A Communicative Act Library in the context of Artificial Institutions*

Nicoletta Fornara[1], Francesco Viganò[1], Marco Colombetti[1,2]

[1]Università della Svizzera italiana, via G. Buffi 13, 6900 Lugano, Switzerland
{nicoletta.fornara, francesco.vigano, marco.colombetti}@lu.unisi.ch,
[2]Politecnico di Milano, piazza Leonardo Da Vinci 32, 20135 Milano, Italy
marco.colombetti@polimi.it

Abstract. In this paper we propose a model of agents' institutional reality that makes us able to define the semantics of a set of communicative acts suitable for agent communication. Taking inspiration from Searle's studies on the construction of social reality, we introduce our notion of Artificial Institution consisting of: an ontology of the concepts defined by the institution, a set of authorizations for the performance of actions on institutional reality, a set of conventions for the performance of such actions, and a set of norms. In particular we define the Basic Institution that introduces all the basic concepts necessary to define communicative acts in terms of their effects on commitments, and a Communicative Act Library whose syntax is compatible with the one proposed for FIPA-ACL. Finally we introduce and discusse our proposal for the formalization of norms.

1 Introduction

The main goal of this paper is to propose a model of institutional reality that makes us able to define a set of communicative acts suitable for agent communication. Following Searle, we assume that by means of performing communicative acts one can affect the institutional reality shared by the sender and the receivers, that is, the set of entities that exist only thanks to the common agreement of the interacting agents. As will be explained in details in Section 2, two fundamental parts of institutional reality are *institutional entities* and *institutional actions*, that is, actions performed within an institution to modify a fragment of social reality. In our framework, communicative acts are regarded as a particular type of institutional actions.

Our approach to the definition of communicative acts is based on social commitment. In the last few years this concept has been used by a growing number of researchers to define the semantics of Agent Communication Languages (ACLs). After the first studies carried out by Singh and by Colombetti [Singh, 1999a, Colombetti, 2000], further investigations have been carried out from an operational point of view [Fornara and Colombetti, 2002; McBurney and Parsons, 2003], following a logical approach [Verdicchio and Colombetti, 2003], and in the field of argumentation studies [Amgoud et al., 2002; Bentahar et al., 2004]. The main advantages of this approach are that commitments are objective and independent of an agent's internal structure, and that it is possible to verify whether an agent is behaving according to the given semantics.

Part of the work presented in this paper is a revised version of [Fornara *et al*., 2004]. The paper is organized as follows. In Section 2 we introduce our concept of artificial institutions and of their main components: institutional entities, institutional actions, and "counts as" relationships between instrumental and institutional actions. In Section 3 we introduce the Basic Institution, which regulates the management of social commitments and allows for the definition of communicative acts as institutional actions. In Section 4 we delineate a simple communicative acts library. In Section 5 we consider a further component of artificial institution, namely norms, that play an important role in the definition of special institutions. Finally in Section 6 we draw some conclusions and delineate some directions for future research.

2 Artificial Institutions

We view a multiagent system (MAS) as a technological extension of human society, by which single persons and human organizations can delegate the execution of institutional actions to the artificial

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

system. Examples of such actions are establishing appointments, signing contracts, and carrying out commercial transactions. For this reason there are strong connections between some aspects of a MAS and some aspects of human society, and therefore the concepts used to model a MAS interaction framework have to reflect some crucial characteristics of their human counterpart. Therefore in creating our model we draw inspiration from philosophical studies about human communication.

Institutional actions can be performed only within *institutions*. In our view, the specification of an institution consists of the following components:

- 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* (see Section 5) that impose obligations and permissions on the agents that interact within the institution.

2.1 The core ontology: entities and attributes

The context within which artificial agents operate can be modeled 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 (or more precisely of their users). For example, the color of a book is a natural attribute, while the book's price and its owner are institutional attributes. Natural attributes are assumed to reflect the physical properties of the corresponding entities of the real world, and typically cannot be changed by artificial agents (unless the agent controls a physical robot). On the contrary, institutional attributes can be affected by institutional actions performed by purely software agents.

As will be described in Section 3, *social commitments* are the fundamental institutional entities because they are essential to express the meaning of various communicative acts.

