

A COMMITMENT-BASED APPROACH TO AGENT COMMUNICATION*

Nicoletta Fornara¹ and Marco Colombetti^{1,2}

¹ Università della Svizzera italiana, via G. Buffi 13, 6900
Lugano, Switzerland

{nicoletta.fornara, marco.colombetti}@lu.unisi.ch,

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

In this paper we propose an operational method for the definition of the semantics of Agent Communication Languages based on the notion of social commitment. Our proposal is suitable for open interaction frameworks where agents, designed by independent constructors, dynamically enter and leave different interaction systems. In this type of environments it is crucial to define a standard and commonly accepted semantics for the exchanged messages. We give an operational specification of commitment and introduce temporal propositions for the representation of the contents of commitments within an object-oriented paradigm. Then we use operations on commitments to define the meaning of a set of communicative acts that is complete with respect to Searle's taxonomy of illocutionary acts, and give an example of the use of communicative acts in an interaction protocol.

INTRODUCTION

The definition of a commonly accepted standard semantics for Agent Communication Languages (ACLs) has become more and more important in recent years thanks to the advent of actually existing *open environments* such as the Internet, and, hopefully in the near future, the Semantic Web (Berners-Lee et al. 2001). These kinds of environments are characterized by the fact that heterogeneous agents created by independent designers, having different and sometimes conflicting interests, are free to dynamically enter and leave various interactions systems. To achieve this result it is necessary to identify and formally define a set of basic concepts necessary to

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

describe the rules that govern the interaction of different systems and in particular, as a first step, the library of communicative acts available for communication. Starting from the assumption that language is the fundamental component of every interaction, in this paper we propose a semantics for ACLs, based on the notion of *social commitment* (Castelfranchi 1995) and suitable to be used in open environments.

There are some crucial requirements, considered most important in the literature, that an ACL has to satisfy (Mayfield, Labrou, and Finin 1995; Singh 1998; Labrou, Finin, and Peng 1999). First, semantics has to be *verifiable*. This problem has been thoroughly discussed by Wooldridge (1998), who regards verifiability as the possibility to determine whether or not a system *conforms* to a given ACL. Second, the meaning of communicative acts has to be expressed in an *objective* and *external* way, in order that everybody attributes the same meaning to the messages and that the meaning is expressed using concepts that do not depend on the agent's internal structure. Another crucial characteristic is that an ACL should be *extensible* to let agents (or better, their designers) cope with new situations. Finally, an ACL should be *simple* to be correctly interpreted by agent designers, and enough *expressive* to be used in all common situations.

All existing proposals for the definition of a standard agent communication language are based on *Speech Act Theory* (Austin 1962; Searle 1969). This approach to the study of communication is so successful because viewing language as a form of action makes it possible to treat communicative acts and other types of action in a uniform way. In literature it is possible to identify many approaches to the definition of standard ACLs. The early approaches like KQML (Finin, Labrou, and Mayfield 1997) and FIPA ACL (Foundation for Intelligent Physical Agents 1997; 2002) define the meaning of a communicative act using agent's *mental states*, like beliefs, desires and intentions. This approach may be adequate in *cooperative* multiagent system, but is problematic when the multiagent system is composed by *competitive*, heterogeneous agents made by different vendors, where it is impossible to trust other agents completely or to make strong assumptions about their internal way of reasoning. To overcome these problems more recent approaches, like ours, take into account social, objective consequences of performing a communicative act and express the meaning of a communicative act in terms of commitments directed from one agent to another (Singh 1999; Colombetti 2000; Yolum and Singh 2001; Fornara and Colombetti 2002; 2003; Verdicchio and Colombetti 2003). The notion of commitment has also been used in the field of argumentation studies (Amgoud, Maudet, and Parsons 2002; Bentahar, Moulin, and Chaib-draa 2004). In another approach the meaning of a communicative act is implicitly defined as the role it plays in a given set of conversation protocols (Pitt and Mamdani 1999). A problem with this approach, however, is that any change in the

set of accepted protocols is going to affect the meaning of communicative acts. Recently other approaches have been proposed (Dignum 2003; Viroli, Ricci, and Omicini 2004), but are still at an initial stage of development.

