the domain of the institution and therefore the translation rules depend on e.g. the ontology for that domain.

Because even the concrete translations of norms do not have a direct counterpart in the institution, another translation is needed on this level to indicate how the norm is implemented. For instance, the norm to pay when you have bought a product can be implemented by restricting the available actions in the institution after the buying action to just the paying action. However, one might also implement this norm by not allowing an agent to leave the institution before he has paid (in case he bought something). This means that the agent can still perform all kinds of actions, but always has to pay at some time.

The way norms are specified and used in institutions as described above could be depicted as follows:



The picture also shows that the context in which the institution functions plays a role in the specification and implementation of the norms in the institution. Although this is an important topic, we will only say a few words about this relation in this paper and leave a more thorough discussion of this relation for another paper. The main focus of this paper is to point out the differences in levels of abstraction and possible relations between the levels denoted by the arrows in the figure above.

The rest of this paper is organized as follows. In the next section we will say a few words about the influence of norms on agents. In section 3 we will describe the issues pertaining to the translation of abstract norms into concrete norms. In section 4 we describe some ways in which concrete norms can be implemented in the institution. In section 5 we draw some conclusions and indicate the areas for further research.

2. Agents and Norms

In each social system one can study the social interaction between the parties in the system from two different viewpoints. One can look at the interaction from the individual's point of view or from the society's point of view. In the first viewpoint one is mainly interested in the way an individual handles the interaction. How does it react to signals from other parties? Does it have a fixed protocol to handle interactions or does it reason about the signals? What is the "best" way to react to a signal, etc.?

From the society's point of view one is mainly interested in properties of the interaction itself. E.g. is the interaction "fair"? Is it "efficient"? Although the two viewpoints deal with different concerns, they are of course not completely independent. If an interaction protocol that is defined from the society's point of view requires a party to perform some action in order to respond then, of course, one should make sure that all parties are capable to perform this action. E.g. if a request for information should be answered with the information (if available) in some (belief) database then the parties should be able to access their database. (this becomes more complicated if the requirement is that the information can be deduced from the database!).

In the rest of this paper we will take the society's point of view concerning the role of norms in the interaction between agents. In particular, we will take the institution's point of view. However, before getting into that viewpoint, in this section, we will briefly discuss the requirements that we can or should impose on the agents interacting within the institution.

Whenever norms are used within the context of multi-agent societies and institutions the question arises how the agents should cope with these norms. Elsewhere I have argued that it is beneficial for an agent society if agents are aware of the norms and have the capability to reflect upon them (e.g. [5,2]). Although I still support this point of view it does not mean that electronic institutions that function according to a set of norms can only operate with agents that are aware of these norms. As we will see later all norms are in the end implemented in the institution through some restrictions on possible behavior or as triggers on some unwanted behavior. I.e. the norms are no direct, explicit components of the interactions defined in the institution, but function as one of the requirements on these interactions.

One can design agents in such a way that they will always behave according to these restrictions (or take into account the possible consequences of their behavior). The agents do not have to “know” that these restrictions follow from specific norms. Therefore the overhead in reasoning about the possible consequences of their behavior is (almost) zero. This means that these agents can be relatively simple and small. Of course, one might also represent the norms explicitly in the institution (next to the interactions). In that case one can (for the same institution) also design agents that are more advanced and are able to reason about the norms and their implementation. These agents might exhibit the same behavior in most circumstances as the simple agents, but will be able to make better choices in special circumstances. They know when it pays off to violate a rule.

One could compare this to a situation in which people have to follow the norms of the company in which they work. They might just obey the rules that implement the norms without ever thinking (or knowing) about the norms. E.g. the norm of working 8 hours each day can be implemented by the company by requiring the employees to arrive at 8:00 o'clock and leave at 17:00 (having a lunch break between 12:00 and 13:00). Employees might just follow these rules without ever reasoning about the norm. However, if at one time they have to go to the doctor before going to work they might be in trouble to get to work in time. A person that is able to reason about the norms will be able to see that in this case it is more important to go to the doctor than to be at work in time (and might still conform to the norm through working after 17:00). See [9] for more discussion on this difference of treating norms.

So, institutions that will function according to some explicit norms do not force a huge overhead on the agents that operate within that institution. They only create the possibility for more intelligent behavior of the agents that are capable of coping with such an overhead.