2.2 The core ontology: institutional actions

Institutional actions are particular types of actions [Colombetti and Verdicchio, 2002] that are crucial for the formalization of communicative interactions taking place in open interaction frameworks. The effect of institutional actions is to change institutional attributes, which exist only thanks to common agreement. Therefore, agents cannot perform such actions by exploiting causal links occurring in the natural world, as it would be done to open a door or to remove a physical object. Rather, as we shall see, institutional actions are performed on the basis of a particular construct: the *counts as* relation.

Because of their intrinsic social nature, a crucial condition for the actual performance of institutional actions is that they must be public, that is, made known to the relevant agents by means of some action that can be directly executed by an artificial agent. It is therefore natural to assume that all institutional actions in a multiagent system are performed by sending suitable messages to the relevant agents as will be discussed in next sections.

We define institutional actions by specifying their preconditions and postconditions, therefore abstracting from the way in which such actions are concretely carried out. More precisely, an institutional action, that will be generically represented as *iaction(parameters)*, is characterized by:

- an *action name* followed by a possibly empty list of parameters;
- a possibly empty set of (ontological) *preconditions*, that specify the values that certain institutional attributes must have for the action to be meaningful (for example, the institutional action of opening an auction is meaningful only if the auction is not already open);
- a nonempty set of *postconditions*, that specify the values of certain institutional attributes after a successful performance of the action.

As we will see later in Section 4, communicative acts are a particular type of institutional actions.

2.3 Instrumental actions

As we have already remarked, an institutional action is performed by executing an instrumental action, conventionally associated to the institutional action. In the human world such instrumental actions vary from certain bodily movements (raising one's arm to vote), to the use of specific physical tools (waving a white flag to surrender), to the use of language (say "the auction is open" to open an auction). In a system of artificial agents, it is natural to assume that all institutional actions are performed by means of a single type of instrumental actions, namely exchanging messages.

For the purposes of the current treatment, a message consists of: a *message type*, a *sender*, one or more *receivers*, and a *content*. The action of exchanging a message will be represented with the following notation:

exchMsg*(message type, sender, receivers, content)*

Note that here sender and receivers are just fields of a message. That such fields correctly represent the agent that actually sends the message and the agents to which the message is delivered has to be guaranteed by the underlying message transport system.

2.4 The "counts as" relation

Which is the relation that binds the performance of an instrumental action to the performance of an institutional action? Following Searle [Searle, 1995], the construction of social reality in the human world is possible thanks to constitutive rules of the form "*X counts as Y in C"*; in the particular case where *X* and *Y* are actions, performing an action of type *X* in context *C* counts as performing an action of type *Y* . Similarly in artificial systems the "counts as" relation can be used to bind the performance of a message exchange to the corresponding institutional action (in particular a communicative act) if certain contextual conditions are satisfied.

Conventions

In order to be able to model the connection between *X* and *Y* we introduce the notion of *convention*, that is, an agreement about which type of message is bound to a given type of institutional action. In Section 4 some useful conventions will be defined for the performance of institutional actions. In our model the definition of a convention has the following generic form:

exchMsg*(message type, sender, receivers, content)* $=_{conv}$ *iaction(parameters)*

Contextual conditions: C

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 about the agent that sends the message, about the agents that receive the message, and about the state of the system in relation to the content of the message 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 can be given only 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:

Auth(agent_id, iaction(parameters), conditions)

Our notion of authorization should not be confused with the notion of permission. The distinction we make between these two concepts is similar to the one between institutionalized power and permission proposed by Jones and Sergot in [Jones and Sergot, 1996]. While authorizations are necessary conditions

for the valid performance of institutional actions, permissions (like obligations) are brought about by norms (see Section 5), that is, by rules that affect the normative positions of the agents in the system. The crucial difference between authorizations and permissions is highlighted in the cases when they are not granted. If an agent is not authorized to perform an institutional action, a performance of the corresponding instrumental action does not count as a performance of the institutional action (the institutional action is thus *not executed*). On the contrary, if an authorized agent performs an institutional action without permission, the institutional action is successfully performed, but the agent violates a norm and may be sanctioned for its behavior.

