A Semantics for Agent Communication Languages based on commitments and penalties

Leila Amgoud   Florence Dupin de Saint-Cyr

Institut de Recherche en Informatique de Toulouse (I.R.I.T.)–C.N.R.S.
Université Paul Sabatier, 118 route de Narbonne,
31062 Toulouse Cedex 4, France
{amgoud, bannay}@irit.fr

Abstract. In complex multi agent systems, the agents may be heterogeneous and possibly designed by different programmers. Thus, the importance of defining a standard framework for agent communication languages (ACL) with a clear *semantics* has been widely recognized. The semantics should meet two important objectives: *verifiability* and *flexibility*. Classical proposals (*mentalistic semantics* and *social semantics*) fail to meet these objectives.
In this paper we propose a semantics for ACL which is both verifiable and flexible. That semantics is social in nature since it is based on commitments. Two kinds of commitments can be distinguished: i) the commitments made by the agent itself during the dialogue such as promises, and ii) the commitments made by the others to the agent such as requests or questions. To these commitments, we associate two kinds of penalties for sanctioning the agents which do not respect its commitments. The first kind of penalty ensures that the agent is honest whereas the second kind of penalty ensures that the agent is cooperative. These penalties make the semantics verifiable.
Moreover, our semantics is flexible since it is not given on the basis of particular speech acts, but on the well-known *categories* of speech acts identified by Searle in [11, 12]. This makes it more general than the existing ones.

Key words: Agent Communication Languages, Commitment, Penalties.

1 Introduction

When building multi-agent systems, we take for granted the fact that the agents which make up the system will need to communicate and to engage in the different types of dialogue identified by Walton and Krabbe in [15], using a communication language (ACL).
The definition of an ACL from the syntactic point of view (the different *speech acts*[1] that agents can perform during a dialogue) poses no problems. The situation is different when semantics is taken into account. Any communication

[1] The speech acts are also called elsewhere *illocutionary* acts or *performatives*.

language must have a well-defined semantics. Given that agents in a multi-agent system may be independently designed by different programmers, a clear understanding of semantics is essential. Moreover, the semantics should meet two important *properties*: *verifiability* and *flexibility*.
Although a number of agent communication languages have been developed, obtaining a suitable formal semantics for ACLs which satisfies the above objectives remains one of the greatest challenges of multi-agent theory.
There are mainly two categories of semantics: the *mentalistics* semantics and the *social* ones. The mentalistic semantics, used in KQML [7] and FIPA [8], is based on a notion of speech act close to the concept of illocutionary act as developed in speech act theory [4, 12]. The basic idea behind this semantics is to define the conditions under which a given speech act can be played. Unfortunately, the conditions are based on the mental states of the agents and this makes the semantics not verifiable since one cannot access to those mental states to check whether the conditions are really satisfied or not. Consequently, such category of approaches violates one of the important properties of a semantics.

In the second category of semantics, called *social* and developed in [5, 13, 14], primacy is given to the interactions among the agents. The semantics is based on social *commitments* brought about by performing a speech act. For example, by affirming a data, the agent commits on the truth of that data. After a promise, the agent is committed carrying it out. There are several weak points of this approach and the most relevant one is the fact that the concept of commitment itself is ambiguous and its semantics is not clear. According to the performative act, the semantics of the commitment differs. (See section 5 for more details)

Another limitation of both approaches is the fact that they are neither modular nor general. Since they are defined on particular speech acts, if a new speech act is needed for a particular dialogue, then the semantics should be extended.

Our aim is to define a semantics which prevents the shortcomings of the existing approaches while keeping their benefits. Thus, we propose a semantics which satisfies the two main properties. Our approach is social in nature since it is based on the notion of commitments. An agent may have two kinds of commitments:

1. The commitments made by the agent itself during the dialogue such as promises and that it should respect. Such commitments allow the evaluation of the degree of *honesty* of the agent. Indeed, an agent which respects its commitments is said *honest*.
2. The commitments for replying to requests made by the other agents. For instance if the agent receives a question from another agent, then it should answer. Such commitments ensures the *cooperativity* of the agent.