The approach to the definition of the semantics of ACLs presented in this paper is based on an analysis of the primitive notion of commitment, and is an amended version of the work presented in (Fornara and Colombetti 2002; 2003). We give an operational specification of commitment as an abstract data type, and analyze its life cycle as a trajectory through a suitable state space. To express the content of commitments, we introduce the notion of *temporal proposition* (a more detailed analysis of this concept can be found in (Colombetti, Fornara, and Verdicchio 2004)). Commitments are then used to define the meaning of basic types of communicative acts corresponding to Searle's taxonomy of illocutionary acts (Searle 1979). In this way, ACL messages are given a formal semantics, thus eliminating any ambiguity in their meaning. Moreover, as suggested in (Fornara and Colombetti 2003), it becomes possible to check the soundness of conversation protocols defined in terms of communicative acts. Another important feature of our proposal is that it leads to the definition of a communicative act library that is extensible and allows for the reuse of various components. Starting from a small set of basic operations on commitments, it is possible to define the meaning of simple communicative acts, which in turn can be used to define a new layer of more complex communicative acts, and so on.

This paper is structured as follows. In Section 2 we give the operational specification of the components necessary to define the semantics of an Agent Communication Language. In Section 3 we define the meaning of the main communicative act types. In Section 4 we illustrate the use of commitments in a realistic interaction protocol. Finally, in Section 5 we discuss some relevant aspects of our work and compare it with other approaches.

OPERATIONAL SPECIFICATION OF MAIN CONCEPTS

The semantics proposed in this paper is given by describing the effects that sending a message has on the social relationship between the sender and the receiver of the message. To do so we rely on a concept that is unambiguous, objective, and independent of the agents' mental states, that is, social commitment.

Commitments

To represent the network of commitments binding the interacting agents we use *commitment objects* that have an internal structure, a life cycle, and a set of methods available for their manipulation. The internal structure of a commitment object consists of the following fields:

- a unique commitment *identifier (id)*;
- a reference to the commitment's *debtor*, that is, the agent that has the commitment;
- a reference to the *creditor*, that is, the agent relative to which the debtor is committed;
- the commitment's *content*, that is, a temporal proposition (see below) describing a state of affairs or a course of action, to which the debtor is committed relative to the creditor;
- the commitment's *condition*, that is introduced to treat conditional commitment; it is a temporal proposition (see below) that represents the condition that has to be satisfied in order for the commitment to become active. There is a constraint that forces the time interval of the condition to precede the time interval of the content.
- a *state* used to keep track of the dynamic evolution of the commitment and that can assume the following values: *unset*, *cancelled*, *pending*, *active*, *fulfilled*, *violated*. In particular the *unset* state has been introduced to express the meaning of directive communicative acts like requests and the *pending* state is necessary to express conditional commitments.
- a *timeout*, which is relevant only in the case of *unset* commitments, and will therefore be treated as an optional parameter. It represents the time limit for the debtor of an *unset* commitment to accept, fulfil or reject it. As will be discussed below after it is elapsed the activation of rule 7 transforms the commitment to *cancelled*.

Commitment objects will be represented with the following notation:

$C_{id}(state, debtor, creditor, content|condition[,timeout])$.

The life cycle of a commitment object is described by the finite state machine in Figure 1. The state of a commitment can change as an effect of the invocation of its basic methods (solid lines) or of environmental events (represented by dotted lines labelled with the name of the related *update rule* described in Table 1), that is, of events that change the truth-value of a commitment's conditions or content. Update rules formalize event driven routines that are automatically invoked when the temporal proposition associated with the content or the condition of commitment changes its truth-value from undefined to true or false. In particular rule 5 is necessary to cover situations when an agent performs the requested action without accepting the request explicitly. (In Table 1 we denote with *T* a temporal proposition whose truth value is true).