3. From Abstract Values to Concrete Norms

3.1. Statutes: Objectives, Values and inherited norms

The characteristic of institutions is that they enforce a certain behavior between members of the group of persons (or agents) that carry out their interactions within the jurisdiction of the institution. This behavior must be of a type that it contributes to a certain pre-specified objective and set of values and norms. At the most abstract level these three elements can be found in the statutes of the institution. The statute will indicate the main objective of the institution, it will indicate the values that direct the fulfilling of this objective and it indicates the norms given by society or the group under whose responsibility the institution falls.

For example the statutes of the National Organization for Transplants (ONT) ([17]) in Spain states that it's main objective is to increase the number of organ donations and the subsequent increase in available organs for transplants. This objective by itself might lead to all kinds of speculations on the way organ donation is stimulated

and whether the objective is in any way restricted by ethical values. Therefore the ONT states also that it operates according to the regulation of the national health system and that it strives to distribute the donated organs in the most appropriate and correct way according to the current technical knowledge and according to the ethical principles of equality. (This is my interpretation of the objectives and might not be completely accurate. However, for the current discussion this rough translation from the Spanish text will suffice).

The ONT states that it operates according to the regulations of the national health system. It therefore “inherits” the norms and values of this system as well and they restrict the ways in which the objectives of the ONT can be reached. The above shows that institutions hardly ever operate in isolation and therefore frequent references are made to other regulations and institutions.

The latter part of the objectives of the ONT actually indicates the values according to which the ONT operates. I.e. to distribute the donated organs in an “equal” way among potential recipients. Where “equal” is defined by both technical and ethical standards. This value will also play a role in the regulations that will determine the actual process according to which the transplantation should be performed.

At the highest level of abstraction the values fulfill the role of norms in the sense that they determine the actions that we should or should not take in a certain situation. **Values** are beliefs that we have about what is important, both to us and to society as a whole. A value, therefore, is a belief (right or wrong) about the way something should be. It should be apparent that values, by definition, always involve judgements (since they tell us how something should be). In short, the values we hold are general behavioral guidelines. They tell us what we believe is right or wrong, for example, but they do not tell us how we should behave appropriately in any given social situation. This is the part played by **norms** in the overall structure of our social behavior.

So, how could we specify and relate the concepts of value, (inherited) norms and objectives?

The norms and values that are inherited from other institutions can be described at the top level of the institution explicitly. These norms have the highest priority and will “overwrite” any norms specified by the institution itself if they are contradictory. One could model this using a prioritized deontic logic, where default rules are used. In this paper we will not get into the theory of hierarchies of norms, but refer to [15] for this aspect. We will discuss a more concrete representation of the norms itself in the next section on abstract norms.

The values of the institution can be described as desires and/or goals. E.g. the value of “equal distribution of donated organs” can be described as “V(equal_distribution)”. However, besides a formal syntax, this does not provide any meaning to the concept of “value”. In the context of institutions the meaning of the values is defined by the norms that contribute to this value. (See [14] for a more in depth discussion about the aspects related to values). A norm contributes to a value if fulfilling the norm always leads to states in which the value is more fully accomplished than the states where the norm is not fulfilled. E.g. a norm contributing to the above value would be that the age of the potential recipients cannot be taken into account when determining the distribution of the organs. So, each value has attached to it a list of norms that contribute to that value. The total list of norms together “defines” the meaning of the value in the context of the institution.

Finally, the objectives of the institution can be represented as the goal of the institution. As far as the institution has control over the actions of the agents, acting within that institution, it will try to ensure that they perform actions that will lead to the overall goal of the institution. This means that agents performing an official role within the institution (such as administrator, broker, auctioneer, etc.) will have a goal that is equal to some sub-goal of the overall goal of the institution. All the visiting agents will be expected to at least not obstruct the institution to obtain its main goal. And actually all the norms are directed such that the objectives, values and inherited norms of the institution are obtained (or preserved).