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. The concept of a role is very broad: for example, it is possible to regard social commitments as institutional entities that define two roles: the *debtor* of the commitment and its *creditor*. This fact appears to be general; that is, roles are defined relative to an institutional 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*):

Auth(ientity.role, iaction(parameters), 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. All agents that play a role in the institutional reality affected by the performance of the act. For example if the institutional action operates on a commitment its debtor and/or its creditor (depending on who is the actor of the action) have to receive the message; in the case that the institutional action open an auction, the participants to the auction have to be the receivers of the message.

Conditions about the state of the system. All the preconditions of the institutional action associated to the performance of the exchange of the message must be satisfied.

Figure 1 schematically depicts how the “counts as” relation works: an actual exchange of a message counts as the performance of a particular institutional action if there is a convention that binds the two acts and the relevant contextual conditions are satisfied.

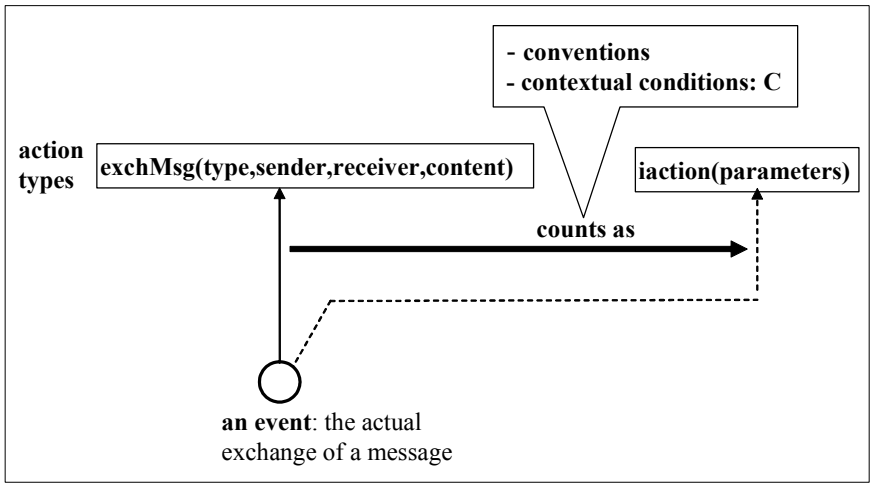

Fig. 1. The "counts as" relation.

3 The Basic Institution

The concepts defined so far are our tools for describing the institutions that characterize an open interaction framework. We start from what we call *Basic Institution*, which is necessary to define the semantics of an Agent Communication Language. Further institutions, that we call *special institutions*, can then be defined to model the aspects of institutional reality typical of certain application domains. For instance, for electronic commerce applications it will be necessary to model the institutions of ownership, money, business transactions, auctions, and so on. A formal specification of the English auction as a special institution is given in [Fornara et al., 2004].

The Basic Institution provides for a definition of communicative acts in terms of their effects on commitments. To do so it is necessary to define an ontology of commitment, the institutional actions necessary to operate on commitments, and the authorizations to perform such actions. As we shall see, all institutional actions on commitment can be performed by means of the basic communicative act: the declaration. In Section 4 we will define a set of conventions necessary to perform some type of communicative acts directly by the exchange of a suitable message.

3.1 The Ontology of Commitment

We regard a commitment as an institutional 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. We assume that a given the debtor, the creditor, and the content a commitment could be univocally identified. Commitments will be represented with the following notation:

Comm(state, debtor, creditor, content)

The content of a commitment can be represented by means of a *temporal proposition* (for a detailed treatment of temporal propositions see [Fornara and Colombetti, 2003, Colombetti et al., 2004]), that is used to represent an *action*, a *proposition*, or a *referential expression* referred to a specific interval of time. At every time instant, a temporal proposition has a truth value, which can be *undefined*, *true*, or *false*. We perceive that commitments to actions and commitments to propositions have different aspects [Walton and Krabbe, 1995], but a detailed treatment of these is beyond the scope of this paper.

The state of a commitment undergoes a life cycle, described by the state diagram of Figure 2, and can change as an effect of the execution of institutional actions (solid lines) or of environmental events (dotted lines). Relevant events are due to the change of the truth-value of the commitment's content.

In our approach commitments can be created by individual agents through the execution of communicative acts, or can be created by norms and undertaken by an agent in virtue of its role in an institution.

