

A Logical Model for Agent Communication Languages

Mario Verdicchio¹

*Department of Electronics and Information
Politecnico di Milano
Milan, Italy*

Marco Colombetti²

*Politecnico di Milano
Milano, Italy
University of Lugano
Lugano, Switzerland*

Abstract

The current ACL proposals show some shortcomings with respect to the definition of their semantics. Our paper aims at tackling those issues by defining an ACL semantics as a specification of the analytical effects of agent communicative acts. We analyze agent communication in terms of concepts taken from Speech Act Theory, as several researchers have already done, but move away from the mainstream view of artificial agent research, as we define communicative acts in terms of changes at the level of social relationship between agents. We take commitment to be a primitive concept underlying the social dimension of multiagent systems, and define a basic artificial institution whose aim is to provide agents with the means to affect the commitment network that binds them to each other.

1 Introduction

Our work aims at defining a sound and effective framework for communication between agents in open multiagent systems. To date, all major proposals have shared the assumption that agent communication is to be dealt with in terms of speech acts, a notion that comes from philosophy of language. We think that this approach is so widespread mainly because AI research often deals with planning and rational action, and considering communication as a

¹ Email: mario.verdicchio@polimi.it

² Email: marco.colombetti@polimi.it

form of action makes it natural to embed a communicative component in an agent’s architecture. Our work has been developed from this standpoint, but we intend to tackle some issues that have not been dealt with in a satisfactory way, possibly hindering the rise of a universally accepted ACL standard. In a previous work [2], we have highlighted some shortcomings in the definition of communicative acts in FIPA ACL. We give FIPA the credit for separating the utterance level of a communicative act (the process of sending a message) from the level of illocution (the mental states characterizing the sender). Nevertheless, FIPA’s definition refers to agents’ mental states, and thus fails to grasp the institutional nature of illocution, which is where our work kicks off from. Instead of dealing with mental states, we analytically define communicative acts in terms of changes at the level of social relationship between agents. We take *commitment* to be a primitive concept that underlies the social structure of a multiagent system, and describe communicative acts as actions brought about by an agent to affect the network of commitments that bind the agent to the others. Agents communicate by performing message exchanges, which *count as* communicative acts only if some particular conditions hold. These conditions include the authorizations the agents must be granted to be able to perform such acts. We then need to define *artificial institutions* that not only are responsible for issuing and managing such authorizations, but also affect communication by determining the social relationship between the communicating agents. For instance, for an order issued to agent y by agent x to be felicitous, the two agents must be part of a hierarchical organization, empowering x to give y orders: such a hierarchy provides for a suitable institutional ground for the execution of orders. By the term *institution* we mean a set of shared rules that regulate the management of a fragment of social reality ([9],[6]). From our perspective, we assume that to play a part in multiagent systems, an agent must be officially recognized as a member of the “society of agents”. To achieve this, the agent has to undergo a suitable registration procedure, which enables it to interact with other registered agents according to the *Basic Institution*, that is, the institution that sets the general rules of agent interaction.

In Section 2 we give a brief description of the formalism we use to describe the communicative acts that agents bring about. Agents are enabled to do so as they are members of the Basic Institution, whose structure is described in Sections 3 to 6. Section 3 is dedicated to the *core ontology*, which describes commitments, the actions by which agents can manipulate them, the *ontological presuppositions* of these actions, and the *truth conditions* of the sentences that form the commitments’ content and determine their state. Section 4 deals with the *authorizations* that agents are granted by the Basic Institution, and Section 5 briefly tackles the issues related to the *norms* established by an institution. Section 6 illustrates the message exchanges that agents have to perform to affect the social network of commitments. Section 7 gives a complete view on how all the components of the Basic Institution affect agent

communication by representing the performance of basic communicative acts in a general form. Section 8 takes into account those communicative acts, like orders and declarations, that need further *special institutions* to be carried out successfully; we also illustrate how to execute all communicative acts in the form of a declaration. Finally, in Section 9 we draw our conclusions.

