

Denotational Semantics for Agent Communication Languages

Frank Guerin and Jeremy Pitt

Intelligent and Interactive Systems
Department of Electrical & Electronic Engineering,
Imperial College of Science, Technology & Medicine,
Exhibition Road, London, SW7 2BZ.
+44 20 7594 6331

{f.guerin,j.pitt}@ic.ac.uk

ABSTRACT

The dilemma encountered in the design of an agent communication language (ACL) for an open society is that it should be based on externally observable phenomena yet it should capture something of the intuitions behind the high level abstractions typically found in internal mental states. Our solution treats an ACL message as a declarative statement that is given a procedural interpretation by a denotational semantics. This defines a speech act as a function between states. These states are *social states* which store public information including expressed mental attitudes and control variables. Expressed mental attitudes are externally observable and capture the conventional public meaning of communication. The variables control the flow of conversation in a protocol. We conclude firstly that since the denotational semantics is based on externally observable phenomena, it is possible to verify compliance and prove properties of protocols. Secondly, since the semantics is more expressive than behavioural specifications, it lays the foundation for high-level communication between intelligent agents.

1. INTRODUCTION

We are concerned with specifying communication languages for agents in an open society. In such systems the internal states of agents are not accessible so the semantic definition should be based on externally observable phenomena. This poses problems for semantics based on declarative statements about mental states. In previous work [1], we analysed an auction protocol based on a procedural approach, using expected replies for external semantics and further specifications for decision-making based on state variables and mental attitudes. This work revealed firstly the existence of a richer state description commonly known to all agents participating in the auction, which was changed by their speech acts. Secondly, the speech acts had consequences beyond just the intended replies. Agents were publicly committed to certain other acts (auctioneer commits to sell, bidder commits to buy). However, these state variables and commitments were specified in the logical specification of the protocol and not accommodated in the general semantic framework. This paper shifts all the public knowledge into an external semantics for agent communication. It therefore advances the semantic framework while being complementary to and consistent with our

work on intentional specifications of internals. We treat an ACL message as a declarative statement that is given a procedural interpretation by means of denotational semantic evaluations. We introduce the notion of expressed mental attitudes which are externally observable and capture the conventional public meaning of communication. We define a *social state* which holds all public information and is explicitly modified by the semantic definitions of speech acts. The social state stores propositions and control variables and defines commitments for future acts.

We begin with an analysis of the chief problems with existing ACLs and list some desiderata. We then describe our semantic framework. This is followed by the formal semantics of the framework which includes an auction protocol specification. Finally we evaluate the framework and outline areas for further research. Since the formal semantics are based on externally observable phenomena, it is possible to verify compliance and prove properties such as the possible outcomes of protocols. Furthermore, since the semantics is more expressive than behavioural specifications, it lays the foundation for intelligent communication between agents in the long term.

2. PROBLEMS AND REQUIREMENTS

Agent communication languages (ACLs) are typically designed by either a *procedural* or *declarative* approach [4]. In the *procedural* approach agents exchange procedural directives, a message defines the appropriate response from the recipient. Behaviour based approaches extend this to larger sections of conversation, identifying certain patterns of communication (auction, contract-net, negotiation) that arise in many situations and designing protocols for them. Speech acts are given a behavioural meaning in the context of a conversation in that only certain responses are appropriate at any stage of a conversation. Procedural communication languages are more oriented toward computer implementation, but there are some difficulties:

1. Behavioural meanings are low level and are more appropriate for computational entities which lack the rationality to plan their own behaviour. Multiagent systems are supposed to provide higher level abstractions than traditional distributed programming [8].
2. If the protocol is the unit of communication and speech acts are defined only in terms of possible replies then communication becomes an ordered exchange of meaningless tokens and the language is not sufficiently expressive.
3. By specifying exactly how a speech act should be used in a conversation there is a lack of flexibility to use it in different situations and one needs a library of predefined protocols to handle every possible conversation.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

AGENTS'01, May 28-June 1, 2001, Montréal, Québec, Canada.