We assume that when a commitment object is declared, the constructor of the class creates an empty commitment object, $Ci()$. We represent the invocation of a method by

the name of the object followed by a dot and by the name of the method with its parameter list. Commitments are created and manipulated through the following basic operations:

- *Make commitment.* By invoking the method $mc(a,b,P,Q)$ with arbitrary debtor a , creditor b , content P , and condition Q , a new unset commitment object is created:

$$C_i().mc(a,b,P,Q[,to]) \rightarrow C_i(unset,a,b,P|Q[,to])$$

- *Set commitment.* The method $sc(s)$ changes the current state of an existing commitment object (whichever it is “-”) to state s :

$$C_i(-,a,b,P|Q).sc(s) \rightarrow C_i(s,a,b,P|Q)$$

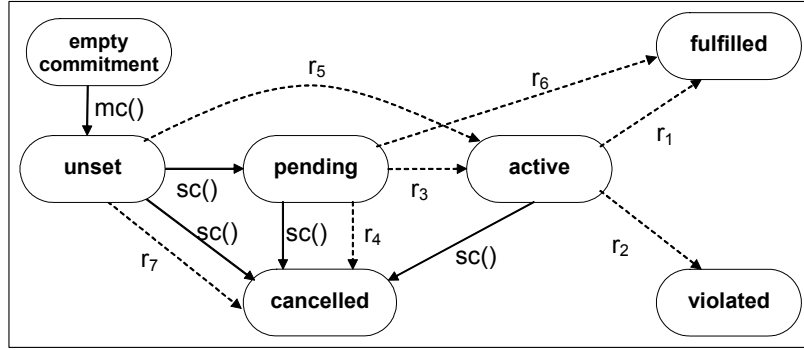


FIGURE 1. The life-cycle of commitments.

Basic operations should not be viewed as actions that are directly performed by agents. Rather, they are low-level primitives used to implement operations on commitment objects, more specifically, agents manipulate commitments through a communicative act library. The operations introduced are necessary for the definition of the library of communicative acts presented in Section 3, other ones, for example for the manipulation of content or condition can be introduced in order to be able to define new type of communicative acts.

In this paper we do not tackle the crucial problem of formalizing the concept of an electronic institution (Esteva et al. 2001), where interactions among communicative agents take place. Anyway regarding the complete definition of the ontology of the commitment we think that it is possible to list a reasonable set of basic authorizations that have to be taken into account when new communicative acts are defined using operations on commitment objects. Such basic authorizations are:

- any agent can create an *unset* commitment with arbitrary debtor and creditor;
- the debtor of an *unset* commitment can set it to either *pending* or *cancelled*;
- the creditor of an *unset*, *pending*, or *active* commitment can set it to *cancelled*.

These basic authorizations may be modified or new ones may be introduced on the basis of the particular electronic institution where the interaction actually takes place.

Finally note that we defined the conditions under which commitments are fulfilled or violated, but we are not concerned with the management of violations, for instance in terms of sanctions, because this aspect lies beyond the definition of ACL semantics.

TABLE 1. Update Rules.

event	action	rule
$P.truth_value() = 1$	$C_i(a,a,b,P T) \rightarrow C_i(f,a,b,P T)$	1
$P.truth_value() = 0$	$C_i(a,a,b,P T) \rightarrow C_i(v,a,b,P T)$	2
$Q.truth_value() = 1$	$C_i(p,a,b,P Q) \rightarrow C_i(a,a,b,P Q)$	3
$Q.truth_value() = 0$	$C_i(p,a,b,P Q) \rightarrow C_i(c,a,b,P Q)$	4
$P.truth_value() = 1$	$C_i(u,a,b,P T) \rightarrow C_i(a,a,b,P T)$	5
$P.truth_value() = 1$	$C_i(p,a,b,P Q) \rightarrow C_i(f,a,b,P Q)$	6
$curr_time > t$	$C_i(u,a,b,P Q,t) \rightarrow C_i(c,a,b,P Q)$	7
<i>legend: u = unset, p = pending, c = cancelled, a = active, f = fulfilled, v = violated</i>		

Temporal propositions

The definition of an ACL is strictly related to the specification of a content language used to express the content of messages. None of existing proposals for content language specification (Foundation for Intelligent Physical Agents 2004) provide for a standard treatment of temporal aspects. We think that this is a crucial aspect especially if the ACL semantics is defined in terms of commitments, because commitments often specify deadlines for the execution of actions. Therefore in this section, even if we do not intend to define a new content language, we propose a representation of commitment content and condition that explicitly takes time into account. A temporal proposition object consists of the following fields:

- a *statement* in a suitable language which may state that: (i) a certain state of affairs holds; (ii) an action has been performed; (iii) a specific commitment with certain specific attributes holds;
- the *truth value* of the statement, which may be true (1), false (0) or undefined (\perp);
- a *time interval*, which may go from a single instant to the entire life of the system, relative to which the statement is considered;

- and a temporal *mode*, which specifies whether the statement should be true for the whole time interval (\forall) or on at least an instant of the time interval (\exists).