2 Communicative acts

An agent must be registered in order to interact with other registered agents within the context of the Basic Institution. We assume that all registered agents play a specific role (*RegAgt*) within the Basic Institution (*BI*). If a is a registered agent, the formula

$$Role(a, BI, RegAgt)$$

holds. Agents are able to bring about events of different types. We reify events, that is, we treat them as individuals belonging to specific event types. In the human world, events that are intentionally brought about by one or more actors are called actions. In the world of artificial agents, distinguishing between what is intentional and what is not is problematic. If we regarded agents as intentional systems, that is, as systems entertaining mental states like beliefs, desires, and intentions (see for example [12]), it would be plausible to distinguish between events and intentional actions, but we would also bring in a number of philosophical and practical problems that are far from being solved. We therefore prefer to adopt a more cautious approach, not considering artificial agents as intentional systems. We shall simply assume that the responsibility for certain events can be ascribed to specific agents. An event whose responsibility can be ascribed to a specific agent can be denoted with the term ‘action’. When agent a brings about an event (or performs an action) e of type τ , the formula

$$Done(e, a, \tau)$$

holds. For the sake of simplicity, we may use the ‘n-dash’ character to express existential quantification, as in the examples below:

$$\begin{aligned} Done(e, -, \tau) &\triangleq \exists x Done(e, x, \tau); \\ Done(e, -, -) &\triangleq \exists x \exists \tau Done(e, x, \tau). \end{aligned}$$

To illustrate the axioms dealing with events, we enrich our *Semantic Language* (SL) with CTL*-like operators [3], relying on an intuitive understanding of these operators. We assume that the logical model relies on a discrete-time frame with a tree-like structure, infinite both in the future and in the past, in which every state has a unique predecessor and at least one successor, and there is at most one walk between any pair of states. In such a frame, a *path* is a sequence of states, each of which is the predecessor of the following one. A path is infinite only in the future, as it has a starting state without a predecessor. We assume that on every path an event can happen only once, as

stated by the following axiom:

$$(UE) \quad Done(e,-,-) \rightarrow G^- X^- \neg Done(e,-,-) \wedge AG^+ X^+ \neg Done(e,-,-).$$

Here we use temporal operators, whose intuitive meanings are as follows:

- A for all paths;
- G^+ always on the current path, at the current state and in the future;
- G^- always on the current path, at the current state and in the past;
- X^+ at the next state on the current path;
- X^- at the previous state on the current path.

A formal account of the temporal logic in which we embed our communication framework can be found in [11].

We now consider *communicative acts*, that is, actions that agents bring about to communicate with other agents. We regard them as institutional acts performed by way of exchanging messages. In particular, we focus our attention on those events that agents are authorized to perform by the Basic Institution.

We define the semantics of our communication framework as a specification of the analytical effects of events of message exchange. Past proposals like FIPA's relate these effects to the agents' mental states, and thus, in our opinion, raise some problems.

To define communicative acts independently of the message level without relying on mental states, we adopt an approach dealing with *commitments*, already followed in other scientific papers like [10], [1], and [13]. As commitments are part of the social reality defined and regulated by the Basic Institution, we give a more detailed description of this institution to illustrate how it enables agents to perform communicative acts.

Some work in the direction of defining institutions in general has already been carried out in the field of multiagent systems (see for example [4]). For our purposes, we define an institution as comprised of four fundamental components:

- core ontology** the ontology of the fragment of reality the institution is intended to regulate;
- authorizations** the specification of the institutional effects that each member of the institution is empowered to bring about;
- norms** the obligations and the permissions imposed by the institution to its members;
- conventions** the specification of the concrete events that conventionally bring about institutional effects.

The Basic Institution regulates the general aspects of all kinds of agent interaction, including commitments, as we explain in the sequel.

3 Core ontology of the Basic Institution

An institution is intended to regulate a fragment of social reality; its core ontology defines the set of entities such fragment is comprised of.

3.1 Commitment

In our approach, the core ontology of the Basic Institution defines the concepts of *commitment*, *precommitment*, and a set of operations for *commitment manipulation*. Commitment is a primitive concept underlying the relations between agents that play the role of *RegAgt* within the Basic Institution. More precisely, a commitment holds in a state in which an agent (the *debtor*) is bound, relative to another agent (the *creditor*), to the fact that some proposition (the *content*) is true. The content of a commitment is a sentence of a *Content Language* (CL) represented as a first-order term of SL. The relevant formula is

$$Comm(e,x,y,s),$$

which states that a communicative act e has brought about a situation that binds agent x , with respect to agent y , to the truth of a proposition of the content language, represented by the SL term s . When there is a commitment proposal to an agent (normally, the potential debtor), but it has not been accepted nor refused yet, we say that a *precommitment* holds. Precommitments are represented in the same way as commitments:

$$Prec(e,x,y,s)$$

holds when e has brought about a precommitment between two agents (the potential debtor x and the potential creditor y) to the truth of a CL sentence represented by s . We think that the creation (and the modification) of (pre)commitments are those effects of communicative events that have to be dealt with to define an effective ACL semantics.