3.2. From abstract norms to concrete norms

Having set the most abstract level of normative behavior in the statutes, we will now turn to the description of the norms themselves. Norms are expected, socially acceptable, standards of behaving in any given social situation. In order to check norms and act on possible violations of the norms by the agents within an institution, the abstract norms have to be translated into actions and concepts that can be handled within the institution. Concrete norms pertain to actions that are described in terms of the ontology of the institution and from which therefore the meaning and effect is known or they pertain to situations that can be checked directly by the institute. E.g. a concrete norm might be that a buyer should pay a deposit before he can start bidding at an auction. In this case the “pay deposit” is a basic action and the “bidding” is also a basic action within the institution. The norm that a buyer must be of legal age (>21) might be concrete if the institution has a means to check the age of the buyer directly. However, if the institution and/or the buyer must follow some procedure to proof the age of the buyer it is considered an abstract norm.

3.2.1. Representing Norms

Before we will discuss the categorization of norms in terms of the ways they abstract from the actions and situations in the institution, we have to determine how the norms will be represented. The representation of the norms of course will in a large part determine the level of abstraction and the way they can be related to more concrete norms. E.g. if we represent all obligations by a classical (modal) deontic logic based on propositional logic then every state of affairs is collapsed into a proposition and relations between states of affairs might be difficult to represent.

For this paper we will assume that norms are expressed as either static or dynamic deontic logic formulas. The deontic logic is a temporal, relativized and conditional one. I.e. an obligation to perform an action or reach a state can be conditional on some state of affairs to hold, it is also meant for a certain type (or role) of agents and should be fulfilled before a certain point in time. We will not get into a complete formalization of the norms in this paper, but refer to [8] and [7] for the background of these choices.

The consequence of the above is that we assume that norms can be categorized in three types: obligations, prohibitions and permissions. The obligations and prohibitions are mainly seen, in this context, as restricting the possible actions (and subsequent situations) while the permissions are generally used to indicate the conditions in which, otherwise forbidden actions, can be performed.

Using this deontic logic we can formalize a norm such as: “the donor should consent to the transplantation before the transplantation can take place” through the following formula:

$$O_{\text{hospital}}(\text{consent}(\text{donor}) < \text{transplant}(\text{hospital}, \text{donor}, \text{recipient}, \text{organ}))$$

The obligation is directed towards the hospital, assuming that the hospital is responsible for fulfilling it. I.e. it is the responsibility of the hospital to acquire the consent of the donor before the transplantation is performed. In some sense this obligation has an implicit conditional. It only comes into effect if the hospital *intends* to perform a transplantation. We will not elaborate on this aspect, but we just assume that the intention of the agent is an implicit condition of the obligation if the action before which the obligation should be fulfilled is controlled by that agent.

3.2.2. Classifying abstract norms

Of course one can categorize norms in many different ways. In this section we will categorize the norms in the ways they are abstract. I.e. in the way that they generalize from or are more abstract than the concepts that are used to define the functioning of the institution. Subsequently we will indicate how a translation can be made between the abstract norms and more concrete forms. We will also indicate the relations that exist between the abstract and concrete norms (although we do not get to a complete formalization of this aspect).

The abstract norms try to capture many different situations and therefore are “abstract” in several different ways:

1. They are referring to an abstract action that can be implemented in many ways
2. They use terms that are vague and that have to be defined separately
3. They abstract from temporal aspects
4. They abstract from agents and or roles
5. They refer to actions or situations that are not (directly) controllable and/or checkable by the institution

The first category is perhaps the most easy one. E.g. “a living donor should consent to the donation of an organ”. The action of “consent” is might not be an action that is described in terms of the concepts used in the ontology of the institution and is therefore still abstract. It can be implemented by an action of the donor signing a contract or the donor telling it to his family or the donor carrying a codicil, etc. Each of which might be a valid action within the institution. So, the first case of abstraction refers to an abstraction of actions such that they refer to a set of possible ways of performing the action. Each of the possible ways of performing the abstract action (defined within the institution) should be either permitted, prohibited or obliged within the institution.

The translation in this case is a kind of definition of the abstract action in terms of the concrete actions. In the above case one could define this as follows:

$\text{Consent}(\text{donor}) \equiv \text{sign}(\text{donor}, \text{contract}) \cup \text{carry}(\text{donor}, \text{codicil}) \cup \text{tell}(\text{donor}, \text{family})$

i.e. “to consent” is a performing either of the three more specific actions. Important to note is that this definition closes the way consent can be given. In the above case it should be one of the three ways mentioned. One could also define it a bit more liberal by just giving a one way implication:

$\text{Consent}(\text{donor}) \subset \text{sign}(\text{donor}, \text{contract}) \cup \text{carry}(\text{donor}, \text{codicil}) \cup \text{tell}(\text{donor}, \text{family})$

Which means that doing one of the three specific actions “counts-as” a consent, but they are not the only ways to implement consent (see [12] for an account of the “counts-as” operator).