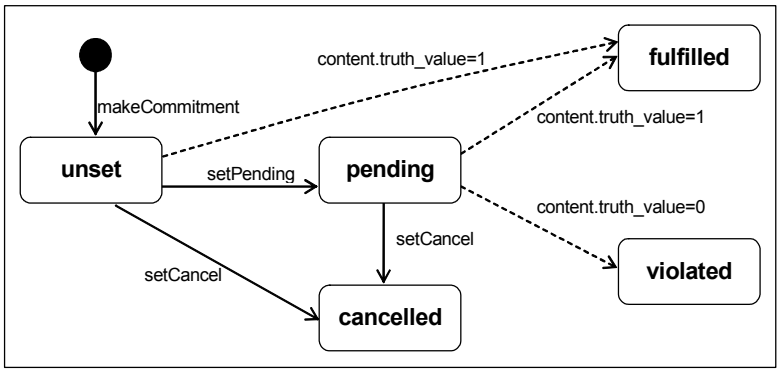

Fig. 2. The life-cycle of commitments.

3.2 Institutional Actions on Commitment

The institutional actions that operate on commitments are defined below; preconditions and effects are described using Object Constraint Language (OCL) [Object Management Group, 2003].

- *name*: ***makeCommitment****(debtor, creditor, content)*
  pre: *not Comm.allInstances→exists(c|c.debtor=debtor and c.creditor=creditor and c.content=content)*
  post: *Comm.allInstances→exists(c|c.state=unset and c.debtor=debtor and c.creditor=creditor and c.contet=content)*

- *name:* ***setCancel****(debtor, creditor, content)*
  pre: *Comm.allInstances→exists(c|(c.state=unset or c.state=pending) and c.debtor=debtor and c.creditor=creditor and c.content=content)*
  post: *Comm.allInstances→exists(c|c.state=cancelled and c.debtor=debtor and c.creditor=creditor and c.content=content)*

- *name:* ***setPending****(debtor, creditor, content)*
  pre: *Comm.allInstances→ exists(c|c.state=unset and c.debtor=debtor and c.creditor=creditor and c.content=content)*
  post: *Comm.allInstances→ exists(c|c.state=pending and c.debtor=debtor and c.creditor=creditor and c.content=content)*

It is often useful to define institutional macro-actions, that is, actions whose execution coincides with the sequential execution of a list of previously defined institutional actions, conceived of as a single transaction. For example:

name: ***makePendingComm****(debtor,creditor,content)* $=_{def}$
makeCommitment*(debtor,creditor,content);*
setPending*(debtor,creditor,content)*

3.3 Authorizations

We define a set of authorizations concerning the creation and the manipulation of commitments. Such authorizations will be associated to the two roles introduced by commitments themselves: the role of *debtor* and the role of *creditor*. Moreover, we assume a universal role, *RegAgt*, which every registered agent plays throughout its lifetime:

– any registered agent can create an *unset* commitment with any other registered agent as debtor or creditor:

Auth(RegAgt, ***makeCommitment****(debtor, creditor, content))*;

– the debtor of an *unset* commitment can set it to pending:

Auth(Comm(debtor, creditor, content).debtor,
setPending*(debtor, creditor, content))*;

– the debtor of an *unset* commitment can set it to *cancelled*:

Auth(Comm(unset, debtor, creditor, content).debtor,
setCancel*(debtor, creditor, content))*;

– the creditor of a commitment can set it to *cancelled*:

Auth(Comm(debtor, creditor, content).creditor,
setCancel*(debtor, creditor, content))*.

These authorizations allow an agent to perform all communicative acts that will be defined in Section 4; they may be modified or extended within special institutions. In general, institutions also define sets of norms to regulate the behavior of agents, in our current view, the Basic Institution does not specify norms; however, norms are introduced by most special institutions. Therefore, in Section 5 we give a detailed description of our approach to norm.

4 A Communicative Act Library

In this section we define a library of communicative acts and the set of conventions that bind the exchange of a message, characterized by a message type and a content structure, to their performance. To be compatible with the syntax of FIPA-ACL [Foundation for Intelligent Physical Agents, 2004a], when possible we name our communicative acts with the FIPA performative that has the closest intuitive meaning.

Communicative act content

A crucial requirement for all content languages compatible with FIPA-ACL is to be able to express at least descriptions of propositions and actions [Foundation for Intelligent Physical Agents, 2004b]. Such feature has been exploited to model content language concepts required by FIPA-ACL in terms of a set of objects that represent actions and propositions without any commitment about how a particular content language is structured ([Cranefield and Purvis, 2001] and [van Aart et al., 2002]). One of the most valuable advantages of this approach is that it is possible to transform such object oriented representation into different content languages.