Defining the truth conditions for sentences in our model is not trivial, mainly because of the branching structure of time. This issue and the relation between SL and CL are explored in more detail in Section 3.4.

3.2 Commitment manipulation

In our treatment, both commitments and precommitments arise from the performance of communicative acts. More precisely, agents bring about communicative events by exchanging messages; a communicative event, under given conditions, counts as a *commitment manipulation action*, which creates a new (pre)commitment or modifies the state of an existing one, thus affecting the relations between agents and, if we view commitment as a deontic state, their obligations towards each other. The core ontology of the Basic Institution allows for five basic operations for commitment manipulation. Commitments can be *made* or *cancelled*, and precommitments can be *made*, *can-*

celled, or *accepted*, that is, turned into actual commitments. Registered agents achieve these results by performing actions of the following types: make commitment ($mc(x,y,s)$), make precommitment ($mp(x,y,s)$), cancel commitment ($cc(e,x,y,s)$), cancel precommitment ($cp(e,x,y,s)$), and accept precommitment ($ap(e,x,y,s)$). Such action types are defined by axioms that describe their constitutive effects, that is, the state of affairs that necessarily hold after a token of the given action type is performed.

We need to introduce some more temporal operators to express axioms in a simpler form:

$\phi\mathbf{U}^+\psi$ (ϕ is true until ψ is eventually true);
 $\phi\mathbf{W}^+\psi \triangleq \mathbf{G}^+\phi \vee \phi\mathbf{U}^+\psi$ (weak until operator);
 $\phi\mathbf{Z}^+\psi \triangleq \phi\mathbf{W}^+\psi \wedge \mathbf{G}^+(\psi \rightarrow \mathbf{G}^+\neg\phi)$ ($\phi\mathbf{Z}^+\psi$ is true if and only if in the future ψ never becomes true and ϕ is always true, or ϕ is true until ψ eventually becomes true and since then ϕ is no longer true).

We are now ready to illustrate the axioms that describe the commitment manipulation actions.

$$(MC) \quad \mathbf{Done}(e,-,mc(x,y,s)) \rightarrow \mathbf{A}(\mathbf{Comm}(e,x,y,s)\mathbf{Z}^+\mathbf{Done}(-, -, cc(e,x,y,s))).$$

Axiom MC (Make Commitment) states that:

- if an agent (not necessarily x or y) performs an action of making a commitment with x as the debtor, y as the creditor, and s as the content,
- then on all paths x is committed, relative to y , to the truth of s ,
- until an agent possibly cancels such a commitment, after which the commitment no longer exists.

Axiom MC only defines the meaning of making a commitment, without specifying in what way and under what conditions an agent can actually create or cancel a commitment, which is dealt with in Section 4.

$$(MP) \quad \mathbf{Done}(e,-,mp(x,y,s)) \rightarrow \mathbf{A}(\mathbf{Prec}(e,x,y,s)\mathbf{Z}^+(\mathbf{Done}(-, -, ap(e,x,y,s)) \vee \mathbf{Done}(-, -, cp(e,x,y,s)))).$$

Axiom MP (Make Precommitment) is analogous to MC.

$$(AP) \quad \mathbf{Done}(e', -, ap(e,x,y,s)) \rightarrow \mathbf{A}(\mathbf{Comm}(e',x,y,s)\mathbf{Z}^+\mathbf{Done}(-, -, cc(e',x,y,s))).$$

Axiom AP (Accept Precommitment) implies that if an agent performs an action of accepting a precommitment brought about by event e with x , y , and s respectively as debtor, creditor, and content, then the acceptance action brings about on all paths a commitment of x , relative to y , to the truth of s , which will stand until it is possibly cancelled. Again, the axiom does not show the conditions under which an agent can accept a precommitment in an actual case.

There are no specific axioms for the actions of cancelling a precommitment (cp) or a commitment (cc), because the analytical effects of these commitment manipulations are already illustrated in the axioms dealing with other actions, whose performance must be presupposed in order to take into account any kind

of cancellation. This issue is dealt with in more detail in the next section.

3.3 *Ontological presuppositions*

It is important to remark that some of the commitment manipulation actions rely on *ontological presuppositions*, that is, they can be performed only if particular states of affairs hold. In fact, only if a precommitment actually exists, it can be refused by the potential debtor, or cancelled by the potential creditor. Moreover, it can be accepted by the potential debtor only if it is not contemporarily cancelled by the potential creditor. Similarly, the ontological presupposition to a commitment's cancellation is the existence of the commitment itself. Formula

$Poss(\tau)$

states that an event of type τ is ontologically possible. The definition of these ontological presuppositions is also part of the core ontology. We assume that making a commitment or a precommitment is always ontologically possible, as specified by the following axioms:

- (PMC) $Poss(mc(x,y,s))$,
- (PMP) $Poss(mp(x,y,s))$.