The second category of norms cannot be implemented directly because the norms use terms that have no precise meaning. E.g. “the extraction of organs from living donors should be limited to those situations where one can *expect a high probability of success* of the transplant”. This is a very common phenomenon in legal texts. “Vague terms” are often explicitly used to allow for interpretation within different contexts. See [11] for more discussion on “vague terms”. In this case it is not specified what it means to “expect a high probability of success”. There are two ways in which these terms can be explicated.

1. through a rule that defines the meaning of the term. E.g. “if the data from the donor's file and that of the recipient's file match for more than 90% and the recipient is otherwise in good condition and the donor's organ is healthy then one can expect *a high probability of success*”.
2. through a procedure that has to be followed in order to determine whether the condition referred to in the term is true. E.g. “One can *expect* a high probability of success if the medical committee of the recipient's hospital has consulted the data from the files of both donor and recipient and declared that the probability of success is higher than 90%”.

The first case actually replaces a condition that cannot be checked directly by the institution by a number of conditions that can be checked by (or within) the institution. In the second case there are no data available to be checked, but a procedure can be followed to determine the value of the condition. It also implies that the *only* (accepted) way to find out the value of the condition is through the given procedure.

Both cases have in common that one cannot proof directly whether the condition referred to in the norm is true or not. In order to proof that the condition is true one has to check some other conditions (either directly or through some prescribed procedure).

The next category of norms refers to norms that abstract away from the temporal elements. Often there is an implicit deadline for obligations, which is implied by the fact that the fulfillment of an obligation is also the fulfillment of a condition for a permission. E.g. the consent of the donor is a condition of the permission to perform a transplantation. Therefore the consent of the donor should be given before the transplant can be carried out. So, the obligation to get the consent does not have a deadline, but the rest of the procedure has to wait for this action. It might be that some deadlines are defined on other parts that indirectly limit the time to fulfill the obligation. E.g. if the whole procedure of the transplantation should be finished within one week (or otherwise the recipient might die) then the consent should be given within this week.

The fourth category of abstraction is about the role or agent for whom the norm holds. There is often an implicit assumption about who is responsible for an action or situation. Therefore the obligations, prohibitions and permissions are stated in general form. They are supposed to affect any agent that wants to perform a transaction through the institution and which is eligible to do so according to the regulations for admission of the institution. E.g. transplants can only be performed in hospitals (and probably heart transplants only in academic hospitals) etc.

The last category of norms refers to actions and situations that are not directly checkable. E.g. "the decision of who is the best recipient for an organ cannot be based on the age of the recipient". Although the norm is clear, it is impossible to check directly on which basis a decision is taken by an agent. This is an internal (mental) action. In order to check this norm, the institution has to devise some constraints and/or procedures that are checkable (or controllable) by the institution and which take care of the fulfillment of the norm. E.g. the institution might withhold all information pertaining to the age of the potential recipients to the decision makers. Alternatively they might let the decision makers sign a contract in which they promise not to let the age of the potential recipients influence their decisions.

The same type of argument holds for situations that are not directly checkable by the institution. E.g. "the donor should be of legal age". The fact is that in order for the institution to uphold this norm, it should be able to *know* whether the state described in it holds or not. Knowing in this context means that the institution has some kind of proof which it deems sufficient in order to claim it knows the state. So, the institution has to have a procedure or another fact which can count as evidence for the state of affairs. This is very similar to the situation in commerce. If the contract states that the buyer should pay as soon as the seller shipped the product then the question becomes how the buyer can make certain that the seller actually shipped the product before he pays. Also in this case there are a number of principles that can be applied to check whether the procedure that is followed actually assures that both parties have enough evidence (knowledge) to act upon. See [1] for more details.