Following this stream of research, we assume that content language expressions are temporal propositions (see Section 3.1) that define the derived property *expressionType*, which can assume the following values: *action*, *proposition*, or *referential expression*. This allows us to define conditions on the kind of expression that can be used as content of each message and as parameter of each communicative act.

Communicative act library

In the sequel, the semantics of communicative acts is partly given in terms of preconditions and postconditions of other institutional action. In fact, most of the communicative acts differ from institutional actions defined by the Basic Institution only because they have additional conditions on the type of the content. Thus, we introduce two new OCL operators, *precondition* and *postcondition*, to evaluate the precondition and postcondition of an institutional action. Applying these operators, we can describe the act of informing, whose main point is to commit the sender of the message to the truth value of a proposition, as follows:

name: ***Inform****(sender, receiver, content)*
pre: ***makePendingComm****(sender, receiver, content).precondition*
and content.expressionType = proposition
post: ***makePendingComm****(sender, receiver, content).postcondition*

The ***Inform*** communicative act can be performed by exchanging a message with the FIPA *inform* message type:

exchMsg*(inform, sender, receiver, content)* $=_{conv}$ ***Inform****(sender, receiver, content)*

The definition of ***Inform*** allows us to define ***Failure***, that is, the act that informs an agent that the performance of the described action has failed and the reason for the failure.

name: ***Failure****(sender, receiver, action, cause)*
pre: *action.expressionType=action and cause.expressionType= proposition and*
Inform*(sender, receiver, not(done action) and cause).precondition*
post: ***Inform****(sender, receiver, not(done action) and cause).postcondition*

The convention that allows agents to perform a ***Failure*** communicative act is:

exchMsg*(failure, sender, receiver,(action, cause))* $=_{conv}$
Failure*(sender, receiver, action, cause)*

An agent may get another agent to perform an action by performing a request, which creates an unset commitment for the receiver of the message to the performance of an action, which can be another communicative act. The formal definition of ***Request*** is the following:

name: ***Request****(sender, receiver, content)*
pre: ***makeCommitment****(sender, receiver, content).precondition and*
content.expressionType=action
post: ***makeCommitment****(sender, receiver, content).postcondition*

By sending a message with the request performative, the sender can perform the ***Request*** communicative act:

exchMsg*(request, sender, receiver, content)* $=_{conv}$ ***Request****(sender, receiver, content)*

Unlike most content languages, temporal propositions allow us to express conditional commitments for the execution of an action *when* or *whenever* the described conditions are met (for details see [Colombetti et al., 2004]). Thus, in our framework ***Request-When*** and ***Request-Whenever*** coincide with ***Request***. Furthermore, using the definition of ***Request***, we can give the semantics of several acts defined by FIPA, like ***QueryIf***, ***QueryRef***, ***Proxy***, ***Subscribe***, and ***Propagate***. For sake of brevity, we report only the semantics of ***QueryIf***, whose point is to get another agent to answer whether a state of affairs holds:

name: ***QueryIf****(sender, receiver, content)* $=_{def}$
Request*(sender, receiver,* ***Inform****(content) or* ***Inform****(not content))*

An agent can agree to perform the requested action by executing ***Agree***, which means setting to pending an unset commitment:

name: ***Agree****(sender, receiver, content)*
pre: ***setPending****(sender, receiver, content).precondition and*
content.expressionType=action

post: ***setPending****(sender, receiver, content).postcondition*

The creditor of an unset commitment can agree to do an action thanks to the following convention:

exchMsg*(agree, sender, receiver, action)* $=_{conv}$ ***Agree****(sender, receiver, action)*

Otherwise, an agent that has received a request (or an act that can be defined in terms of **Request**) can refuse it, by canceling the unset commitment. The definition of the **Refuse** communicative act and the convention for its execution are:

name: ***Refuse****(sender, receiver, content)*
pre: ***setCancel****(sender, receiver, content).precondition and*
content.expressionType=action
post: ***setCancel****(sender, receiver, content).postcondition*

exchMsg*(refuse, sender, receiver, content)* $=_{conv}$ ***Refuse****(sender, receiver, content)*