From a syntactic point of view, these commitments are modelled via a notion of *commitment store* as originally defined by MacKenzie in [9]. Indeed, each agent is supposed to be equipped with a commitment store which will keep track of the different commitments of the agent. These stores are accessible for all the agents.

The semantics of a commitment is *unique* and is given by the sanction, called *penalty*, associated to a commitment in case it is not respected by the agent.

Another interesting feature of our semantics is the fact that it is defined on the well-known *categories* of speech acts defined by Searle in [11, 12]. Indeed, Searle has defined five categories of performatives: the *assertives*, the *directives*, the *commissives* and finally the *expressives*. We will show that we can define a semantics for each category without worrying about the different speech acts it may contain. This makes our semantics flexible and more general than the existing ones.

The proposed semantics is also verifiable since the penalty associated to a move is computed directly from its category and the agent commitment store which is visible to all agents; so there is no need to know what the agent really believes, only the category of the moves and previous moves are taken into account.

The paper is organized as follows: Section 2 introduces the logical language that will be used throughout the paper. In section 3 we present the basic concepts of our semantics. In particular, we define the two kinds of commitments of an agents as well as the two corresponding penalties. In section 4, we define the semantics of the four categories of speech acts defined by Searle (we do not consider the expressive category). Section 5 compares our approach to existing semantics and finally, section 6 concludes with some remarks and perspectives.

2 Background

In this section we start by presenting the logical language which will be used throughout this paper. We will distinguish among *action variables* ($\mathcal{AV}$) and *non-actions variables* ($\mathcal{NAV}$). $\mathcal{AV}$ and $\mathcal{NAV}$ are supposed to be disjoint (i.e. $\mathcal{AV} \cap \mathcal{NAV} = \varnothing$).

$\mathcal{L}$ will denote a propositional language built from $\mathcal{AV} \cup \mathcal{NAV}$. $\vdash$ denotes classical inference and $\equiv$ denotes logical equivalence.

$Arg(\mathcal{L})$ will denote the set of all the arguments that can be constructed from $\mathcal{L}$. An *argument* is a pair (Γ, ϕ) where Γ is a set of formulas of $\mathcal{L}$, called the *support* of the argument and $\phi \in \mathcal{L}$ is its *conclusion*. In [6], Dung has presented a powerful argumentation framework which takes as input a set of arguments and the different conflicts which may exist between them, and returns among all the arguments the "good" ones, called the *acceptable* arguments.

Let $\mathcal{A} = \{a_1, \ldots, a_n\}$ denote the set of agents. A communication language is based on a set of *speech acts*. In what follows, $\mathcal{S}$ denotes the set of speech acts of a given communication language. In [2], a communication language has been defined using the following set of speech acts: {*Question*, *Challenge*, *Assert*, *Accept*, *Refuse*, *Request*, *Promise*, *Retract*, *Argue* }. *Assert* allows the exchange of information such as "the weather is beautiful" and "It is my intention to hang a picture". *Request* is invoked when an agent cannot, or prefer not to, achieve its intentions alone. The proposition requested differs from an asserted proposition in that it cannot be proved true or false. The decision on whether to accept it or not hinges upon the relation it has to the agent's intentions. An agent will make a *Promise* when it needs to request something from another, and has something it doesn't need which can offer in return. In replying to an assertion, a request or a promise, an agent can accept, refuse or ask a question such as "is it the case that p is true ?". Another kind of question is called a challenge. It allows an agent to ask another agent why it has asserted a proposition or requested something, for example "why newspapers can't publish the information X?". The answer to a challenge should be an *argument*. *Argue* allows agents to exchange arguments, and *Retract* allows them to retract propositions previously asserted or requested.

In [11, 12], Searle has identified five categories of speech acts according to what he calls their illocutionary purpose, their psychological direction to the world, their expressed state, and their propositional content. The five categories are as follows:

Assertive speech acts: The speaker claims that some proposition is true. This category contains speech acts like "inform" or "assert". For instance, *Inform: the sky is blue* or *Inform: it is my intention to hang a mirror.* In what follows, we will refer to the set of all speech acts of this category by $\mathcal{S}_1$.