The axioms dealing with the presuppositions to the actions of cancelling and accepting are the following:

- (PCC) $Poss(cc(e,x,y,s)) \leftrightarrow X^- Comm(e,x,y,s)$;
- (PCP) $Poss(cp(e,x,y,s)) \leftrightarrow X^- Prec(e,x,y,s)$;
- (PAP) $Poss(ap(e,x,y,s)) \leftrightarrow X^- Prec(e,x,y,s) \wedge \neg Done(-, -, cp(e,x,y,s))$.

The core ontology also describes the states (*fulfilled, violated, pending*) in which a commitment can be. The definition of these states relies on the *truth conditions* of the content of the commitment, which are formally described in the next section.

3.4 *Sentence meaning and truth conditions*

Before we define the truth conditions of CL sentences, two remarks should be made. First, the truth of a temporal sentence at a given state (the *point of reference*, [7]) can be evaluated only if we know at which state the sentence has been uttered (the *point of speech*). For example, the sentence “I shall pay you within the end of the month” implicitly refers to the end of the current month, which in turn is determined by the state at which the sentence is uttered. Second, branching time brings in a phenomenon known as *contingent future*, which means that at a given point of reference it may be still undetermined if a sentence is going to be true or false. Consider again to the previous example, and assume that the sentence has been uttered, on January 10th, and that no payment has been made as far as January 15th; on January 15th, the sentence is still not settled true nor settled false, and thus it is undefined. Note however

that a sentence that is true (false) at a state, will go on being true (false) at all states in the future of that state. We represent CL sentences as SL terms, which allows us to define CL semantics in SL. More precisely, we first assume that for every SL term s that denotes a CL sentence, there is exactly one SL formula $[s]$ that corresponds to s , which we call the *sentence meaning* of s ; then we define the truth conditions of s in SL.

CL semantics is dealt with by means of the following predicates, whose definitions we call *truth conditions* of a sentence:

$$\begin{aligned} True(e,s) &\leftrightarrow \mathbf{AF}^-(Done(e,-,-) \wedge [s]), \\ False(e,s) &\leftrightarrow \mathbf{AF}^-(Done(e,-,-) \wedge \neg [s]), \\ Undef(e,s) &\leftrightarrow \mathbf{AF}^- Done(e,-,-) \wedge \neg True(e,s) \wedge \neg False(e,s). \end{aligned}$$

Note that the \mathbf{A} path quantifier in front of the \mathbf{F}^- operator is necessary, even if our model structure is not branching in the past, because formula $[s]$ may include operators like \mathbf{F}^+ that need a path to be specified. The truth conditions of sentence s are given with respect to an event e , which does not necessarily correspond to the event of uttering s . Event e is used to set a well-defined temporal reference by which we can evaluate the truth of s . We may then analyze the *logic* of CL-sentences, but the topic lies beyond the scope of this paper. The truth conditions of a CL sentence determine the fulfillment or the violation of the relevant commitment. More precisely, a commitment whose content is s is said to be *fulfilled*, *violated*, or *pending* respectively when s is true, false, or undefined according to the above definitions. The event with respect to which the truth conditions of the content are checked is the one that has brought about the commitment. Here are the axioms that formalize what stated above:

$$\begin{aligned} Fulf(e,x,y,s) &\leftrightarrow Comm(e,x,y,s) \wedge True(e,s), \\ Viol(e,x,y,s) &\leftrightarrow Comm(e,x,y,s) \wedge False(e,s), \\ Pend(e,x,y,s) &\leftrightarrow Comm(e,x,y,s) \wedge Undef(e,s). \end{aligned}$$

Intuitively, every commitment is either fulfilled, or violated, or pending. It is actually possible to prove that

$$\models Comm(e,x,y,s) \rightarrow \mathbf{xor}(Fulf(e,x,y,s), Viol(e,x,y,s), Pend(e,x,y,s)).$$

This means that only one of $Fulf(e,x,y,s)$, $Viol(e,x,y,s)$, or $Pend(e,x,y,s)$ is true in all models in every state in which $Comm(e,x,y,s)$ holds.

The core ontology defines the collection of entities that may exist in the context of an institution, and also the set of possible changes such entities may undergo; as a consequence, the core ontology delimits the set of all possible institutional actions. However, not every ontologically possible action can be actually carried out: it is part of the function of an institution to authorize the execution of a subset of the ontologically possible action.