We assume that the truth value of temporal proposition objects is updated by a suitable “notifier”. In particular: if the mode is ' \exists ' the notifier sets the truth-value to true if the statement becomes true at any point of the time interval, otherwise sets it to false when the time interval expires; if the mode is ' \forall ' the notifier sets the truth-value to false if the statement becomes false at any point of the time interval, otherwise sets it to true when the time interval expires. It is important to remark that the truth value of a temporal proposition object can switch from \perp to 1 or 0, but then cannot change any more. Temporal proposition objects are represented with the following notation:

$P_{id}(\text{statement}, \text{time interval}, \text{mode}, \text{truth_value})$.

In particular temporal proposition objects can be used to represent conditions on the temporal evolution of commitments, an example is given in Section 3.5.

Other Concepts

To define the meaning of declarations we need to introduce some other components that contribute to model the social context of the interaction between agents (a more detailed analysis of these concepts can be found in (Fornara, Viganò, and Colombetti 2004)). These concepts are:

- the notion of domain-specific objects, which represent entities of the application world. Such entities may possess both “natural” and “institutional” attributes; for example, the size of a product being sold is a natural attribute, while its price is an institutional attribute. Natural attributes are assumed to reflect the physical properties of the corresponding entities of the real world; on the contrary, institutional attributes can be affected by the performance of certain communicative acts, in particular by declarations (as discussed below). We assume that each domain-specific object has a value-setting method for each of its institutional properties; for example the method “*setState()*” can be invoked to set the “state” property.
- The notion of *role* useful to be able to refer to agents abstracting from the specific ones that take part in an interaction.
- The notion of *authorizations* associated to roles, used to specify which agent is authorized to perform a particular declaration (see Section 3 for details).

A LIBRARY OF COMMUNICATIVE ACTS

In this section the previously defined operations on commitment will be used to express the meaning of the basic types of communicative acts as identified by Speech Act Theory (Searle 1979). We do not treat expressives, because we think they are not relevant for artificial agents interactions. In the following definitions the symbol “:=” means that the act represented on the left-hand side is actually performed through the invocation of the methods listed on the right-hand side, and the sign “ $=_{def}$ ” means that performing the action represented on the left-hand side is the same as performing the action represented on the right-hand side.

Assertives

According to Speech Act Theory, the point of an assertive act is to commit the sender, relative to the receiver, to the truth of what is asserted. We consider the inform act as our prototypical assertive act. This act is used by agent a to inform agent b that P is the case. In a commitment-based approach, an act of informing can be defined as follows (T is the identically true temporal proposition object):

$$\mathbf{inform}(a,b,P) := \{C_i().mc(a,b,P,T); C_i(unset,a,b,P|T).sc(pending)\}$$

The final result is an active commitment, thanks to the intervention of rule 3.

Directives

As defined in Speech Act Theory, the point of a directive act is to get the receiver to perform an action (usually a communicative act) within an interval of time. We treat request as our basic directive act, and given that an agent can not directly commit another agent we define it as the creation of an unset commitment with the sender as the creditor and the receiver as the debtor. The request by agent a to agent b to bring about P if condition Q is satisfied is defined as:

$$\mathbf{request}(a,b,P,Q[,to]) := \{C_i().mc(b,a,P,Q[,to])\}.$$

The receiver of a request can react in three different ways: it can perform the requested action (like for example when the request is “please close the door”), accept the request, or refuse it, using the suitable commissive act as discussed in the next section. In any case an unset commitment is automatically cancelled after a predefined *timeout* is expired.

Questions are requests to be informed about something. There are questions that are used to get the receiver to commit to the truth value of a given proposition and are treated in detail in (Fornara and Colombetti 2002) like for example “is the book red?”.

