

**ITC**

**Institute for  
Communication  
Technologies**

[www.itc.com.unisi.ch](http://www.itc.com.unisi.ch)

Faculty of  
Communication  
Sciences

Università della  
Svizzera  
italiana

via Buffi 13  
CH-6900  
Lugano

**Technical report**

No. 2

**An Operational Approach to Norms in Artificial Institutions**

Francesco Viganò, Nicoletta Fornara, and Marco Colombetti

Institute for Communication Technologies  
Università della Svizzera Italiana, 2005.



# An Operational Approach to Norms in Artificial Institutions <sup>\*</sup>

Francesco Viganò<sup>1</sup>, Nicoletta Fornara<sup>1</sup>, and Marco Colombetti<sup>1,2</sup>

<sup>1</sup> Università della Svizzera italiana, via G. Buffi 13, 6900 Lugano, Switzerland  
{francesco.vigano, nicoletta.fornara,  
marco.colombetti}@lu.unisi.ch,

<sup>2</sup> Politecnico di Milano, piazza Leonardo Da Vinci 32, Milano, Italy  
marco.colombetti@polimi.it

**Abstract.** The notion of artificial institution is crucial for the specification of open and dynamic interaction frameworks where heterogeneous and autonomous agents can interact to face problems in various fields, like for instance electronic commerce, business-to-business (B2B) applications, and personal assistant applications. In our view the specification of artificial institutions requires a clear standard definition of some basic concepts: the notion of ontology, authorizations, conventions, and norms. In this paper we propose an operational approach to the definition of norms that is mainly based on the generation of commitments. These norms can be employed to verify if the interacting agents are behaving in accordance with the normative specification of the system. In particular we regard norms as event-driven rules that are fired by events happening in the system and then modify commitments affecting the agents having a certain role. We will discuss the crucial differences between the notion of authorization and permission and how the notion of permissions, obligations, and prohibitions can be expressed in our model. We will investigate the connections among the specification of different artificial institutions, in particular how an institution can enrich or further regulate the entities defined in another one. Finally we will briefly present the specification of the Dutch Auction Institution and of the Auction House Institution in order to exemplify the model presented in this paper.

## 1 Introduction

The definition and representation of *norms* and their actual specification is a fundamental component of the formalization of every open and dynamic interaction system [14, 21], that is, a system where heterogeneous and autonomous agents enter and leave

---

<sup>\*</sup> Supported by Swiss National Science Foundation project 200021-100260, "An Open Interaction Framework for Communicative Agents"

dynamically. In such systems norms play a crucial role because they: (i) regulate the behavior of agents, and (ii) create expectations on the behavior of other agents.

In open systems, norms can be analyzed and used from two different points of view: the design of autonomous agents and the design of interaction systems. From the first perspective norms have to be objectively expressed and external (with respect to the internal structure of agents). Moreover, the model used to represent norms has to enable agents to decide whether to fulfill them on the basis of their own goals [15, 13]. From the second perspective norms can be used to verify if agents are correctly behaving, even if, given that agents are autonomous, norms are not sufficient to prevent undesirable behavior. In case of violation of a norm, the authority regulating the interaction system may apply sanctions against the misbehaving agent [7, 21].

In the literature, it is possible to find numerous attempts to specify norms using deontic logic [5] or predicate logic [7, 13], but they have a limited expressiveness; for example they are not able to deal with time, a crucial aspect in normative systems. Finally there are some works, closer to our approach, where norms specify actions to be performed if certain preconditions are satisfied [14, 21].

In this paper we propose an operational approach to the specification of norms based on the notion of *social commitment*, that is, on a concept whose use in the specification of agent communication is becoming increasingly common [6, 8]. Thus an important advantage of our approach is its coherence with an existing specification of an Agent Communication Language (ACL). An agent able to reason on commitments would, as consequence, be able to reason both on the effects of communicative acts and on norms defined within a system. A crucial property of our approach is the possibility to verify if one or more agents have violated the norms of the system. That is, it is possible to keep trace of the commitments generated by the activation of norms for all the agents having some role in the system. By identifying the violations of commitments it is possible to detect the violation of the associated norm.