Directive speech acts: The speaker attempts to get the hearer to do something. This category includes speech acts like "Request", "Question" and "Challenge". For instance, *Request: clean your room!*, *Question: is p true.* In what follows, we will refer to the set of all speech acts of this category by $\mathcal{S}_2$.

Commissive speech acts: The speaker commits to some future course of action. An example of a commissive speech act is "promise", for instance *Promise: I will do it.* In what follows, we will refer to the set of all speech acts of this category by $\mathcal{S}_3$.

Expressive speech acts: The speaker expresses some psychological state. For instance, an agent can say *Im sorry.*

Declarative speech acts: The speaker brings about a different state of the world. For instance, "Propose", "Accept", "Refuse", "Retract" are declaratives. In what follows, we will refer to the set of all speech acts of this category by $\mathcal{S}_4$.

In this paper, we are more concerned with setting a semantics for dialogues between artificial agents. Thus, the category of expressive speech acts will not be considered here.

$\mathcal{S}_1, \mathcal{S}_2, \mathcal{S}_3, \mathcal{S}_4$ is seen as a partition of the set $\mathcal{S}$ of speech acts. This means that $\forall i, j \in [1,4], i \neq j \Rightarrow \mathcal{S}_i \cap \mathcal{S}_j = \varnothing$ and $\mathcal{S}_1 \cup \mathcal{S}_2 \cup \mathcal{S}_3 \cup \mathcal{S}_4 = \mathcal{S}$. Hence, every speech act symbol belongs to one and only one of the four categories $\mathcal{S}_1$ or $\mathcal{S}_2$ or $\mathcal{S}_3$ or $\mathcal{S}_4$.

Definition 1 *A* move *is a formula $a : x$ where $a \in \mathcal{S}$ is a speech act and x is either a propositional formula ($x \in \mathcal{L}$) or an argument ($x \in Arg(\mathcal{L})$). Let $\mathcal{M}$ be the set of moves.*

For instance, $Question : sky\_blue$ is a move meaning that the agent asks whether the sky is blue or not. Here, "Question" is a speech act and sky_blue is a propositional formula (whose value is true or false).

3 The semantics

3.1 Commitment

In the scientific literature, one can find proposals where the semantics of an ACL is defined in terms of commitments. Examples of these are given by Colombetti [5] and Singh [13, 14]. The authors argued that agents are social entities, involved in social interactions, then they are committed to what they say.

In recent years, it has been argued that informal logic has much to offer to the analysis of inter-agent communication. Central in these approaches are the notions of dialogue games and (social) commitments. One rather influential dialogue game is DC, proposed by MacKenzie [9] in the course of analysing the fallacy of question-begging. DC provides a set of rules for arguing about the truth of a proposition. Each player has the goal of convincing the other, and can assert or retract facts, challenge the other player's assertions, ask whether something is true or not, and demand that inconsistencies be resolved.

Associated with each player is a *commitment store*, which holds the commitments of the players during the dialogue. Commitments here are the information given by the players during the dialogue. There are then rules which define how the commitment stores are updated. Take for instance the assertion, it puts a propositional statement in the speaker's commitment store. What this basically means is that, when challenged, the speaker will have to justify his claim. But this does not presuppose that the challenge will come at the next turn in the dialogue.

For our purpose, we adopt this presentation of commitments. In this paper we are not interested in modeling the agent's reasoning, we only consider what is said by each agent. Our purpose is to provide a semantics for the dialogue without worrying about the mental states of the agents.

Each agent is supposed to be equipped with a commitment store, accessible to all agents, which will contain its commitments made during the dialogue. The union of the commitment stores of all agents at turn t can be viewed as the state of the dialogue at turn t. We adopt a commitment store much more structured than the one presented in previous works on dialogue [1, 3]. It keeps tracks of two kinds of commitments:

– The commitments made in the dialogue by the agent itself such as assertions and promises.
– The commitments made by other agents, such as requests, challenges and questions. For instance if an agent a_i makes a request r to another agent a_j, the request (r) is stored in the commitment store of a_j. Hence, a_j is said committed to satisfy this request.

More formally, for each agent a_i, we denote CS_i its commitment store, containing statements and issues to be resolved.