Other questions are used to get the receiver to commit to the value of a given property, for instance “what’s your name?”. In this case the requested act of informing cannot be completely described by the sender (otherwise, why should it ask the question?). The sender provides a “template” for the answer, that is, a temporal proposition object $S(x)$ containing a meta-variable x that the receiver has to replace with a constant value c . A question has therefore the form:

$$\mathbf{request}(a,b,P,T), \quad \text{where } P.\text{statement}() = \mathbf{inform}(b,a,S(x)),$$

given that:

$$\mathbf{inform}(b,a,S(x)) =_{\text{def}} \mathbf{inform}(b,a,S(c)), \text{ for some constant value } c.$$

This definition implies that the performance of the requested inform act with the temporal proposition $S(c)$ as a parameter makes the temporal proposition P true. Indeed, as remarked by Searl (1991) the concept of a question is more general: by a question, an agent may request the execution of a non-assertive communicative act (like a directive, or a commissive). However, our definition above easily generalizes to such cases.

Commissives

The point of a commissive act, as defined by Speech Act Theory, is to commit the debtor, relative to the creditor, to the execution of an action of a given type within a given interval of time. Here we define the basic commissive act of promising. The promise by agent a to agent b to bring about P if condition Q is satisfied is defined as:

$$\mathbf{promise}(a,b,P,Q) := \{C_i().mc(a,b,P,Q); C_i(\text{unset},a,b,P|Q).sc(\text{pending})\}.$$

To make an unconditional promise the constant proposition object T is used as the condition, and thus the pending commitment created by the promise is immediately turned into an active commitment by rule 3. Two types of commissive acts can be performed only in connection with an unset commitment, namely *accept* and *reject* defined as follows:

$$\begin{aligned} \text{preconditions} &: \exists C_i(\text{unset}, b, a, P|Q[, to]) \\ \mathbf{accept}(b,a,C_i(\text{unset},b,a,P|Q[,to])) &:= \{C_i(\text{unset},b,a,P|Q[,to]).sc(\text{pending})\} \\ \text{preconditions} &: \exists C_i(\text{unset}, b, a, P|Q[, to]) \\ \mathbf{reject}(b,a,C_i(\text{unset},b,a,P|Q[,to])) &:= \{C_i(\text{unset},b,a,P|Q[,to]).sc(\text{cancelled})\} \end{aligned}$$

Declarations

Declarations are a special type of communicative acts. Examples of declarations are “I pronounce you man and wife” or “I declare the auction open”. The point of a declaration is to bring about a change in the world, obviously not in the physical or natural world but in an institutional world (Searle 1995), (Colombetti and Verdicchio 2002), that is, a conventional world relying on common agreement of the interacting agents (or, more precisely, of their designers). Declarations actually change the institutional world simply in virtue of their successful performance. In our interaction framework, to treat declarations we introduce objects with *institutional properties*, that is, conventional properties that result from common agreement, like for example the ownership of a product. Such properties can be affected by declaration acts. It is however necessary to identify which agents are *authorized* or *empowered* to perform a given declaration act in the system. Typically, authorizations are granted to agents in virtue of the *role* they play in an interaction, and thus authorizations are naturally associated to roles. To do so, we need to introduce a construct to express that an agent having a given role in the interaction system is empowered to bring about an institutional change of a given kind on an object O_k :

preconditions: $authorized(role_i, O_k.setProp_j()) \wedge a.role() = role_i$
declare $(a, O_k.prop_j = x) := \{O_k.setProp_j(x)\}.$

Proposals

Proposals do not belong to the basic taxonomy of illocutionary acts and can be defined as a combination of a directive and a commissive act. Proposals deserve special attention because they are crucial in many interesting application fields, like for example electronic commerce. A propose act can be defined as the parallel execution of a request and a promise, as denoted by the symbol \parallel :

propose $(a, b, P, Q[, to]) =_{def} \mathbf{request}(a, b, P, Q[, to]) \parallel \mathbf{promise}(a, b, Q, S)$
where $S.statement() = C_i(pending, b, a, P|Q)$

Note that in the above definition the temporal proposition object S will become true only if a commitment C_i with the values of the attributes as indicated will be created.

A SAMPLE APPLICATION

In this section we formalize a standard interaction protocol within the previously defined framework. To deal with a significant example, we analyze a conversational protocol widely used in electronic commerce applications, in particular in some types

of auctions: the protocol of proposals. The example shows how it is possible to use commitment-based definitions of communicative acts to express the meaning of the various messages exchanged in a negotiation and the dynamic evolution of the system states during an interaction.