Finally, following our approach to norms it is possible for an artificial agent to reason on its own actions and to have expectations on the actions of other agents. In fact, it is possible to foresee which norms will be activated by the performance of an action. For example, an agent can decide whether to bid or not to bid in an auction, reasoning on the fact that its bid will activate a norm creating a commitment for the bidder to pay a given amount of money (see Section 5).

This paper is organized as follows. In Section 2 our view of the main components necessary for the specification of artificial institutions is presented. Among those components norms play a crucial role; our operational model of norms based on commitments is discussed in Section 3. In Section 4 the connections among the specifications of different artificial institutions are investigated and in Section 5 our model is clarified through an example. Finally in Section 6 we draw some conclusions and delineate some directions for future work.

## 2 Artificial Institutions

*Artificial institutions* are a software model, inspired by human institutions [18], used as technological extension of human society. They can be employed to perform, by

means of artificial agents, certain computational intensive tasks such as: establishing appointments, signing contracts, and carrying out commercial transactions for example by means of an auction. As will be also shown in Section 5 a concrete open interaction framework for heterogeneous agents is the result of the reification of the abstract specification of a set of artificial institutions.

Given that the system has to be *open*, that is, heterogeneous agents made by different designers may enter and leave the system dynamically, it makes crucial that the concepts used for its specification have to be clearly and unambiguously defined and commonly accepted as a standard by all agent designers.

In our view, the specification of an artificial institution consists of the following components [10]:

- the *core ontology*, that is, the definitions of the institutional concepts introduced by the institution and of the institutional actions that operate on them;
- a set of *authorizations* specifying which agents are authorized to perform the institutional actions;
- a set of *conventions* for the concrete performance of institutional actions;
- a set of *norms* that impose obligations, prohibitions and permissions for the agents that interact within the institution.

## 2.1 The core ontology

The context within which artificial agents operate can be modelled as consisting of a set of entities that can have both *natural* and *institutional attributes*, that is, attributes that exist only thanks to the common agreement of the interacting agents. For example, the size of a book is a natural attribute, while its price and its owner are institutional attributes.

Institutional attributes can be changed by particular types of actions: *institutional actions*. Therefore, agents cannot perform such actions by exploiting causal links occurring in the natural world, as it would be done to open a door. Rather institutional actions are performed on the basis of *conventions* on the exchange of messages.

We define institutional actions by specifying:

- an action name followed by a possibly empty list of parameters: *iaction(param)*;
- a possibly empty set of (ontological) *preconditions*, which specify the values that certain institutional attributes must have;
- a nonempty set of *postconditions*, which specify the values of certain institutional attributes after a successful performance of the action.

## 2.2 Authorizations and Conventions

Institutional actions are performed by exchanging messages. To specify what kind of message implements a specific institutional action we define conventions in the following form:

$$exchMsg(msg\_type, sender, receivers, content) =_{conv} iaction(param)$$

By itself, a convention is not sufficient to guarantee the successful performance of an institutional action by the exchange of the appropriate message: indeed, some additional conditions must be satisfied.

**Conditions on the sender of the message.** In general, an agent must be authorized to perform an institutional action; for example, only the auctioneer can open an auction by sending a suitable message to the participants. Moreover an authorization typically holds if certain conditions about the state of the system, expressed by suitable Boolean expressions, are satisfied. For example, it may be established that an auction is validly opened only if there are at least two participants. Assuming that every agent in the interaction system has an identifier (*agent\_id*), authorizations will be represented with the following notation:

$$\text{Auth}(\text{agent\_id}, \text{iaction}(\text{param}), \text{conditions})$$

In the specification of an interaction system it is useful to express authorizations in term of the roles filled by agents, in order to abstract from the concrete agents that are actually involved in an interaction. For example, the authorization to open and close an auction is granted to the agent that fills the role of the auctioneer, independently of its individual identity. Usually roles are defined relative to an institutional entity, for example the role of participants and auctioneer are defined relative to the auction entity. We can then abstractly define the authorization to perform a specific institutional action (with given parameters) associating it to a role defined in the context of a specific institutional entity (*ientity*):

$$\text{Auth}(\text{ientity.role}, \text{iaction}(\text{param}), \text{conditions})$$

In a concrete interaction, the authorizations associated to roles need to be transformed into authorizations of an actual agent in the system. Such transformation can be obtained searching among all the institutional entities in the system the ones that match the description given through the parameters of the institutional action, and then creating a concrete authorization for each agent having the role indicated in the abstract authorization.