Copyright 2001 ACM 1-58113-326-X/01/0005...\$5.00.

For example, if the meaning of *request* is defined behaviourally as an intention to reply with *agree* or *refuse*, then a helpful intelligent agent could not suggest an alternative source in response to a request for a service it is incapable of providing. To enable such cooperation, the semantics of request should capture the intuitive meaning i.e. the requestor expresses a desire. The behaviour based approach is adequate for simple reactive agents, but a more meaningful semantics is vital for rational agents to be able to plan their communications intelligently.

The *declarative* approach is akin to natural language. Communicative acts are given a meaning compatible with human intuitions and agents are expected to have a sufficiently advanced level of rationality to interpret received acts and plan new acts. The semantics of acts is typically specified in terms of mental attitudes of participants [2],[5] (also called intentional semantics). Such languages are expressive but have the following drawbacks:

1. It is difficult to verify compliance with a semantics which is based on internal states if the internal states are inaccessible.
2. An agent's autonomy is limited. The example most frequently cited [8] is the inability of FIPA [2] agents to operate in a setting where sincerity cannot be taken for granted.
3. Meaning is specific to a certain mental state thus the speech acts are not flexible for use in a different context. Ideally the meaning of an act should depend on the context.
4. It is not clear how the agent's reasoning should be designed so that desired behaviour results, especially when protocol specification is attempted. Many such languages specify protocols without considering the speech act semantics [2].

A significant advancement to the declarative approach is presented in Singh [7], where semantics are declarative, describing social commitments which are externally observable thus making verification possible.

To summarise, the semantic definition of an ACL should have the following features:

1. Formal definition. If it is specified informally, with a natural language description, it is possible that different vendors will take different interpretations of the specification and produce agents which cannot interoperate.
2. Verifying compliance with the semantics should be possible.
3. State what a speech act means rather than how it should be used. This makes the language more expressive and makes speech acts flexible so they can be used in different contexts.
4. Intuitive or human-like semantics. To facilitate high level communication for intelligent agents.
5. Context dependence. The context in which a speech act is uttered should be a factor in determining its meaning.
6. Based on conventional meaning (public perspective) rather than the inferences of an individual (private perspective).
7. Design Autonomy. Compliant agents should not be forced to adopt a certain architecture, e.g. they should not be forced to have an explicit representation of beliefs, desires or intentions.
8. Consideration of the social aspects of communication and the perspective of the society of agents. It should be possible to specify (observable) commitments to actions.

An ACL should also have a formal protocol description language which can cope with multi-party protocols, nested protocols, parallel conversations and different agent roles within a conversation. There should be a precise relationship between the semantics of protocols and the semantics of the speech acts of which it is composed. We address these issues in the next section.

3. SEMANTIC FRAMEWORK

An ACL message is treated as a declarative statement that is given a procedural interpretation by means of denotational semantic evaluations. We define a speech act as a function from context onto context [3]. Hamblin (cited in Gazdar [3]) suggests that the context can take the form of a list of propositions representing commitments. We extend this to include propositions representing expressed mental attitudes, control variables for the conversation and the history of speech acts. We distinguish between an agent's publicly expressed mental attitudes and its personal internal mental attitudes. These can be different (if agents are not sincere). For example, an agent may express a desire to have an action performed (when it does have that desire in its internal state) in order to test the willingness of another agent to comply. This means that an agent does not need to hold a mental attitude as a precondition to expressing it. Thus meanings are specified from a public perspective, rather than the private perspective of a single individual whose personal inferences are subjective. This public meaning together with social relations captures the conventional meaning of acts within the society.