4 Authorizations of the Basic Institution

Consider again the example of ordering: if we want to follow the example of human societies, in which issuing orders is not effective unless some special institutional framework (e.g.: a military hierarchy) is defined, we can define the Basic Institution so that it does not authorize agents to create commitments as creditors. We then need to introduce a predicate to state that an agent is authorized to bring about an event of a certain type. In general, to perform an action of type τ , agent x must be authorized to do so, that is, the formula $Auth(x,\tau)$

must hold. A reasonable set of authorizations concerning the creation and the manipulation of commitments can be defined in the form of an axiom as follows:

$$(ABI) \quad Role(x, BI, RegAgt) \wedge Role(y, BI, RegAgt) \rightarrow \\ Auth(x, mc(x, y, s)) \wedge Auth(x, mp(y, x, s)) \wedge Auth(x, cp(e, x, y, s)) \wedge \\ Auth(x, cp(e, y, x, s)) \wedge Auth(x, cc(e, y, x, s)) \wedge Auth(x, ap(e, x, y, s)).$$

This formula means that: as a *debtor* a registered agent is authorized to make a commitment with any registered agent as the creditor, and accept or cancel an existing precommitment; as a *creditor* a registered agent has the authorization to make precommitments with any registered agent as the debtor, and cancel an existing commitment or precommitment.

All the communicative acts that comply with the authorizations defined above are institutional actions, whose execution is authorized by the Basic Institution. On the contrary, as ordering corresponds to performing a make commitment action as a creditor, it is not authorized by the Basic Institution, and thus some further institution is required for an act of such type to be carried out.

5 Norms of the Basic Institution

Authorizations define the institutional powers of agents; in general, however, the exercise of such powers is further regulated by a set of norms. Consider, for example, a scientific society: the president of the society typically has the power to call the general meeting of the society's members; the president, however, is also obliged to call such a meeting at least once a year, and may be allowed to call it more often if he or she has good reasons to do so. At the present stage of our research, we still do not know what kind of norms should be regarded as part of the Basic Institution. Consider for example the adjacency pair made up by a request and its acceptance or refusal. At some institutional level, we might want to dictate that agents should react to all requests by producing an acceptance or a refusal. But is this rule to be regarded as a norm of the Basic Institution? Or does it belong to a special institution, like for example an "Institution of Dialogue"? We feel that more

work has to be done in order to clarify this issue; in particular, practical applications will have to be analyzed to get a better understanding of the systems of norms involved in the functioning of real multiagent systems.

6 Conventions of the Basic Institution

An institutional action is performed through the execution of some lower level act that conventionally counts as a performance of the institutional action; obvious examples are offered by communicative acts, which are performed by executing lower level acts of message exchange. As a consequence, institutional actions require a set of conventions for their execution. In our approach, the institutional actions of the Basic Institution are commitment-manipulations actions, conventionally realized by the exchange of ACL messages, as described in the following axioms. In other words, these axioms define a library of basic communicative acts in terms of commitment manipulation.

In Section 4 we have illustrated the six authorizations that registered agents are granted by the Basic Institution. Now we have to deal with the structure of the messages that agents exchange in order to perform those authorized commitment manipulation actions. We view a message as a pair made up by a *type indicator* and a *body*. Type indicators (corresponding to KQML's performatives [5]) are constant symbols taken from a finite set whose definition is part of SL. The body can be a CL sentence represented in our semantic language by a first-order term. In the case of acceptance or refusal messages, the body is comprised of a more complex structure, that is, a tuple of elements $(\langle e, x, y, s \rangle)$ that identifies an existing (pre)commitment. For every message type we need to introduce a functor that specifies the relevant type of the action that an agent performs when exchanging a message of such kind. This approach is best explained in terms of an example. Suppose that agent x sends a message to agent y to inform y that σ is the case. The exchange of such a message is an event of type $inform(x, y, s)$, where $inform$ is a three-place functor denoting the type of the message, x and y denote the sender and the receiver of the message respectively, and s is a term corresponding to SL formula σ . When event e is an exchange of a message of type $inform$ and content s , sent by agent x to agent y , the formula

$$Done(e, x, inform(x, y, s))$$

holds. This event, under given conditions, implies the performance of a commitment manipulation action. In other words, the meaning of the message is defined as the effect that exchanging such a message has on the network of commitments binding the sender and the receiver. The correspondence between the message exchange event type and the commitment manipulation action type is defined by a convention of the relevant institution, and is formally stated by means of the formula

$$CountAs(\tau, \tau'),$$

which means that an action of type τ conventionally counts as an action of type τ' . Here we define the communicative acts by means of which agents carry out the commitment manipulation actions authorized by the Basic Institution.