**Conditions about the receivers of the message.** Messages realizing institutional actions should be received by all agents that are affected by the performance of the act. For example, in the case of an institutional action of auction opening, all the participants have to be the receivers of the message.

**Conditions about the state of the system.** In general, the exchange of a conventional message successfully performs the associated institutional action only if certain conditions on the state of the system are satisfied.

### 3 Norms

The execution of institutional actions by an authorized agent often needs to be regulated by a system of norms imposing which actions should or should not be performed. When an agent is planning its following action, it has to consider whether to satisfy the constraints imposed by norms defined in the interaction systems by different institutions. For example, in our formalization of the Dutch Auction the auctioneer not only is authorized to declare a price for the good on sale, but it is also obliged to declare a

price lower than the previous one or to close the auction. At the same time, a norm of the Auction House forbids to the auctioneer to declare a price lower than a *reservation price*.

Norms allow us to describe the expected behavior of the interacting agent from an external point of view. In doing so, norms play an important role in a multi-agent system, in that they make an agent's behavior at least partially predictable and allow agents to coordinate and plan their actions according to the expected behavior of the others, as studied in [16, 1]. In our framework such desired behavior is expressed through a set of commitments that represent the obligations and prohibitions of agents created or cancelled by norms. Before presenting our formalization of norms, we need to introduce the concepts of commitment and how we describe events occurring in a multi-agent system.

### 3.1 Commitments

In this paper we give only a short description of our model of commitment, which is assumed to be the fundamental entity of what we call the *Basic Institution*. For a complete treatment see [10, 4].

We regard a commitment as an entity with the following attributes: a *debtor*, a *creditor*, a *content*, and a *state*, used to keep track of the temporal evolution of the commitment. Commitments are represented with the following notation:

$$Comm(state, debtor, creditor, content)$$

The content of a commitment is represented by means of a *temporal proposition* (for a detailed treatment see [9, 4]). At every time instant, a temporal proposition has a truth value, which can be *undefined*, *true*, or *false*. When the content of a commitment is no longer *undefined*, as a consequence of the occurrence of a domain event, the state of that commitment is automatically set to *fulfilled* if the content has become *true*, otherwise it is set to *violated*.

In our framework every agent is authorized to create a commitment by performing the *makeCommitment* institutional action, whose successful performance creates an *unset* commitment. The debtor of an *unset* commitment may refuse it by executing *setCancel*, or it may undertake the proposed commitment by executing *setPending*. We represent a refused commitment by means of the *cancelled* state, whereas an accepted commitment is depicted with the *pending* state. The creditor of *pending* or *unset* commitment can always set it to *cancelled*. Another institutional action, used in the example in Section 5, is *makePendingComm*, which creates a *pending* commitment and whose execution coincides with the sequential execution of *makeCommitment* and *setPending*.

Summarizing, commitments undergo the life cycle described in [8] by reacting either to institutional actions performed by agents or to domain-dependent events, which modify the state of temporal propositions.

### 3.2 Events

As we will see, norms are event-driven rules that when are fired by events happening in the system modify commitments affecting the agents having a certain role. The Unified

Modeling Language (UML) [2] models four kinds of events: *signals*, *calls*, *passing of time* and *change in state*. Inspired by UML notation for signals, here we propose to model type of events as stereotyped classes.

We represent events as entities of the system having attributes that provide information about the state transition that caused it. Moreover we assume that the system returns the time at which an event token has occurred, by means of the *time\_of()* operator. In our formalization we have singled out three main categories of events:

- a *TimeEvent*, that occurs when the system reaches a certain *instant\_of\_time*;
- a *ChangeEvent*, that happens when an institutional entity changes in some way. This kind of event type can be specialized further:
  - an *InstitutionalPropertyChange*, registered when an attribute has changed its value;
  - an *InstitutionalRelationChange*, occurring when a new relation is *created* or an existing one between the institutional entity and another one is *dropped*;
  - an *InstitutionalStateChange*, occurring when an entity modifies its type in a given taxonomy (e.g., when an auction from *unset* becomes *open*);
- an *ActionEvent*, that happens when an agent perform an action (an interesting type of this kind of events is *ExchangeMessage*, that represents the act of sending a message).