Definition 2 *A* commitment store *CS_i associated to an agent a_i is a pair:*

$$CS_i = \langle A_i, O_i \rangle$$

where

– *$A_i \subseteq \mathcal{M}$ stands for the set of moves of agent a_i whose speech acts are in $\{\mathcal{S}_1, \mathcal{S}_3, \mathcal{S}_4\}$.*
– *$O_i \subseteq \mathcal{M}$ stands for the set of moves received by agent a_i and whose speech acts are in $\mathcal{S}_2$.*

The disctinction between the two kinds of commitments is necessary because statements expressed by the agent itself commit its honesty whereas statements expressed by other agents commit its cooperativity. They are the different by nature and, in our opinion, should not be mixed.
From the above definition, it is also clear that assertive, commissive and declarative speech acts are related to the first kind of commitments (honesty), whereas directive speech acts are related to the second kind of commitments.

3.2 Penalties

Each time an agent makes a speech act, it makes a commitment. So, it is natural to associate to each commitment a penalty which sanctions the agent if it does not respect this commitment. For the sake of simplicity, we consider that the cost associated to a move depends only of its category.

Definition 3 (Penalty) *Let α_1, α_2, α_3 and α_4 be numbers in $[0, +\infty]$ associated respectively to the categories $\mathcal{S}_1$, $\mathcal{S}_2$, $\mathcal{S}_3$ and $\mathcal{S}_4$.*
A penalty *associated with a category $\mathcal{S}_i$ is the number α_i.*

We distinguish between two kinds of bad behaviors: i) to not be honest, and ii) to not cooperate. Hence, to each commitment store a pair of costs is defined. Formally, the cost associated to the commitment store $CS_i = \langle A_i, O_i \rangle$ of agent a_i at time t, is a pair $\langle c(A_i), c(O_i) \rangle$, where $c(A_i)$ is the cost associated to the statements of agent a_i which are violated, and $c(O_i)$ is the cost associated to the requests made by other agents and to which agent a_i has not answered yet. These costs are computed by summing penalties of moves. Deciding which moves are to take into account depends on their categories.

With such a semantics for an agent communication language, a "rational" protocol should be defined in such a way that each agent aims to minimize the costs associated to its commitment store. This notion of protocol is beyond the scope of this paper.

4 The semantics of the categories of speech acts

In this section we will examine how the penalty associated to a move of each category can be computed, i.e., we will describe when a given commitment is violated by the agent.

4.1 Categories: Assertives ($\mathcal{S}_1$) and Declaratives ($\mathcal{S}_4$)

The assertives and the declaratives behave exactly in the same way, but apply on two distinct languages. Indeed, for the assertives, the content of the moves are formulas of $\mathcal{L}$ built only on the set of non-action variable ($\mathcal{NAV}$), whereas the ones of declaratives are built on action variables ($\mathcal{AV}$).
During a dialogue, an agent may claim that some statement is true. This agent is then, not allowed to contradict itself during all the dialogue otherwise it should pay a penalty.

Storage All the assertive moves are stocked in the commitment store of the agent, exactly in the part A_i. When an assertive move comes as a response to a question or a challenge, for instance, this question or challenge is removed from the set O_i. However, if it is rather a response to a kind of commissive speech act like a promise, then the promise is removed from A_i.
$\forall a \in \mathcal{S}_1$, there are two cases, either the associated assertion is a formula or it is an argument:

- $\forall x \in \mathcal{L}$, if agent a_i commits $a : x$ then
  - $a : x$ is *added* to A_i, and
  - $\forall (a' : y) \in O_i$, such that $a' \in \mathcal{S}_2$ and $y \in \mathcal{L}$ and $x \vdash y$, $\quad a' : y$ is *removed* from O_i, and
  - $\forall (a' : y) \in A_i$, such that $a' \in \mathcal{S}_3$ and $y \in \mathcal{L}$ and $x \vdash y$, $\quad a' : y$ is *removed* from A_i.
- $\forall x = \langle \Gamma, \varphi \rangle \in Arg(\mathcal{L})$, if agent a_i commits $a : x$ then
  - $a : x$ is *added* to A_i, and
  - $\forall (a' : y) \in O_i$, such that $a' \in \mathcal{S}_2$ and $y \in \mathcal{L}$ and $\varphi \vdash y$, $\quad a' : y$ is *removed* from O_i, and
  - $\forall (a' : y) \in A_i$, such that $a' \in \mathcal{S}_3$ and $y \in \mathcal{L}$ and $\varphi \vdash y$, $\quad a' : y$ is *removed* from A_i