A protocol is based on the set of the communicative acts previously defined as operations on commitment objects. It is described by an *interaction diagram*, that is, a graph whose nodes represent system states, and whose edges represent certain types of state transitions. In an interaction diagram, state transitions correspond either to communicative acts performed by the interacting agents, or to environmental events strictly related to the interaction. A more detailed description of how to define interaction protocols using the proposed framework and how to prove their *soundness* can be found in (Fornara and Colombetti 2003).

To each state is associated a *content*, consisting of commitment objects and temporal proposition objects holding at that particular state. Such objects are computed in the following way. The diagram has a distinguished *initial state (start)*, whose content is a set of empty commitments. If state s_j is reached from state s_i by a communicative act, then the content of state s_i must satisfy the preconditions of the communicative act, and the content of s_j is obtained by modifying the content of s_i according to the definition of the communicative act. If on the contrary state s_j is reached from state s_i by an environmental event, then the content of s_j is obtained by modifying the content of s_i according to the relevant update rule. States with no outgoing edges are classified as final; when a final state is reached, the interaction ends.

The example shows the dynamic evolution of the system states when a propose act is performed by agent a addressing agent b . Figure 2 reports the interaction diagram, while the content of each state is described in Table 2.

CONCLUSIONS

In this paper we have presented an operational definition of a commitment-based semantics for agent communicative acts, and shown, with an example, that this semantics is enough expressive to be used to specify interesting interaction protocols. Moreover the proposed semantics is objective and does not make any assumption about the internal structure of the agents that may use it. Keeping trace of the dynamic evolution of a system on the basis the actions performed by the agents is possible to verify if an agent is behaving in accordance with its commitments. Another important feature of our proposal is that is extensible, new acts can be defined using the presented notion of commitment, and allows the modular reuse of various components.

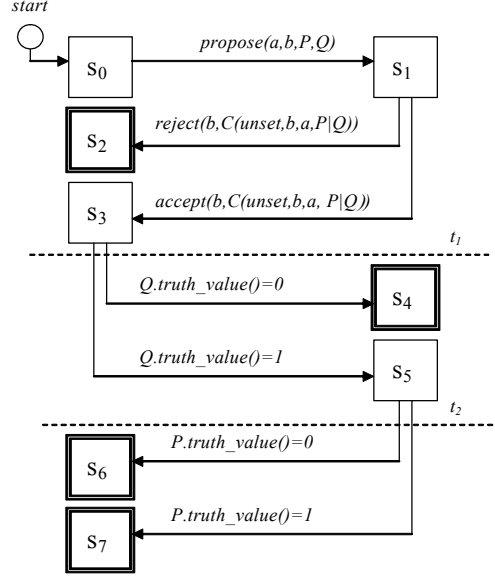


FIGURE 2.

Interaction diagram of the proposal protocol (see Table 2 for state contents).

In the presented model temporal aspects are only sketched, we plan to develop this point in the future. In future studies we plan also to extend the model in order to cover some other communicative acts, in particular orders and commands treated in hierarchical contexts. Moreover the expressiveness of the model may be increased with the addition of new methods to the commitment class to let interacting agents negotiate the content and the condition of a commitment object.

One important difference between our approach and other existing ones based on the notion of commitment (Singh 1999; Colombetti 2000; Yolum and Singh 2001; Verdicchio and Colombetti 2003) is given by its operational connotation. This is an important step for the realization of practical applications. Anyway the logic of commitment presented in (Verdicchio and Colombetti 2003) has strong connections with the model presented in this paper (for a thorough comparison see (Colombetti, Fornara, and Verdicchio 2004)). Moreover our approach is strongly related to the one proposed by Yolum and Singh (Yolum and Singh 2001). However, between the two approaches there are some significant differences. First, our proposal provides a more complete account of how different types of communicative acts can be defined in terms of operations on commitments. Second, and more important, we also provide a commitment-based analysis of directive communicative acts, like requests, relying on

TABLE 2.

Contents of the states of the interaction diagram of the proposal protocol.