The definition of event types allows us to describe *event templates*, that are, event types with some restrictions on certain attributes that describe a set of possible event occurrences. Event templates are used in the *on* section of a norm: when an event matches the given descriptor, the corresponding norm is fired and its variable *e* is filled with the event that activated it.

Such a description can be straightforwardly transformed in a specification based on *preconditions* and *post-conditions*, that is, a description of the state of the system before and after the event has happened. The possibility to specify events through preconditions and post-conditions is crucial for artificial agents to reason on whether the performance of a certain action makes an event happen. Therefore, an agent is able to infer which norms will be activated by performing an action, which is an important property of a system of norms.

### 3.3 The definition of norms

We regard norms as event-driven rules that create or cancel commitments affecting a set of agents that enact a specified role within the institution. At an abstract level, a norm is part of the definition of an artificial institution; its instances then regulate and are bounded by the organization that reifies the institution. When an agent fills a role in an institution, we assume that it accepts that norms create commitments binding the agent to a pseudo-agent representing the institution, which we call an *institutional agent*. Such agent allows us to keep trace of commitments created by a certain instance of institution, which also means that commitments created by norm of an institution can be canceled only by norms defined by the same institution; this is because only the creditor of a pending commitment can set it to cancelled [10]. Furthermore, if two

commitments are in conflict, an agent can decide which is more important with respect to its own policy (see [13]) by reasoning about which institutional agents have created such commitments.

A norm is defined within an institution, observes an entity of an institution and is activated by an event concerning such an entity. Typically, interesting event types are not only communicative acts like in [21], but also the filling of a role by an agent, a value change of an institutional attribute, the reaching of certain instant of time, and so on.

When a norm is activated, if certain contextual conditions are met, it is applied to a collection of *liable agents*, which are described by a suitable selection expression; in general, the collection of liable agents corresponds to the set of agents that play a given role in the institution. For every liable agent, the norm creates or cancels a set of commitments of the agent toward the institutional agent.

The general structure of a norm can be described as follows:

```
within context_name: ientity  
on e: event_type  
if contextual conditions then  
  foreach agent in liable agent selection expression  
  do {commitmentActionDescription(agent, inst_agent, parameters)}+
```

In an open multi-agent system by themselves norms are not able to banish violations, because the sincerity and benevolence of agents are not guaranteed. This is because we do not assume that norms are constraints encoded in each agent as in [16] and [1]. In fact, in our opinion agent designers can program their agents to violate norms if they consider this as a means for the fulfillment of their goals. In open systems norms should be used to indicate preferred path for the evolution of the system and to detect violation of commitments created by norms. In this respect, our point of view is close to [15] and [21], where no assumptions are made about the internal structure of agents.

Using norms, institutions can regulate in an uniform way both the communication protocol and protocol-independent normative aspects, like for instance the fulfillment of agreements made during the interaction. Norms can be used to specify protocols, because they can dictate that in certain circumstances an agent ought to send a given type of message, or react to a message in a specific way. At the same time, norms can forbid the execution of institutional actions, in particular communicative acts, even if they are authorized. Furthermore, in correspondence of events that conclude the interaction process, norms can instantiate commitments to noncommunicative actions, like the payment of the purchased goods at the price negotiated during the interaction. Examples of these types of norms will be discussed in Section 5.

Many studies have been devoted to the analysis of the relationship holding between norms and commitments. For example in [14] commitments are defined as a specialization of norms, while in [3] and [19] norms are a special kind of commitments, called *metacommitments*.

From our point of view, norms are not themselves commitments, but rules that manipulate commitments of the agents engaged in an interaction. In fact, norms are associated to roles rather than to individual agents, and strictly speaking they cannot be

fulfilled or violated. Indeed, what can be fulfilled or violated is not a norm, but rather a commitment created by the application of a norm.

By creating a new commitment whenever a norm is applied, we can compute how many times a norm has been violated or fulfilled by counting how many commitments instantiated by that norm are violated or fulfilled. This is important because we consider that a normative system should allow one to detect not only the presence of violations, but also how often they occur.