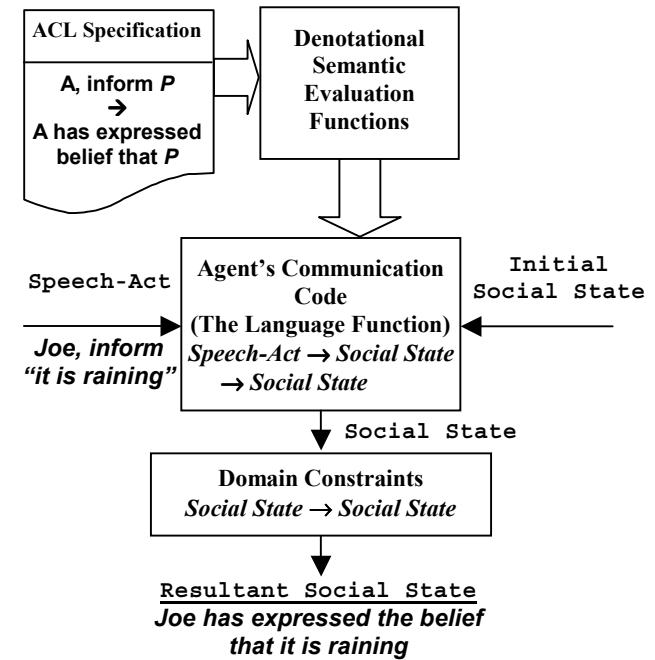


Figure 1 Communication Model

3.1 Public Information

We represent the context as a *social state* which is used to hold information that is externally visible and known to all participants in a social interaction. Within the social state is the state of *persistent social relations* (dealing with long term commitments for example) and the *conversation states* (dealing with all public information relating to the current conversations). Each conversation state can be further broken down into a set of

⁽¹⁾variables and ⁽²⁾propositions describing the conversation, a set of ⁽³⁾contingent social relations and the ⁽⁴⁾history of speech acts in the conversation. We distinguish between contingent and persistent social relations as follows: contingent relations are solely concerned with the permissions and obligations (to perform subsequent speech acts in the current conversation) that arise as a result of the current conversation state. Each agent keeps a copy of the social state and updates it as speech acts are sent and received (both sender and receiver change state). The conversation variables and propositions control the conversation, although they are labelled with beliefs, desires and intentions, they need not have any direct relationship with the internal states of the participants. The use of intentional labels makes the state description more intuitive for the agent designer. Agents play certain roles in a conversation and these roles may define permissions and obligations at any point in the protocol. Conversation variables keep track of which agents occupy which roles at any time during the conversation.

As shown in Figure 1, The ACL specification is treated like a computer program which is turned into a function (compiled) by the denotational semantics. The function maps a speech act to a change in the state of the society. The ACL specification has three parts: ⁽¹⁾the *Converse Function* gives permissions and obligations for subsequent speech acts based on the current conversation state; ⁽²⁾the *Protocol Semantics* gives the meanings of speech acts in context of the current protocol and ⁽³⁾the *Speech-Act Semantics* gives the protocol independent elements of meaning. For example, announcing may have a protocol independent meaning which makes some fact public as well as an additional meaning which depends on whether it is used in an auction or a negotiation. Protocol specific meanings may override the speech act semantics. These three functions are incorporated in the *language function*. The language function matches incoming speech acts with the correct conversation state (based on the conversation identifier parameter) and updates the history in that state.

We do not define the appropriate replies to an act along with the semantics of the act. We consider the separation between the converse function and the semantic function to be important since the determination of the possible replies to an act may not depend on that act alone, but also on the history of previous acts.

Domain constraints further modify the social state. For example, we could specify for some domain that if an agent *A* expresses a desire for an agent *B* to perform an action, then the agent *B* is obliged to do it. Domain constraints are external to the ACL specification, this allows for the same ACL to be used in different domains with different constraints. These rules could appear in the converse function or in the specification of protocols, but this would lead to a specification which is locked in context.