Informing is defined as committing to the truth of the message body, which, we suppose, is comprised of an arbitrary CL sentence. More precisely, when agent x exchanges with agent y an *inform* message with content s , agent x commits, relative to y , to the truth of s :

$$(CAInf) \quad CountAs(inform(x,y,s),mc(x,y,s)).$$

We assume that the body of a *request* message is comprised of an *action expression*, which indicates the requested action's type, its actor, and possibly a temporal constraint. More precisely, in the request message we use a term that represents the abstract syntax of an action expression, which is not to be confused with its concrete form, which belongs to a specific CL. In our SL, an action expression may have the form $Done(x,\tau)\mathbf{B}^+\Gamma$, in which \mathbf{B}^+ means intuitively “before” ($\phi\mathbf{B}^+\psi \triangleq \neg(\neg\phi \mathbf{U}^+ \psi)$), and Γ is an SL formula referring to a particular time-point. Here we are assuming that for every time-point expression of a CL (e.g. “Wednesday”, “Christmas”, etc. if we consider a CL using English terms referring to time) there exists an SL formula which becomes periodically true according to the time-point it is indicating (the period is one week for “Wednesday”, one year for “Christmas”, and so on). Thus, in this case the body of a request message is comprised of the term $[Done(x,\tau)\mathbf{B}^+\Gamma]$, in which $\lceil \cdot \rceil$ is a function which, given an SL formula ϕ , returns the relevant SL term $\lceil \phi \rceil$. If we denote the term above with s , and then define the type $request(x,y,s)$ to denote events by which agent x requests s from agent y , the semantics of *request* messages is conventionally defined as below:

$$(CAReq) \quad CountAs(request(x,y,s),mp(y,x,s)).$$

The above-mentioned action expression is only an example that works in particular cases. Several more action expressions can be defined to meet different needs, using more complex temporal operators, but we think that such a task is to be tackled only when we deal with the application of our framework to actual cases.

The act of *accepting* is not only defined with respect to requests, but with respect to precommitments in general. We assume that the body of an acceptance message is a tuple which includes all the elements that uniquely identify the relevant precommitment. To denote the event types corresponding to such message exchange we introduce the functor $accept(e,x,y,s)$, whose arguments are the same as those characterizing the precommitment that the sender of the message is accepting, and the relevant convention is as follows:

$$(CAAcc) \quad CountAs(accept(e,x,y,s),ap(e,x,y,s)).$$

The Basic Institution allows registered agents to *cancel* the commitments in which they play the role of the creditor and the precommitments in which

they are the debtor or the creditor. The body of a *cancel* message is supposed to be the same as that of an *accept* message, that is, a tuple which includes all the elements that uniquely identify the relevant (pre)commitment. To denote the event types corresponding to such message exchange we introduce the functor $cancel(e,x,y,s)$, whose arguments are the same as those characterizing the (pre)commitment which is cancelled by the relevant message exchange. The conventions are then defined like this:

$$\begin{aligned} (\text{CACancP}) \quad & \text{CountAs}(cancel(e,x,y,s),cp(e,x,y,s)), \\ (\text{CACancC}) \quad & \text{CountAs}(cancel(e,x,y,s),cc(e,x,y,s)). \end{aligned}$$

A *cancel* message exchange can count as different commitment manipulation actions (*cp* and *cc*), in accordance to what is being cancelled (a precommitment or a commitment, respectively). There is no ambiguity in an actual *cancel* message exchange, as there cannot exist both a precommitment and a commitment with the same arguments, and only an action that cancels an existing object can be successfully carried out, as stated by Axioms PCC and PCP in Section 3.3. These ontological presuppositions contribute to the general definition of institutional actions, as illustrated in the next section.

7 The general representation of communicative acts

We are now ready to give a general definition of communicative acts. The *conventions* of the Basic Institution establish that exchanging a message of given type *counts as* a specific institutional action, provided certain conditions hold. These conditions can be classified in two categories: *ontological presuppositions*, defined by the *core ontology*, and *authorizations*. While authorizations deal with the institutionalized power of agents, ontological possibility concerns the state of affairs that must hold for a communicative act to be possible. All these points are formally expressed by the following axiom that gives a general definition of institutional actions:

$$(\text{IA}) \quad Done(e,x,\tau) \wedge CountAs(\tau,\tau') \wedge Poss(\tau') \wedge Auth(x,\tau') \rightarrow Done(e,x,\tau').$$