Given the above classification we can define a formal language in which we translate the more abstract norms into the concrete ones (like was illustrated in the example above). However, even the concrete norms have no direct implementation in an institution. The obligation to sign a donor contract before a transplantation takes place can be formally represented as: $O(\text{signed}(\text{donor}, \text{contract}) < \text{transplant})$ i.e. in the situation a transplant is taking place it is obliged to have a contract signed by the donor. However, this obligation does not state how this signed contract comes about. In the next section we will spend a few words on this aspect.

4. From concrete norms to institutions

Unlike some concepts like concrete actions or descriptions of objects, norms cannot be translated into one specific concept that represents the norms within an institution. Theoretically, of course, one could represent norms by some data structure containing a description of the norm. However, this description itself would not influence the interactions within the institution, unless some entity takes the descriptions into account

and makes its actions in each situation dependent upon them. However, from the point of view of the institution we do not want to depend on this to happen. We would like the norms to have a pro-active effect. Therefore the norms are translated for some part into restrictions on behavior and for the other part into triggers on unwanted behavior of the agents interacting in the institution. From the point of view of the institution one would like to create a situation in which it is impossible for the agents to violate any of the specified norms, because in that situation the institution could guarantee that all the transactions are performed according to the objectives and values indicated in its statutes. Usually the actions that are taken by the institution in order for the values, objectives and norms to be fulfilled are stated in the *policies* of the institution.

For prohibitions this comes down to (trying to) make it impossible for the agents to perform an action that it is not allowed to perform. E.g. if a person is not registered at an auction he is not allowed to bid at the auction. The latter can be enforced by just preventing a person to enter the auction before he is registered.

Norms that are indicating that some behavior is wanted in certain situations (e.g. paying after buying something) are translated as much as possible as constraints on unwanted behavior. The difference is that the institution cannot actually force an agent to pay (this is the autonomous decision of the agent), but it can control the fact that the agent cannot leave before it has paid. By translating as many norms as possible into behavior that can be controlled by the institution, the institution can enforce as many norms as possible. See also [10] for this issue.

Finally, norms that indicate the permission of actions under certain conditions can be translated to checks on the conditions whenever an agent tries to perform the action. If the conditions are not fulfilled, the agent is prevented from performing the action (if possible).

However, not all norms can always be translated in this way. First of all, there might be actions that are not under the (direct) control of the institution. For instance, the bidders in an auction are not allowed to make agreements or communicate outside the auction about price arrangements. However, the institution cannot really enforce this prohibition. Secondly, trying to control all possible behaviors to enforce normative behavior can result in a very inefficient or cumbersome functioning of the institution. So, in some cases the institution might decide for the sake of flexibility and efficiency to not enforce a norm through restricting behavior but react on violations of the norms.

So, in all cases where the institution does not directly control the behavior of the agents and thus does not enforce the fulfillment of the norms it should have procedures to check violation of the norm and triggers to react on such a violation.

For each concrete norm we thus need a kind of rule specifying either the part of the procedures within the institution that enforce that norm or the events and triggers that signal violation of the norm and the reaction of the institution. For an example such as “before a transplant can commence the donor should give consent” the rules for the two possible implementations could look as follows:

1. $O(\text{signed}(\text{donor}, \text{contract}) | \text{transplant}) := \text{sign}(\text{donor}, \text{contract}) < \text{transplant}$ stating that the sign action is taking place before the transplant. The use of a temporal relation indicates that this is a property for all possible ways the procedures are implemented in the institution.

2. $O(\text{signed}(\text{donor}, \text{contract}) | \text{transplant})) := \{ \text{PRE}(\text{check}(\text{signed}(\text{donor}, \text{contract})), \text{transplant}), \neg \text{signed}(\text{donor}, \text{contract}) \rightarrow \text{request}(\text{donor}, \text{sign}(\text{donor}, \text{contract})) \}$

This states that as precondition of the transplant the institution checks whether the donor signed the contract. In case she did not sign, this triggers a request to sign the contract. (So, the last implication should be read as a trigger of an action).

Of course the above rules are still quite sloppy and contain many points that should be made more precise and clear. However, they serve here mainly as indication how such rules could be specified in a formal way. Further formalization and development of this specification is left for the full paper.

5. Conclusions

In this paper we have discussed the main issues around the implementation of norms in institutions. The main observation is that norms are specified in regulations that are (on purpose) on a high abstraction level. The level of abstraction is high on purpose in order not to be dependent on a circumstantial implementation of the norm. They should be stable for many situations and for a relative long time. Therefore it is obvious that the norms do not have concrete handles for their implementation.