3.2 Private Information

The ACL specification describes how the social state is to change as a result of speech acts and thereby defines the permissions and obligations for participants to perform subsequent acts. It remains to describe how the agents update their internal information states and how they select the next act to perform. Agents in a heterogeneous system may be implemented in diverse ways and may not even have an explicit representation of desires and intentions, therefore the description here is purely informative. The internal code of the agent can be implemented as two functions: ⁽¹⁾The *add* function has as inputs the internal state,

String Lists

Domain $l \in \text{String-List} = \text{String} \rightarrow \text{Tr}$

Operations

newlist : variables

newlist = $\lambda s. \text{false}$

checklist : $\text{String} \rightarrow \text{String-List} \rightarrow \text{Tr}$

checklist = $\lambda s. \lambda l. l(s)$

updatelist : $\text{String} \rightarrow \text{String-List} \rightarrow \text{String-List}$

updatelist = $\lambda s. \lambda l. [s \mapsto \text{true}]l$

Variables

Domain $v \in \text{variables} = \text{Id} \rightarrow (\text{Nat} + \text{String})$

Operations

newvars : variables

newvars = $\lambda i. (\text{one}, \text{"null"})$

access : $\text{Id} \rightarrow \text{variables} \rightarrow (\text{Nat} + \text{String})$

access = $\lambda i. \lambda v. v(i)$

update : $\text{Id} \rightarrow (\text{Nat} + \text{String}) \rightarrow \text{variables} \rightarrow \text{variables}$

update = $\lambda i. \lambda n. \lambda v. [i \mapsto n]v$

History

Domain $v \in \text{History} = \text{seq} \rightarrow \text{Speech-Act}$

Operations

newHistory : History

newHistory = $\lambda i. (\text{null}, \text{null}, \text{null}, \text{null}, \text{null})$

accessHistory : $\text{seq} \rightarrow \text{History} \rightarrow \text{Speech-Act}$

accessHistory = $\lambda i. \lambda v. v(i)$

updateHistory : $\text{seq} \rightarrow \text{Speech-Act} \rightarrow \text{History}$

updateHistory = $\lambda i. \lambda n. \lambda v. [i \mapsto n]v$

Conversation State

Domain $s \in \text{Conversation-State}$

= $\text{variables} \times \text{String-List} \times \text{String-List} \times \text{History}$

Operations

newconversation : Conversation-State

newconversation = $(\text{newvars}, \text{newlist}, \text{newlist}, \text{newHistory})$

Speech Act

Domain $a \in \text{Speech-Act} =$

$\text{Name} \times \text{Name} \times \text{perf} \times \text{content} \times \text{cid} \times \text{seq}$

Cases of Speech Acts

Domain $c \in \text{Speechactcase} =$

$\text{Speech-Act} \rightarrow \text{Conversation-State} \rightarrow \text{Conversation-State}$

Operations

nostatechange : Speechactcase

nostatechange = $\lambda a. \lambda s. s$

errorstate : Speechactcase

errorstate = $\lambda a. \lambda s. (\text{error}, \text{error})$

Figure 2 Semantic Algebras

social state and speech act and outputs a new internal state. This function updates the agent's mental state when a speech act arrives. ⁽²⁾The *select* function has as inputs the social state and the internal state and outputs a new internal state. The agent reasons about what acts to perform next and adds them to the its intentions, obviously the permissions and obligations defined in the social state should be a major factor in determining the output if the agent wishes to be ACL compliant. The *add* and *select* functions are implemented after the public information is updated.

3.3 Development Method

The expressive power of natural language is achieved through words or phrases that can take different meanings depending on the context in which they are uttered. This gives flexibility to each natural language word or phrase, and an economy to the number of words needed. Meaning is built from a group of words or phrases in a particular order and in a particular pragmatic context, rather than from a single all-meaningful performative. This concept of building complex meaning from simple building blocks is what we would like to borrow from natural language. Our ACL building blocks are ⁽¹⁾speech act specification, ⁽²⁾protocol specification (via protocol semantics and converse functions) and ⁽³⁾domain constraints. The speech acts themselves should carry as much useful meaning as possible, while still being flexible, so as they lend themselves to use in many different scenarios. They should not be specific to one protocol, one domain or one ACL. This design philosophy follows through in all specifications, for example, domain specific aspects should be kept out of protocol specifications. The specifications for generic speech acts and protocols can be published. An ACL can then be developed by downloading speech acts and protocols, and if necessary, designing new ones. ACL specifications can be published for designers and also so that foreign agents can compile them into code.