As an example, let us consider the act of informing. Suppose that the formula $Done(e_1,a,inform(a,b,s_1))$

holds, that is, agent a informs b that s_1 is the case. From Axiom CAInf we derive $CountAs(inform(a,b,s_1),mc(a,b,s_1))$, and thus we determine the commitment manipulation action that corresponds to such a message exchange. We suppose that formulae

$$\begin{aligned} & Role(a,BI,RegAgt) \text{ and} \\ & Role(b,BI,RegAgt) \end{aligned}$$

hold, that is, both a and b are registered agents in the Basic Institution. Given these premises and Axiom ABI, which illustrates what actions registered agents are authorized to perform, we have (among other consequences)

$Auth(a, mc(a, b, s_1))$.

Axiom PMC states that it is always ontologically possible to perform a make commitment action, thus, from the premises above and Axiom IA we can derive

$Done(e_1, a, mc(a, b, s_1))$,

which, thanks to Axiom MC, gives

$Comm(e_1, a, b, s_1)$.

8 Special institutions, declarations, and performatives

This section deal with communicative acts, like orders and declarations, which cannot be performed thanks to the Basic Institution only.

8.1 Special institutions

Ordering is a kind of communicative act that falls out of the context defined by the Basic Institution. We can consider orders like requests that cannot be refused. Following our approach, a request brings about a precommitment, while an order directly creates a commitment in which the addressee of the message is the debtor. Given these premises, analogously to the definition of requests, we may take also the body of an order message to be an action expression, and postulate that:

(CAOrd) $CountAs(order(x, y, s), mc(y, x, s))$.

However, the Basic Institution does not authorize agents to directly create commitments of which they are the creditor. This means that even if agent a exchanges a message with agent b containing an order to s_2 , that is, even if

$Done(e_2, a, order(a, b, s_2))$

holds, there exists no axiom that we can use to derive

$Done(e_2, a, mc(b, a, s_2))$.

To introduce orders in our communication framework, however, we can define a special institution that grants agents that play particular roles in it the authorization to give other agents commands. Let us call such an institution SI , and suppose that it introduces the following authorization axiom:

(ASI) $Role(x, SI, high) \wedge Role(y, SI, low) \rightarrow Auth(x, mc(y, x, s))$.

Axiom ASI states that agents that play the *high* role in SI are authorized to issue orders to agents that play the *low* role within the same institution. Now, if we assume that agent a and agent b play roles *high* and *low* in SI respectively, we can derive

$Done(e_2, a, mc(b, a, s_2))$, and then, thanks to Axiom MC,

$Comm(e_2, b, a, s_2)$.

8.2 *Declarations*

As is well known, a *declaration* is a speech act that, under appropriate conditions, makes its content true. The content of a declaration must represent an institutional fact of some sort, and the agent that makes the declaration has to be authorized to bring about such an institutional fact; besides, it must be ontologically possible to bring about the fact. As an example, consider the act of opening a society's meeting; such an act can be performed by the president of the society by making a suitable declaration. Indeed, the fact that a meeting is open is institutional; the president of the society is authorized to open the meeting; and finally, to successfully open a meeting by declaration, it must not be the case that the meeting is already open. All this presupposes a special Institution of Associations, which in particular will include a definition of meetings in its core ontology, and authorize the president to open a meeting. Additional norms to regulate the exercise of institutional powers can also be defined. With respect to conventions, in our opinion, declarations can be regarded as a universal convention for the performance of all sorts of institutional actions. The logical definition of declarations is particularly interesting, because it states that

$$(CADecl) \quad CountAs(declare(\tau),\tau).$$

This definition implies that declaring an action of type τ counts as the actual performance of an action of type τ , provided such a performance is ontologically possible and the actor is authorized to perform actions of type τ , as stated below:

$$Done(e,x,declare(\tau)) \wedge Auth(x,\tau) \wedge Poss(\tau) \rightarrow Done(e,x,\tau).$$

8.3 *The performative execution of communicative acts*

It is a remarkable fact that, once declarations are defined, it becomes possible to realize all communicative acts as declarations. Suppose for example that agent x exchanges with agent y a specific message of type *declare* whose body, expressed as a suitable content language sentence, means "I commit to the truth of '2 + 2 = 4' ". The formula

$$Done(e,x,declare(mc(x,y,'2 + 2 = 4')))$$

represents such a message exchange. We derive

$$Done(e,x,mc(x,y,'2 + 2 = 4'))$$

from our premises and a number of axioms, and then, by Axiom MC, we have that

$$Comm(e,x,y,'2 + 2 = 4')$$

holds.