$P(b \text{ do } \beta, \text{ now } \dots t_2, \exists, \perp)$ $Q(a \text{ do } \alpha, \text{ now } \dots t_1, \exists, \perp)$ $S(C_i(\text{pending}, b, a, P Q), \text{ now } \dots t_1, \exists, \perp)$ $t_1 < t_2$		
s	<i>reason of the action</i>	<i>content</i>
s_0	<i>start</i>	$C_i()$ $C_j()$
s_1	$\text{propose}(a, b, P, Q)$	$C_i(\text{unset}, b, a, P Q)$ $C_j(\text{pending}, a, b, Q S)$
s_2	$\text{reject}(b, C_i(\text{unset}, b, a, P Q))$ notifier rule 4	$C_i(\text{cancelled}, b, a, P Q)$ $S.\text{state}()=0$ $C_j(\text{cancelled}, a, b, Q S)$
s_3	$\text{accept}(b, C_i(\text{unset}, b, a, P Q))$ notifier rule 3	$C_i(\text{pending}, b, a, P Q)$ $S.\text{state}()=1$ $C_j(\text{active}, a, b, Q S)$
s_4	$Q.\text{truth_value}=0$, rule 4 rule 2	$C_i(\text{cancelled}, b, a, P Q)$ $C_j(\text{violated}, a, b, Q S)$
s_5	$Q.\text{truth_value}=1$, rule 3 rule 1	$C_i(\text{active}, b, a, P Q)$ $C_j(\text{fulfilled}, a, b, Q S)$
s_6	$P.\text{truth_value}=0$, rule 2	$C_i(\text{violated}, b, a, P Q)$ $C_j(\text{fulfilled}, a, b, Q S)$
s_7	$P.\text{truth_value}=1$, rule 1	$C_i(\text{fulfilled}, b, a, P Q)$ $C_j(\text{fulfilled}, a, b, Q S)$

the notion of unset commitment. Recently other proposals take in consideration the use of the notion of commitment, similar to the one proposed in this paper, for the specification of agent interactions (Amgoud, Maudet, and Parsons 2002; McBurney and Parsons 2003) but with respect to our proposal they mainly lack a complete and general treatment of the dynamic evolution of the state of commitments and of the basic authorizations for their manipulation.

REFERENCES

- Amgoud, L., N. Maudet, and S. Parsons. 2002. An argumentation-based semantics for agent communication languages. In F. V. Harmelen (Ed.), Proceedings of the European Conference on Artificial Intelligence (ECAI-2002), Lyon, France, pp. 38–42. IOS Press.
- Austin, J. L. 1962. *How to Do Things With Words*. Oxford: Oxford University Press.

- Bentahar, J., B. Moulin, and B. Chaib-draa. 2004. Commitment and argument network: A new formalism for agent communication. In F. Dignum (Ed.), *Advances in Agent Communication, International Workshop on Agent Communication Languages, ACL 2003, Melbourne, 2003*, Volume 2922 of *LNAI*, pp. 146–165. Springer.
- Berners-Lee, T., J. Hender, and O. Lassila. 2001, May. The Semantic Web. *Scientific American*.
- Castelfranchi, C. 1995. Commitments: from individual intentions to groups and organizations. In V. Lesser (Ed.), *Proceedings of the First International Conference on Multi-Agent Systems*, San Francisco, CA, pp. 41–48. MIT Press.
- Colombetti, M. 2000. A commitment-based approach to agent speech acts and conversations. In *Proc. Workshop on Agent Languages and Communication Policies, (Held at 'Agents 2000')*, Barcelona, pp. 21–29.
- Colombetti, M., N. Fornara, and M. Verdicchio. 2004. A social approach to communication in multiagent systems. In J. A. Leite, A. Omicini, L. Sterling, and P. Torroni (Eds.), *Declarative Agent Languages and Technologies*, Volume 2990 of *LNAI*, pp. 191–220, to be published. Springer.
- Colombetti, M. and M. Verdicchio. 2002. An analysis of agent speech acts as institutional actions. In *Proc. First International Joint Conference on Autonomous Agents and MultiAgent Systems AAMAS 2002*, New York, pp. 1157–1164. ACM Press.
- Dignum, F. (Ed.) 2003. *Advances in Agent Communication, International Workshop on Agent Communication Languages, ACL 2003, Melbourne, 2003*, Volume 2922 of *LNAI*. Springer.
- Esteva, M., J. Rodriguez-Aguilar, C. Sierra, P. Garcia, and J. L. Arcos. 2001. On the formal specification of electronic institutions. *Lecture Notes in Computer Science 1991*, 126–147.
- Finin, T., Y. Labrou, and J. Mayfield. 1997. KQML as an agent communication language. In J. M. Bradshaw (Ed.), *Software Agents*, Chapter 14, pp. 291–316. AAAI Press / The MIT Press.
- Fornara, N. and M. Colombetti. 2002. Operational specification of a commitment-based agent communication language. In *Proc. First International Joint Conference on Autonomous Agents and MultiAgent Systems AAMAS 2002*, pp. 535–542. ACM Press.
- Fornara, N. and M. Colombetti. 2003. Defining interaction protocols using a commitment-based agent communication language. In *Proc. Second International Joint Conference on Autonomous Agents and MultiAgent Systems AAMAS 2003*, pp. 520–527. ACM Press.
- Fornara, N., F. Viganò, and M. Colombetti. 2004. Agent communication and institutional reality. In *Proc. of the AAMAS 2004 Workshop on Agent Communication*, pp. to be published.
- Foundation for Intelligent Physical Agents. 1997, October. FIPA 97 specification part 2: Agent communication language. Version 2.0.
- Foundation for Intelligent Physical Agents. 2002. FIPA Communicative Act Library Specification. <http://www.fipa.org>.
- Foundation for Intelligent Physical Agents. 2004. FIPA Content Language (CL) Specifications. <http://www.fipa.org>.
- Labrou, Y., T. Finin, and Y. Peng. 1999. Agent communication languages: The current landscape. *IEEE Intelligent Systems* 14 (2), 45–52.