### 3.4 Obligations, Permissions, Prohibitions and Authorizations

In our framework commitments are used to represent all deontic relationships between agents, including as a special case the deontic relationships undertaken by the debtor through communicative acts [10]. In particular, commitments toward institutional agents are used to represent *obligations* and *prohibitions*. In general, we perceive obligations and prohibitions as commitments undertaken by an agent enacting a role within an institution, toward the institution itself; more precisely, obligations are commitments to perform an action of a given type, and prohibitions are commitments not to perform an action of a given type. Furthermore, we interpret the absence of positive or negative commitments to the execution of an action of a given type as *permissions*.

Usually in agent literature the concept of authorization is not distinguished from permission or the former encompasses the latter [5]. Coherently with the concept of institutionalized power of [12], we distinguish between the notions of authorization and permission. The main difference between authorization and permission resides in the effects of the action. Whereas the former represents a necessary condition for the execution of institutional action, permission represents the need to regulate the performance of authorized actions, but it cannot prevent the effects deriving from the performance of a forbidden act.

Summarizing, an agent's commitments can be of two types: those undertaken by the agent through the performance of communicative acts (see [10]) and those imposed to the agent by norms in virtue of the agent's roles in an institution.

## 4 Connections among different Artificial Institutions

In this section we discuss what relations can exist among different artificial institutions constituting an interaction system. We envision that when a designer starts to specify a new artificial institution, there is at his or her disposal a library of previously defined institutions that can be used to generate a new one. Therefore, interrelations among artificial institutions are crucial to define a modular and incremental specification of new institutions.

In general, an artificial institution can enrich with new institutional attributes the description of entities defined by ontologies pertaining to other institutions. For instance, as we will see in the example reported in Section 5, the Dutch Auction defines a *current price* relative to a product on sale during an auction, while, at the same time, the Auction House associates to the same product a reservation price agreed with the owner of the product. Given that institutional facts exist only thanks to common agreement,

attributes described in different institutions exist only within the scope of the defining institution. When another institution enriches an ontology with new attributes, they do not modify the previous specification. Because agents involved by the former institution might not participate in the second, and thus might not access such attributes, new attributes will be recognized only by those agents filling a role in the new institution.

A second interesting relationship existing among two different institutions regards norms. Whereas a norm can regulate institutional actions described by any institution, its set of liable agents can be constituted only by agents enacting a role in the same institution that defines the norm. Given that agents can fill roles pertaining to several institutions, a norm can affect agent behavior within other institution by creating obligations and prohibitions for a role of that institution to execute an action defined by another institution.

We perceive the distinction between authorization and permission as fundamental, because authorizations and permissions (i.e., the absence of prohibitions) might be relative to different sources. In fact, even if an action is authorized by an institution, the performance of such an action might be forbidden to an authorized agent by another institution. In this case, the agent has the necessary institutional power, but it is not permitted to perform that action by the second institution. If it executes such action, its effects take place, but the agent violates a commitment.

The interaction of norms might raise conflicts between obligations and prohibitions imposed by different institutions. At runtime, when an agent has contrasting obligations and prohibitions, it can reason about which is less expensive to violate, usually considering from what institution each commitment stems from. More research is needed to develop methods able to ensure that the specifications defined by different institutions do not raise conflicts. Further research is also needed to provide a general method for an institution to enlarge the set of authorizations defined by another institution.

## 5 An example: the Dutch auction

In this section we will present two examples of *artificial institutions* in order to exemplify how norms are specified in our framework. This allows us to show how different institutions interact and to clarify the distinction between authorizations and permissions. To this purpose, we shall present our formalizations of the Dutch Auction and of the Auction House holding auctions regulated by the former institution. Due to space limitations, we report a simplified version of our specifications, focusing our attention only on those aspects involved in fixing the price of the product on sale.

The main advantage of our formalization with respect to the one specified by FIPA [11] is that, due to the explicit representation of norms as rules that modify agent commitments, it is possible to model in an uniform way the interaction protocols and the other rules that regulate the interaction framework. Furthermore, when an interaction terminates successfully, agents are explicitly committed by suitable norms to carry out the economic transaction.