In order for the norms to be implemented we have therefore to develop a formalism in which we can explicitly specify how the norms are translated into concrete norms applied for the application of the institution (at this moment). Using these concrete norms we can indicate how they are implemented in the institution. Also this latter step is not standard, but depends on efficiency issues and control issues.

The reason to describe each of these translation steps in a formal way is to be able (afterwards) to verify whether the institution still complies to all the norms that were specified in the original regulations. Only when the translation steps are described formally can one verify this compliance in a formal way as well.

The main areas for further research are the description of the formal language to describe the translations and specifically determining the exact semantic relations between the abstract norms and the implemented norms.

6. References

1. R. Bons, F. Dignum, R. Lee, and Y.-H. Tan. *A Formal Analysis of Auditing Principles for Electronic Trade Procedures*. In *International Journal of Electronic Commerce*, Vol.5, Nr.1, 2000, pages 57-82.
2. C. Castelfranchi, F. Dignum, C. Jonker and J. Treur. *Deliberate Normative Agents: Principles and Architectures*, In N. Jennings and Y. Lesperance (eds.) *Intelligent Agents VI (LNAI-1757)*, Springer-Verlag, 2000, pages 364-378.
3. R. Conte and C. Castelfranchi, *Cognitive and Social Action*, UCL Press, 1995.
4. R. Conte and F. Dignum. *From social monitoring to normative influence*, In *Journal of Artificial Societies and Social Simulation* vol. 4, no. 2, 2001.
5. F. Dignum. *Autonomous Agents with Norms*, In *Artificial Intelligence and Law*, Volume 7, 1999, pp. 69-79.

6. F. Dignum. *Agents, Markets, Institutions and Protocols*. In F. Dignum and C. Sierra (eds.) *Agent Mediated Electronic Commerce, The European Agentlink perspective*, LNCS-1991, Springer-Verlag, pages 98-114.
7. M.V. Dignum, J.-J. Ch. Meyer, H. Weigand and F. Dignum. *An Organisational-oriented Model for Agent Societies*, RASTA workshop, Bologna, July 2002.
8. F. Dignum and R. Kuiper. *Specifying Deadlines with Dense Time Using Deontic and Temporal Logic*, In *International Journal of Electronic Commerce*, Volume 3, Number 2, Winter 1998-99, pp. 67-86.
9. F. Dignum, D. Morley, E. Sonenberg and L. Cavedon. *Towards Socially Sophisticated BDI Agents*. In *Proceedings of the fourth International Conference on Multi-Agent Systems (ICMAS-2000)*, Boston, July 2000, IEEE Computer Society, pages 111-118.
10. M. Esteva, J. A. Rodriguez, C. Sierra, P. Garcia, J. L. Arcos. *On the formal specifications of electronic institutions*, In F. Dignum and C. Sierra (eds.) *Agent-mediated Electronic commerce (The European AgentLink Perspective)*, LNAI 1991, pp. 126-147.
11. T. Horgan and M. Potrc, *From Epistemicism to Transvaluationism: Perspectives on Vagueness*, Oxford University Press, forthcoming.
12. A. Jones and M. Sergot. *A Formal Characterisation of Institutionalised Power*. In *Journal of the IGPL*, vol. 4, no. 3, 1996, pp.427-443.
13. J.-J. Ch. Meyer and R. Wieringa (eds.), *Deontic logic in Computer Science*, John Wiley & Sons, 1993.
14. M. Miceli and C. Castelfranchi. *The Role of Evaluation in Cognition and Social Interaction*, In Dautenhahn K. (ed.), *Human Cognition and Social Agent Technology*, John Benjamins, Amsterdam, 1999.
15. H. Prakken and G. Sartor, *Argument-based logic programming with defeasible priorities*. *Journal of Applied Non-classical Logics* 7: 25-75 (1997), special issue on 'Handling inconsistency in knowledge systems'.
16. Von Wright, G.H. *On the Logic of Norms and Actions*. In Hilpinen R. (ed), *New Studies in Deontic Logic*, 3-35, D. Reidel Publishing Company, 1981.
17. <http://www.msc.es/ont/esp/ont.htm>