As described in Section 3, a creditor of a commitment can cancel it if it has not yet reached a final state by performing the ***Cancel*** communicative act:

name: ***Cancel****(sender, receiver, content)*
pre: ***setCancel****(receiver, sender, content).precondition and*
content.expressionType=action
post: ***setCancel****(receiver, sender, content).postcondition*

exchMsg*(cancel, sender, receiver, content)*$=_{conv}$ ***Cancel****(sender, receiver, content)*

An agent can commit itself to the performance of an action by promising it:

name: ***Promise****(sender, receiver, content)*
pre: ***makePendingComm****(sender, receiver, content).precondition and*
content.expressionType=action
post: ***makePendingComm****(sender, receiver, content).postcondition*

exchMsg*(promise, sender, receiver, content)* $=_{conv}$ ***Promise****(sender, receiver, content)*

Promise is not defined in the FIPA communicative act library. In fact, FIPA does not provide an act for communicating that an agent has an unconditional intention to perform an action. Instead, FIPA describe the ***Propose*** act, whose semantics states that the sender will hold the intention to execute an action if the receiver of the message accepts such intention. At present, our model does not provide any support for a commitment that should be accepted by the creditor and, for this reason, the communicative acts related with ***Propose***, like ***Reject-Proposal***, ***Accept-Proposal*** and ***Call-for-Proposal***, are not defined in our library. Furthermore, communicative act ***Confirm*** and ***Disconfirm*** cannot be defined in our model because their semantics is related to the mental state approach of FIPA-ACL. In fact, an agent perform a confirmation (or a disconfirmation) if it knows that the receiver is uncertain about a proposition.

All communicative acts previously described are defined in terms of preconditions and postconditions derived from institutional actions declared in the Basic Institution or in terms of other communicative acts whose definition exploits such basic actions. Thus, the agent that executes a communicative act, the sender of the message, should be authorized to perform the corresponding institutional action defined in the Basic Institution.

Finally, we consider declarations, a kind of communicative act that FIPA-ACL does not define because its semantics needs a model of institutional reality. In fact, according to Searle's Speech Act Theory [Searle, 1969], declarations are the particular category of communicative acts whose point is to bring about a change in the institutional reality in virtue of their successful performance. By definition the content of a declaration describes precisely the institutional changes that it brings about. Therefore we define the declaration institutional action as:

name: ***Declare****(sender, receiver, iaction(parameters))*
pre: *iaction(parameters).precondition*
post: *iaction(parameters).postcondition*

Unlike others communicative acts, which are always authorized for the sender of the message, only agents that are empowered to perform an institutional action are authorized also to declare such action:

*Auth(sender, **Declare**(sender, receiver, iaction(parameters)), Auth(sender, iaction(parameters))*

We introduced the use of the ***Declare*** institutional action for two reasons: for uniformity with the definition of other communicative institutional actions and for similarity with the human way to communicate. We take messages of type *declare* as the fundamental means to perform institutional actions. The convention that binds the exchange of a *declare* message to the performance of the institutional action ***Declare*** is:

$exchMsg(declare, sender, receiver, iaction(parameters)) =_{conv}$
***Declare**(sender, receiver, iaction(parameters))*

Every communicative act defined in this section can be performed also by declaring it. Furthermore, given that such communicative acts are defined in terms of institutional action on commitments, agents can obtain the same communicative effects by declaring the corresponding institutional action.

5 Norms

In a special institution, the execution of an action by an authorized agent often needs to be regulated by another fundamental component of artificial institutions, that is, a system of norms. For example, the auctioneer of an English Auction not only is authorized to declare an auction open, but it is also obliged to do so under certain circumstances. Norms prescribe which institutional actions should or should not be executed, among those that are authorized. In doing so, norms play an important function, 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 [Moses and Tennenholtz, 1995, Barbuceanu et al., 1998]. In particular, we think that 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, to comply with the regulations of a specific institution.

We regard norms as event-driven rules that fire under appropriate conditions and, by doing so, create, update or cancel commitments affecting a predefined set of agents. At an abstract level, a norm is part of the definition of an artificial institution (see Section 2); its instances then regulate and are bound to the organization that reifies the institution. Agents are liable to all the norms associated to the roles they play in an institution.