## 5.1 The Dutch Auction Institution

During a Dutch Auction an auctioneer tries to fix the price of goods. For the sake of simplicity, here we assume that goods cannot be split. An agent taking part in a Dutch Auction can fill the role of *participant*, *auctioneer* or *transaction agent*, the agent that attends to the exchange of money and goods when a price has been accepted. During an auction we assume that a participant cannot be the auctioneer or the transaction agent, while an auctioneer might also fill the role of transaction agent.

After a period of time reserved to the registration of participants, the interaction starts when the *start time* has elapsed and the auctioneer has declared the auction open. Then, the price of the product on sale is initialized, usually higher than the expected final result. We assume that when a new price is declared, during the validity of such price, the auctioneer declares as the *winner* the first participant that accepts the *current price*, and then closes the auction. Otherwise, after the time of validity has elapsed, the auctioneer should declare a new *current price*, lower than the previous one, or close the auction.

The ontology of the Dutch Auction is described by the class diagram reported in Figure 1, where concepts like Agent and Product are assumed to be defined by external ontologies. It is important to observe that the *current price* is defined relative to the Product, which means that the ontology enriches the definition of such an entity by adding a new institutional attribute recognized by the agents involved in the current interaction.

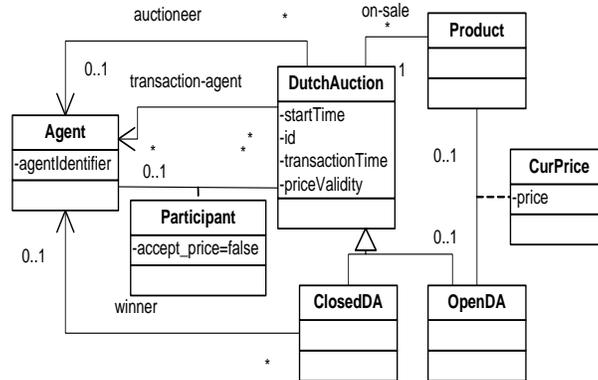


Fig. 1. The Dutch Auction Ontology.

The Dutch Auction ontology also defines a set of institutional actions that allow agents to operate on such institutional entities. For the sake of brevity, we report only the complete specification of the institutional action that sets the current price, where preconditions and effects are described using Object Constraint Language (OCL)[17]. We assume that the price of a Dutch Auction can be changed only if the auction is open

and the previous price is higher than the new one, which becomes the current price of the auction.

*name* : **setCurPrice**(*a\_id*, *p*)  
*pre* : *OpenDA.allInstances*  $\rightarrow$  *exists*(*id* = *a\_id* and *currentPrice.price* > *p*)  
*post* : *OpenDA.allInstances*  $\rightarrow$  *exists*(*id* = *a\_id* and *currentPrice.price* = *p*)

The Dutch Auction Institution defines a set of authorizations for the performance of institutional actions. Some of these authorizations are conditional: for example an auctioneer is authorized to open an auction only if its *start time* has elapsed and if there are at least two agents registered as participants. Here we report only the authorization that allows the auctioneer of a given auction to perform the *setCurPrice* institutional action.

*Auth*(*DutchAuction*<sub>*id*</sub>.*auctioneer*, *setCurPrice*(*id*, *p*))

The behavior of agents that have joined an interaction system regulated by the Dutch Auction Institution is constrained by a normative system, which prescribes what agents should or should not do in correspondence to relevant institutional events. Likewise [10], we have defined a set of norms that regulate both the communicative acts performed by agents and the final exchange of good and money between the *transaction agent* and the *winner* of the auction. Here we report an example of a norm that is activated when an agent is declared to be the winner and creates a commitment for such an agent to pay a given amount of money to the *transaction agent*. Let us denote the institutional agent by the *DutInstAgent* identifier. Then, the formalization of such a norm is as follows.

**within** *a*: *OpenDA*  
**on** *e*: *InstitutionalRelationChange*(*a*, *winner*, **created**)  
**if true then**  
  **foreach** *agent* **in** *a.winner*  
  **do**  
    **makePendingComm**(*agent*, *DutInstAgent*,  
    (*give*(*agent*, *a.transaction\_agent*, *a.curPrice.price*),  
    < *now*, *now* + *a.transactionTime* >,  $\exists$ )

A similar norm commits the *transaction agent*, toward the Institution, to transfer the ownership of the product on sale to the winner of the auction.