Associated Cost: Costs associated to assertives concern non-honesty. An *assertive* move $a : x$, where $a \in \mathcal{S}_1$ and $x \in \mathcal{L} \cup Arg(\mathcal{L})$, is a *violated commitment* if the part A_i of agent a_i's commitments store allows to deduce that the formula x is false or to deduce that one formula of the support of the argument x is false. Hence, the cost associated to an assertive is given by:
$\forall (a : x) \in A_i$, if $a \in \mathcal{S}_1$, there are two cases, either the assertion x is a formula or it is an argument:

- $\forall x \in \mathcal{L}$, $c(a : x) = \begin{cases} \alpha_1 \text{ if } A_i \vdash \neg x, \\ 0 \quad \text{otherwise} \end{cases}$
- $\forall x = \langle \Gamma, \varphi \rangle \in Arg(\mathcal{L})$, $c(a : x) = \begin{cases} \alpha_1 \text{ if } \exists \varphi \in \Gamma, A_i \vdash \neg \varphi, \\ 0 \quad \text{otherwise} \end{cases}$

4.2 Category $\mathcal{S}_2$: Directives

The agent attempts to get another agent to do something. For instance, Request, Question, Challenge are Directives.

Storage Directives of Agent a_i towards Agent a_j are stored in the part O_j of the commitment store of agent a_j, if there is not already an answer to it in A_j (which can be either assertive or declarative):
$\forall a \in \mathcal{S}_2$, the directive $a : x$ of Agent a_i towards Agent a_j is *added* to O_j if $\nexists (a' : y) \in A_j$, such that $a' \in \mathcal{S}_1 \cup \mathcal{S}_4$ and:

- $y \in \mathcal{L}$ and $y \vdash x$
- or $y = \langle \Gamma, \varphi \rangle \in Arg(\mathcal{L})$ and $\varphi \vdash y$.

Associated Cost Costs associated to directives are only penalties of non-cooperativity. It means that a directive is violated if it is in the part O_i of Agent a_i (agent a_i has not already answered to it).
$\forall (a : x) \in O_i$ if $a \in \mathcal{S}_2, \forall x \in \mathcal{L}$, $c(a : x) = \alpha_2$

4.3 Category $\mathcal{S}_3$: Commissives

The agent commits to some future course of action. For instance, Promise is in the Commissives. The most general form of promise is "if φ is true then I swear that ψ will be true", "if you do that, I will kill you". It can be encoded by the move $Promise : \varphi \rightarrow \psi$.

Storage Commissives of Agent a_i are stored in the part A_i of its commitment store if the associated formula is not already true in A_i (either because there is an assertive or a declarative move which makes it true):
$\forall a \in \mathcal{S}_3$, the commissive $a : x$ of Agent a_i is *added* to A_i if $\nexists (a' : y) \in A_i$, such that $a' \in \mathcal{S}_1 \cup \mathcal{S}_4$ and:

- $y \in \mathcal{L}$ and $y \vdash x$
- or $y = \langle \Gamma, \varphi \rangle \in Arg(\mathcal{L})$ and $\varphi \vdash y$.

Associated Cost Costs associated to commissives are only costs of non-honesty. As long as a commissive is in the commitment store, it costs a penalty since it means that this commissive has not been made.
$\forall (a:x) \in A_i$ If $a \in \mathcal{S}_3, \forall x \in \mathcal{L}$, $c(a:x) = \alpha_3$