- Mayfield, J., Y. Labrou, and T. Finin. 1995. Desiderata for agent communication languages. In *AAAI Spring Symposium on Information Gathering from Heterogeneous, Distributed Environments*.
- McBurney, P. and S. Parsons. 2003. The posit spaces protocol for multi-agent negotiation. In F. Dignum (Ed.), *Advances in Agent Communication, International Workshop on Agent Communication Languages, ACL 2003, Melbourne, 2003*, Number 2922 in LNAI, pp. 364–382. Springer Verlag.
- Pitt, J. and A. Mamdani. 1999. A protocol-based semantics for an agent communication language. In D. Thomas (Ed.), *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI-99-Vol1)*, S.F., pp. 486–491. Morgan Kaufmann Publishers.
- Searle, J. 1979. *Expression and Meaning*. Cambridge, United Kingdom: Cambridge University Press.
- Searle, J. R. 1969. *Speech Acts: An Essay in the Philosophy of Language*. Cambridge, United Kingdom: Cambridge University Press.
- Searle, J. R. 1991. Conversation. In H. Parret and J. Verschueren (Eds.), *(On) Searle on Conversation*, pp. 7–29. Amsterdam: J. Benjamins Publishing Company.
- Searle, J. R. 1995. *The construction of social reality*. New York: Free Press.
- Singh, M. P. 1998. Agent communication languages: Rethinking the principles. *IEEE Computer* 31 (12), 40–47.
- Singh, M. P. 1999. A social semantics for agent communication languages. In *Proceedings of IJCAI-99 Workshop on Agent Communication Languages*, pp. 75–88.
- Verdicchio, M. and M. Colombetti. 2003. A logical model of social commitment for agent communication. In *Proc. Second International Joint Conference on Autonomous Agents and Multi-Agent Systems AAMAS 2003*, pp. 528–535. ACM Press.
- Viroli, M., A. Ricci, and A. Omicini. 2004. A semantics for the interaction of agents with coordination artifacts. In R. Trappl (Ed.), *Cybernetics and Systems 2004*, Volume 2, Vienna, Austria, pp. 564–569. Austrian Society for Cybernetic Studies.
- Wooldridge, M. 1998. Verifiable Semantics for Agent Communication Languages. In Y. Demazeau (Ed.), *Proceedings of the Third International Conference on Multi-Agent Systems (ICMAS'98)*, Paris, France, pp. 349–356. IEEE Press.
- Yolum, M. P. and P. Singh. 2001. Synthesizing finite state machines for communication protocols. Technical Report TR-2001-06, Department of Computer Science, North Carolina State University.