A norm is defined within an institution, observes an entity of an institution, is activated by an event concerning such an entity, and then fires if certain contextual conditions are met. Typically, interesting event types are the filling of a role by an agent, a value change of an institutional attribute, the reaching of a certain instant of time, and so on.

When a norm fires, it is applied to a collection of liable agents, that 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 updates or cancels a set of commitments. 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** selection expression*
***do** commitmentActionDescription {; commitmentActionDescription}**

For example, in our formalization of the English Auction (see [Fornara et al., 2004]) the following norm creates an obligation for the auctioneer to open the auction when there are more than 2 participants and the start time has elapsed:

***within** a: UnsetEnglishAuction*
***on** e: TimeEvent(a.startTime)*

if *a.participant.sizeOf() > =2* ***then***
foreach *agent* ***in*** *a.auctioneer*
do *makePendingComm(agent, a.organization, (openAuction(a.id),[now, now+δ],∃))*,

where δ is the time allowed to the auctioneer to fulfill the obligation.

Many studies have been devoted to the analysis of the relationship holding between norms and commitments, which is often perceived as a fundamental aspect of institutions [Esteva et al., 2001] and organizations [Castelfranchi, 1995]. For example in [Lopez y Lopez and Luck, 2003] commitments are viewed as a specialization of norms, while in [Castelfranchi, 1995] and [Singh, 1999b] 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; they do not have a debtor or a creditor, and strictly speaking they cannot be fulfilled or violated. Indeed, what can be fulfilled or violated is not a norm, but a commitment created by the application of a norm. There are, in conclusion, two types of commitments: the ones created by individual agents through the execution of communicative acts, and the ones created by norms and undertaken by an agent in virtue of its role in an institution.

6 Conclusions

In this paper we have defined what we mean by the term "Artificial Institution", a description of the basic concepts that constitute agent interaction systems. We regard an institution as constituted by an ontology, a set of authorizations, a set of conventions, and a systems of norms. In particular, we have described the Basic Institution, the institution that defines the concept of commitment, that we assume as the fundamental entity to describe ACL semantics.

In our approach, agents can modify institutional reality by exchanging suitable messages that, thanks to the definition of conventions, count as the performance of institutional actions. We propose a Communicative Act Library, composed by a set of communicative actions and a set of conventions for their execution. Furthermore, we have discussed the crucial role played by declarations to allow agents to perform every institutional action, even communicative acts. In doing so, the interacting agents can execute actions to which no specific convenion is associated.

Finally, we have discussed the function of norms as means for regulating agents' behaviour and the relationship between norms and commitments. We report a brief example of how the interaction protocol of the English Auction can be described in terms of a set of norms.

We believe that our approach helps clarifying the strict relationships holding between language, institutional reality, and interaction rules in a MAS. Moreover, we believe that the adoption of an operational modelling style makes our proposal reasonably easy to implement. In fact, we plan to implement our framework as an extension of JADE in the near future.

References

[Amgoud et al., 2002] L. Amgoud, N. Maudet, and S. Parsons. An argumentation-based semantics for agent communication languages. In F. V. Harmelen, editor, *Proceedings of the European Conference on Artificial Intelligence* (ECAI-2002), pages 38–42, Lyon, France, July 2002. IOS Press.

[Barbuceanu et al., 1998] 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.

[Bentahar et al., 2004] J. Bentahar, B. Moulin, and B. Chaib-draa. Commitment and argument network: A new formalism for agent communication. In F. Dignum, editor, *Advances in Agent Communication, International Workshop on Agent Communication Languages, ACL 2003*, Melbourne, 2003, volume 2922 of LNCS, pages 146–165. Springer, 2004.

[Castelfranchi, 1995] 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.

[Colombetti, 2000] M. Colombetti. A commitment–based approach to agent speech acts and conversations. *In Proc. Workshop on Agent Languages and Communication Policies, 4th International Conference on Autonomous Agents (Agents 2000)*, pages 21–29, Barcelona, Spain, 2000.

[Colombetti and Verdicchio, 2002] M. Colombetti and M. Verdicchio. An analysis of agent speech acts as institutional actions. In C. Castelfranchi and W. L. Johnson, editors, *Proc. AAMAS 2002*, pages 1157–1166, Bologna, Italy, 2002.