## 5.2 The Auction House Institution

The Auction House is an institution constituted by a set of *employee* agents, holding several auctions and regulating the commercial relation with the *owner* of the product on sale. The Auction House defines only one role, *employee*, that allows agents to fill both the roles of auctioneer and transaction agent, while the *owner* role, together

with Product, is imported by external ontologies. In Figure 2 we report a class diagram representing the ontology defined by the Auction House.

In order to obtain simpler and shorter OCL expressions, in this paper we will assume that an Auction House runs at most one Dutch Auction. When an agent decides to sell a product through an auction, it reaches an *agreement* with the Auction House concerning the minimum price at which the product may be sold. Such an institutional fact does not require the agreement of the participants: in fact, participants are not even assumed to know about the existence of a reservation price (not to mention its actual value current value). We regard reservation price as an institutional attribute associated to the product and representing a private agreement established between the Auction House and the *owner* of the product on sale.

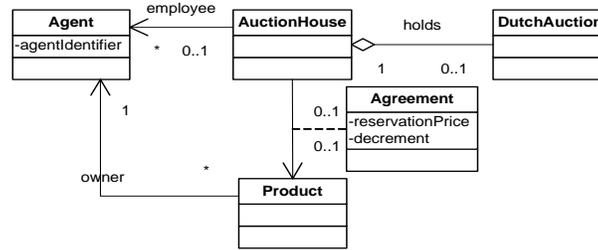


Fig. 2. The Auction House Ontology.

A norm of the Auction House is related to the *agreement* stipulated between the owner of the good and the auction house and is activated when an employee becomes the auctioneer. This norm commits the auctioneer to not declare a price lower than agreed reservation price.

```

within ah: AuctionHouse
on e: InstitutionalRelationChange(ah.DutchAuction, auctioneer, created)
if true then
  foreach agent in ah.employee → select(
    emp | e.involved → contains(emp))
  do
    makePendingComm(agent, DutInstAgent(not setCurPrice(a.id, ?p)
      [?p > ah.agreement.reservationPrice], < now, now + time_of(e1 :
        InstitutionalStateChange(ah.dutchAuction, OpenDA, ClosedDA)) >, ∀))
  
```

This norm constrains the auctioneer's behavior by imposing that the current price cannot be lower than the reservation price, although any price would be legal from the point of view of the Dutch Auction. In fact, the Dutch Auction authorizes the auctioneer to set a new current price, imposing through the ontological preconditions of the *setCurPrice* that it should be lower than the previous one, but not further constrains are

imposed. This example shows clearly that authorizations and permissions may differ when they are relative to different sources.

## 6 Conclusions

In this paper we have presented a model for the specification of artificial institutions. In particular, we have focused on an operational definition of norms as event-driven rules that modify agent commitments. In doing so, we have discussed the fundamental role played by norms, which allow us to express obligations and prohibitions in terms of commitments. Indeed, norms can represent in a unified way both interaction protocols and other normative aspects. Finally we have shown, through an example, how an interaction system can be specified in terms of norms defined by different artificial institutions. In particular, we have discussed how a designer may define a new interaction framework by using several artificial institutions and what connections might exist between them.

Several research questions are still open, and will be tackled in our future work. In particular, we are interested in developing methods for discovering inconsistencies among different artificial institutions. This point is particularly useful to develop tools supporting the generation of new institutions by reusing previously defined institutions. In particular, we are interested in verifying during the specification phase whether norms may create obligations to perform unauthorized actions, or under what conditions two norms may generate conflicting commitments. Finally, we plan to devise an explicit representation of the sanctions connected to the violation of a commitment.