We conclude that the exchange of the declaration message described above has the same effect as the exchange of a message of type *inform*, as analyzed

in Section 7. This means that messages of type *inform* are not strictly necessary, because the same result can be obtained by a declaration, performed in the context of the Basic Institution. Given that the same line of reasoning can be applied to all types of messages, it turns out that all communicative acts can be realized through the use of a single type of messages, namely declaration messages. Carrying out a communicative act by declaration corresponds to the well-known *performative execution* of the communicative act [8]. In our communication framework, all communicative acts that boil down to commitment-manipulation action can be carried out in performative form, thanks to the authorizations granted by the Basic Institution. This fact has an important practical side: it is possible to define a full ACL starting from one single type of messages, that is, declaration messages; all other types of messages can then be regarded as useful abbreviations of declaration messages, but do not increase the expressive power of the ACL.

9 Conclusions

In this work we have presented a framework for agent communication that is being developed to deal with the current ACL proposals's shortcomings. In our approach, we have assumed that agent communication is carried out by means of the performance of communicative acts; such acts are performed by exchanging conventional messages, and have institutional effects defined in the context of given institutions. We take the social structure of a multi-agent system to rely on commitments, whose creation and manipulation are regulated by the Basic Institution. All agents, in that they are registered as such at this institution, interact according to its rules, that is, their message exchanges affect the network of commitments that bind agents to each other. The Basic Institution not only defines the set of entities the social structure comprises, but also specifies the institutional effects agents are authorized to bring about, the conventions that agents must follow and the preconditions that must hold so that such effects can be achieved, and possibly the permissions and the obligations agents are imposed. The Basic Institution enables agents to carry out several types of communicative acts, such as informing or requesting, but it is not enough for dealing with orders and declarations, whose felicitous performance needs further special institutions that provide for suitable hierarchical organizations and additional authorizations. In particular, we have shown that all kinds of communicative act can be carried out in the form of a declaration, that is, they can be executed in a performative way. Our future work will mainly deal with special institutions, to tackle role compatibility issues that may rise from the possibility of having an agent registered in one or more institutions other than the basic one.

References

- [1] Colombetti, M., “A Commitment-Based Approach to Agent Speech Acts and Conversations”, Proceedings of the Workshop on Agent Languages and Communication Policies, 4th International Conference on Autonomous Agents (*Agents 2000*) (2000), 21–29.
- [2] Colombetti, M., and M. Verdicchio, “An Analysis of Agent Speech Acts as Institutional Actions”, Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems (*AAMAS 02*) (2002), (3)1157–1166.
- [3] Emerson, E. A., and J. Y. Halpern, ‘*Sometimes*’ and ‘*Not Never*’ Revisited: On Branching Versus Linear Time Temporal Logic, *Journal of the ACM* **33** (1986), (1)151–178.
- [4] Esteva, M., J. A. Rodriguez, C. Sierra, P. Garcia, J. L. Arcos, “On the Formal Specifications of Electronic Institution”, F. Dignum and C. Sierra, eds., Agent-mediated Electronic Commerce (The European AgentLink Perspective), *LNAI 1991*, Springer (2001), 126–147.
- [5] Finin, T., Y. Labrou, and J. Mayfield, “KQML as an Agent Communication Language”, J. Bradshaw, ed., *Software Agents*, MIT Press (1995).
- [6] Jones, A., and M. J. Sergot, *A Formal Characterisation of Institutional Power*, *Journal of the IGPL* **4** (1996), (3)429–445.
- [7] Reichenbach, H., “*Elements of Symbolic Logic*”, MacMillan, New York (1947).
- [8] Searle, J. R., and D. Vanderveken, “*Foundations of Illocutionary Logic*”, Cambridge University Press, Cambridge (1985).
- [9] Searle, J. R., “*The Construction of Social Reality*”, Free Press, New York (1995).
- [10] Singh, M. P., *Agent Communication Languages: Rethinking the Principles*, *IEEE Computer* **31** (1998), 40–47.
- [11] Verdicchio, M., and M. Colombetti, “A Logical Model of Social Commitment for Agent Communication”, Proceedings of the Second International Joint Conference on Autonomous Agents and Multiagent Systems (*AAMAS 03*) (2003), to appear.
- [12] Wooldridge, M., “*Reasoning about rational agents*”, MIT Press, Cambridge (2000).
- [13] Yolum, P., and M. P. Singh, “Commitment Machines”, J. J. Meyer and M. Tambe, eds., *Intelligent Agents VIII*, 8th International Workshop, ATAL 2001, *LNCS 2333*, Springer (2002), 235–247.