[Colombetti et al., 2004] M. Colombetti, N. Fornara, and M. Verdicchio. A social approach to communication in multiagent systems. In J. A. Leite, A. Omicini, L. Sterling, and P. Torroni, editors, *Declarative Agent Languages and Technologies*, volume 2990 of LNAI, pages 191–220. Springer, 2004.

[Cranefield and Purvis, 2001] S. Cranefield and M. Purvis. Generating ontology-specific content languages. *In Proceedings of the Workshop on Ontologies in Agent Systems, 5th International Conference on Autonomous Agents*, 2001

[Esteva et al., 2001] M. Esteva, J. A. Rodrìguez-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* (LNAI 1991), pages 126–147. Springer, 2001.

[Fornara and Colombetti, 2002] N. Fornara and M. Colombetti. Operational specification of a commitment-based agent communication language. In C. Castelfranchi and W. L. Johnson, editors, *Proc. AAMAS 2002*, pages 535–542, Bologna, Italy, 2002.

[Fornara and Colombetti, 2003] N. Fornara and M. Colombetti. Defining interaction protocols using a commitment-based agent communication language. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proc. AAMAS 2003*, pages 520–527, Melbourne, Australia, 2003.

[Fornara et al., 2004] N. Fornara, F. Viganò, and M. Colombetti. Agent Communication and Institutional Reality. *In Proc. of the AAMAS 2004 Workshop on Agent Communication,* pages 4-18, New York, USA, 2004.

[Foundation for Intelligent Physical Agents, 2004a] Foundation for Intelligent Physical Agents. FIPA Communicative Act Library Specification. 2004 (http://www.fipa.org).

[Foundation for Intelligent Physical Agents, 2004b] Foundation for Intelligent Physical Agents. FIPA Content Language (CL) Specifications. 2004. (http://www.fipa.org).

[Jones and Sergot, 1996] A. Jones and M. J. Sergot. A formal characterisation of institutionalised power. *Journal of the IGPL*, 4(3):429–445, 1996.

[Lopez y Lopez and Luck, 2003] F. Lopez y Lopez and M. Luck. Modelling Norms for Autonomous Agents. In E. Chavez, J. Favela, M. Mejia, and A. Oliart, editors, *Proceedings of Fourth Mexican International Conference on Computer Science*, pages 238–245, 2003.

[McBurney and Parsons, 2003] P. McBurney and S. Parsons. Posit spaces: a performative model of e-commerce. In J. S. Rosenschein, T. Sandholm, M. Wooldridge, and M. Yokoo, editors, *Proc. AAMAS 2003*, pages 624–631, 2003.

[Moses and Tennenholtz, 1995] Y. Moses and M. Tennenholtz. Artificial social systems. *Computers and AI*, 14(6):533–562, 1995.

[Object Management Group, 2003] Object Management Group, OMG . Object Constraint Language Specification 1.4. http://www.omg.org/, 2003.

[Searle, 1969] J. R. Searle. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge University Press, Cambridge, United Kingdom, 1969.

[Searle, 1979] J. R. Searle. *Expression and Meaning*. Cambridge, United Kingdom: Cambridge Univesity Press, 1979.

[Searle, 1995] J. R. Searle. *The construction of social reality*. Free Press, New York, 1995.

[Singh, 1999a] M. P. Singh. A social semantics for agent communication languages. In *Proceedings of IJCAI-99 Workshop on Agent Communication Languages*, pages 75–88, 1999.

[Singh, 1999b] M. P. Singh. An ontology for commitments in multiagent systems: Toward a unification of normative concepts. *Artificial Intelligence and Law*, 7:97–113, 1999.

[van Aart et al., 2002] C. J. van Aart, R. F. Pels, G. Caire and F. Bergenti. Creating and Using Ontologies in Agent Communication. *Proceedings of Workshop on Ontologies in Agent Systems AAMAS 2002*, 2002.

[Verdicchio and Colombetti, 2003] M. Verdicchio and M. Colombetti. A logical model of social commitment for agent communication. In J. S. Rosenschein, T. Sandholm, M.Wooldridge, and M. Yokoo, editors, *Proc. AAMAS 2003*, pages 528–535, Melbourne, Australia, 2003.

[Walton and Krabbe, 1995] D. N. Walton and C. W. Krabbe. Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning. SUNY Press, Albany, NY, 1995.