## References

1. M. Barbuceanu, T. Gray, and S. Mankovski. Coordinating with obligations. In K. P. Sycara and M. Wooldridge, editors, *Proceedings of the 2nd International Conference on Autonomous Agents (Agents'98)*, pages 62–69, New York, 1998. ACM Press.
2. G. Booch, J. Rumbaugh, and I. Jacobson. *The Unified Modeling Language User Guide*. Addison-Wesley, Reading, Massachusetts, USA, 1 edition, 1999.
3. C. Castelfranchi. Commitments: From Individual Intentions to Groups and Organizations. In V. Lesser, editor, *Proc. First International Conference on Multi-Agent Systems*, pages 528–535, San Francisco, USA, 1995. AAAI-Press and MIT Press.
4. M. Colombetti, N. Fornara, and M. Verdicchio. A Social Approach to Communication in Multiagent Systems. In J. Leite et al., editor, *Declarative Agent Languages and Technologies*, volume 2990 of *LNAI*, pages 121–150. Springer, 2004.
5. F. Dignum. Autonomous agents with norms. *Artificial Intelligence and Law*, 7(1):69–79, 1999.
6. F. Dignum, editor. *Advances in Agent Communication: International Workshop on Agent Communication Languages, ACL 2003, Melbourne, Australia, July 14, 2003.*, volume 2922 of *LNAI*. Springer Verlag, 2004.
7. M. Esteva, J. A. Rodríguez-Aguilar, C. Sierra, P. Garcia, and J. L. Arcos. On the Formal Specification of Electronic Institutions. In F. Dignum and C. Sierra, editors, *Agent Mediated Electronic Commerce, The European AgentLink Perspective*, volume 1991 of *LNAI*, pages 126–147. Springer, 2001.

8. N. Fornara and M. Colombetti. A Commitment-Based Approach to Agent Communication. *Applied Artificial Intelligence an International Journal*, 18(9-10):853–866, 2004.
9. N. Fornara and M. Colombetti. Protocols Specification using a Commitment-based ACL. In F. Dignum, editor, *Advances in Agent Communication*, volume 2922 of *LNAI*, pages 108–127. Springer, 2004. (Extended version of N. Fornara and M. Colombetti. Defining interaction protocols using a commitment-based agent communication language. In Proceedings 2nd International Joint Conference on Autonomous Agents and MultiAgent Systems (AA-MAS 2003), pages 520-527. ACM Press, 2003.).
10. N. Fornara, F. Viganò, and M. Colombetti. Agent Communication and Institutional Reality. In van Eijk et al. [20], pages 1–17.
11. Foundation for Intelligent Physical Agents. FIPA Dutch Auction Interaction Protocol Specification. <http://www.fipa.org>, 2001.
12. A. Jones and M. J. Sergot. A formal characterisation of institutionalised power. *Journal of the IGPL*, 4(3):429–445, 1996.
13. L. Kagal and T. Finin. Modeling conversation policies using permissions and obligations. In van Eijk et al. [20], pages 123–133.
14. F. López y López and M. Luck. Modelling Norms for Autonomous Agents. In E. Chavez, J. Favela, M. Mejia, and A. Oliart, editors, *Proceedings of the 4th Mexican International Conference on Computer Science*, pages 238–245, 2003.
15. F. López y López, M. Luck, and M. d’Inverno. Normative Agent Reasoning in Dynamic Societies. In N. R. Jennings, C. Sierra, L. Sonenberg, and M. Tambe, editors, *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004)*, pages 535–542, New York, USA, 2004. ACM Press.
16. Y. Moses and M. Tennenholtz. Artificial social systems. *Computers and AI*, 14(6):533–562, 1995.
17. Object Management Group. UML 2.0 OCL Specification. <http://www.omg.org/>, 2003.
18. J. R. Searle. *The construction of social reality*. Free Press, New York, 1995.
19. M. P. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7:97–113, 1999.
20. R. van Eijk, M. Huget, and F. Dignum, editors. *Developments in Agent Communication*, volume 3396 of *LNAI*. Springer Verlag, 2005.
21. W. W. Vasconcelos. Norm Verification and Analysis of Electronic Institutions. Workshop on Declarative Agent Languages and Technologies (DALT), 3rd International Joint Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2004), 2004.

Technical report No. 1

**Agent Communication and Institutional Reality**

N. Fornara, F. Viganò, M. Colombetti

2004

**ITC**

**Institute for  
Communication  
Technologies**

[www.itc.com.unisi.ch](http://www.itc.com.unisi.ch)

Faculty of  
Communication  
Sciences

Università della  
Svizzera  
italiana

via Buffi 13  
CH-6900  
Lugano