4.4 Summary

To sum up, the penalties associated to a commitment store CS_i of an agent a_i at turn t are:

cooperation penalty:

$$p_c(CS_i) = \sum_{(a:x) \in O_i} \alpha_2$$

honesty penalty:

$$p_h(CS_i) = \sum_{(a:x) \in A_i \quad s.t. \quad a \in \mathcal{S}_1 (resp. \mathcal{S}_4) \quad and \quad x = \langle \Gamma, \varphi \rangle \in Arg(\mathcal{L}) \text{ and } \exists \psi \in \Gamma \text{ s.t. } A_i \vdash \neg \psi} \alpha_1 (resp. \alpha_4)$$

$$+ \sum_{(a:x) \in A_i \quad s.t. \quad a \in \mathcal{S}_1 (resp. \mathcal{S}_4) \quad and \quad x \in \mathcal{L} \text{ and } A_i \vdash \neg x} \alpha_1 (resp. \alpha_4)$$

$$+ \sum_{(a:x) \in A_i \quad s.t. \quad a \in \mathcal{S}_3} \alpha_3$$

Or equivalently, (since we suppose that $c(a:x)$ only depends on the category of a):

$$p_c(CS_i) = \sum_{(a:x) \in O_i, a \in \mathcal{S}_2} c(a:x)$$

and

$$p_h(CS_i) = \sum_{(a:x) \in A_i \begin{cases} s.t. \quad a \in \mathcal{S}_1 \cup \mathcal{S}_4 \\ or \quad and \begin{cases} or\, x = \langle \Gamma, \varphi \rangle \in Arg(\mathcal{L}) \text{ and } \exists \psi \in \Gamma \text{ s.t. } A_i \vdash \neg \psi \\ x \in \mathcal{L} \text{ and } A_i \vdash \neg x \end{cases} \\ s.t. \quad a \in \mathcal{S}_3 \end{cases}} c(a:x)$$

5 Related work

The first standard agent communication language is KQML [7]. It has been developed within the Knowledge Sharing Effort, a vast research program funded by DARPA (the US Defense Advanced Research Projects Agency). More recently, the Foundation for Intelligent Physical Agents (FIPA) has proposed a new standard, named ACL [8]. Both KQML and ACL have been given a mentalistics semantics. The semantics is based on a notion of speech act close to the concept of illocutionary act as developed in speech act theory [4, 12].
Such semantics assumes, more or less explicitly, some underlying hypothesis in particular, that the agents are sincere and cooperative. While this may be well fitted for some special cases of interactions, it is obvious that some dialogue types listed by Walton and Krabbe in [15] are not cooperative. For example, assuming sincerity and cooperativity in negotiation may lead to poor negotiators. Another more important limitation of this approach is the fact that it is not verifiable since it is based on the mental states of the agents. Our semantics does not refer at all to the mental states of the agents. Moreover, it does not treat particular speech acts as it is the case with this mentalistic approach.

In the second approach, called *social* and developed in [5, 13, 14], primacy is given to the interactions among the agents. The semantics is based on social *commitments* brought about by performing a speech act. For example, by affirming a data, the agent commits on the truth of that data. After a promise, the agent is committed carrying it out. There are several weak points of this approach and we summarize them in the three following points:
1. The definition of commitments complicates the agent architecture in the sense that it needs an ad hoc apparatus. The commitments are introduced especially for modeling communication. Thus agents should reason not only on their beliefs, etc but also on the commitments. In our approach, we didn't introduce any new language to treat commitments. We call a commitment any information stocked in a commitment store. Handling these commitments (to add a new commitment, to retract a commitment, to achieve a commitment, to violate a commitment) is done directly on the commitment store.
2. The level at which communication is treated is very abstract, and there is a considerable gap to fill in order to bring the model down to the level of implementation. However, the semantics presented in this paper can be implemented easily.
3. The concept of commitment is ambiguous and its semantics is not clear. According to the speech act, the semantics of the commitment differs. For example:

Inform: by affirming a data, the agent commits on the truth of that data. The meaning of the commitment here is not clear. It may be that the agent can justify the data or can defend it against any attack, or the agent is sincere.

Request: According to Colombetti [5] after a request, the receiver precommits to carry it out. This idea is in our opinion drawn from the notion of protocol. The protocol specifies for each act the set of allowed replies. So after a request, generally the receiver can accept it, reject it or propose another alternative.

The last approach developed by Pitt and Mandani in [10] is based on the notion of protocol. A protocol defines what sequences of moves are conventionally expected in the dialogue. The meaning of a speech act then equates with the set of the possible following answers.
However, protocols are often technically finite state machines. This turns out to be too rigid in several circumstances. Current research aims at defining flexible protocols, which rely more on the *state of the dialogue*, and less on dialogue history. This state of dialogue is captured through the notion of commitment.

6 Conclusion

This paper has introduced a *general* semantics of any agent communication language. The semantics is general in the sense that it can be used with any set of speech acts since it is defined on categories of speech acts rather than the speech acts themselves. Indeed, for the first time, we have a semantics of an ACL which is defined independently from the syntax of that ACL. This makes the semantics *flexible* and general.
The new semantics is social in nature and is based on the notion of *commitments*. However, our modeling of this notion is very different from the existing approaches. Indeed, the commitments are modeled as in MacKenzie's system DC [9]. A commitment is any information stocked in the *commitment store* of the agent. Two kinds of commitments have been distinguished: the commitments to satisfy or to respect what the agent has already said, promised, etc. and the commitments to answer to other agents.
Unlike in the existing approaches, the semantics of a commitment is *unique* and is given via the notion of penalty that should be paid if an agent does not respect that commitment.
Our semantics goes somewhat beyond the existing approaches in giving an operational, verifiable and flexible semantics. An extension of this work would be to simplify the protocols by extending our semantics by rules which were generally defined in the protocol itself. For instance, in the semantics, we can sanction agents which repeat, the same move several times during a dialogue. Another extension would be to refine the penalties by introducing granularity in the categories of speech acts. The idea is, maybe, to associate different penalties for the speech acts in order to capture the idea that some commitments are harder to violate than others. The same idea applies also to the content of the moves.

References

1. L. Amgoud, N. Maudet, and S. Parsons. Modelling dialogues using argumentation. In E. Durfee, editor, *Proceedings of the 4th International Conference on Multi-Agent Systems (ICMAS00)*, pages 31–38, Boston, USA, July 2000. IEEE Press.

2. L. Amgoud, N. Maudet, and S. Parsons. An argumentation-based semantics for agent communication languages. In F. Van Harmelen, editor, *Proceedings of the European Conference on Artificial Intelligence (ECAI-2002)*, pages 38–42, Lyon, France, July 2002. IOS Press.
3. L. Amgoud, S. Parsons, and N. Maudet. Arguments, dialogue, and negotiation. In W. Horn, editor, *Proceedings of the European Conference on Artificial Intelligence (ECAI-2000)*, pages 338–342, Berlin, Germany, August 2000. IOS Press.
4. J. L. Austin. How to do things with words. In *Clarendon Press, Oxford Uk*, 1962.
5. M. Colombetti. A commitment-based approach to agent speech acts and conversations. In *Proceedings of the Workshop on Agent Languages and Conversation Policies. I4th International Conference on Autonomous Agents*, pages 21–29, 2000.
6. P. M. Dung. On the acceptability of arguments and its fundamental role in non-monotonic reasoning, logic programming and n-person games. *Artificial Intelligence*, 77:321–357, 1995.
7. T. Finin, Y. Labrou, and J. Mayfield. Kqml as an agent communication language. In *J. Bradshaw, ed. Softaware agents, MIT Press, Cambridge.*, 1995.
8. FIPA. Agent communication language. In *FIPA 97 Specification, Foundation for Intelligent Physical Agents.*, 1997.
9. J. MacKenzie. Question-begging in non-cumulative systems. *Journal of philosophical logic*, 8:117–133, 1979.
10. J. Pitt and A. Mamdani. A protocol based semantics for an agent communication language. In *Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI'99)*, pages 486–491, 1999.
11. J. R. Searle. Speech acts. In *Cambridge University Press*, 1969.
12. J. R. Searle and D. Vanderveken. Foundations of illocutionary logic. In *Cambridge University Press*, 1985.
13. M. P. Singh. Agent communication languages: Rethinking the principles. In *IEEE Computer*, pages 40–47, 1998.
14. M. P. Singh. A social semantics for agent communication languages. In *IJCAI'99 Workshop on Agent Communication Languages*, pages 75–88, 1999.
15. D. N. Walton and E. C. W. Krabbe. *Commitment in Dialogue: Basic Concepts of Interpersonal Reasoning.* State University of New York Press, Albany, NY, 1995.