4. DENOTATIONAL SEMANTICS

We now present a method for defining the language function precisely. Essentially we treat propositions in the social state as strings and define a language function mapping a speech act onto a state change. We follow the notation of Schmidt [6] where a function f is written with the lambda calculus abstraction $\lambda a.e$ and the function argument appears after the function. To evaluate the function, the argument replaces occurrences of a in e . For example if $f(y) = y^2$ then we write $f = \lambda a.a^2$ and $f(y) = \lambda a.a^2 y$. If f is a function, then $[x \mapsto y]f$ denotes the function that maps x to y , but behaves exactly like f for any other argument. \downarrow_i denotes the operation such that $(a_1, a_2, \dots, a_n) \downarrow_i = a_i$. Let $a \rightarrow b \square c$ take the value b if a is true, or c if a is false. For $x \in R$ and $y \in S$, we tag the members of each set so they can be distinguished: $\text{in}R(x) = (\text{zero}, x)$ and $\text{in}S(y) = (\text{one}, y)$. To remove the tags for any $m \in R+S$, the expression:

cases m of $\text{is}R(x) \rightarrow f(x) \square \text{is}S(y) \rightarrow g(y)$ end

evaluates to $f(x)$ when $m = (\text{zero}, x)$ and evaluates to $g(y)$ when $m = (\text{one}, y)$. ($a \text{ strequals } b$) returns true if a and b are identical strings and cons concatenates strings. For brevity, we omit the abstract syntax definitions and the standard semantic algebras for truth values, strings, etc. Many of the trivial evaluation functions dealing with low-level string manipulations are omitted. We first present a simplified version of the semantics and illustrate how it works with a simple ACL containing a single speech act which is not part of any protocol. We then move on to the complete semantics and animate an ACL containing an auction protocol.

4.1 Simplified Definition (Speech Acts Only)

We wish to define the language function as a mapping from speech act to state to state. Therefore we must first define the domains for speech acts and states, see Figure 2. The domains ⁽¹⁾*seq*, ⁽²⁾*Name*, ⁽³⁾*perf*, ⁽⁴⁾*content* and ⁽⁵⁾*cid* are domains of ⁽¹⁾sequence numbers for conversations, ⁽²⁾agent names, ⁽³⁾performative names, ⁽⁴⁾message contents and ⁽⁵⁾conversation

identifiers respectively. A speech act has six parts: sender, receiver, performative, content, conversation identifier and sequence number within this conversation. The conversation state has four parts: ⁽¹⁾variables (string or numerical values) which may define roles or other relevant information for the execution of a conversation, ⁽²⁾expressions describing the mental attitudes expressed by speakers (treated as strings), ⁽³⁾contingent social relations (described in section 3.1) and ⁽⁴⁾the conversation history. The *speechactcase* domain is necessary to handle more than one speech act as will be seen in section 4.1.3. In the simplified version of the evaluation functions (Figure 3) we only consider the speech act's effect on the conversation state with no protocol.

L: Language \rightarrow
 $\text{Speech-Act} \rightarrow \text{Conversation-State} \rightarrow \text{Conversation-State}$
 $\mathbf{L}[\text{[speech-act-semantics } S]] = S[[S]] \text{ nostatechange}$

S: Speech-Act-Semantics $\rightarrow \text{Speechactcase} \rightarrow \text{Speechactcase}$
 $S[[S_1; S_2]] = \lambda c. S[[S_2]] (S[[S_1]]c)$
 $S[[[F \mid M]]] = \lambda c. \lambda a. \lambda q. a \downarrow 3 \text{ strequals } F[[[F \mid]]]$
 $\rightarrow M[[M]] a q \square c a q$

M: Semantics \rightarrow
 $\text{Speech-Act} \rightarrow \text{Conversation-State} \rightarrow \text{Conversation-State}$
 $M[[[M_1; M_2]]] = \lambda a. \lambda q. M[[M_2]] a (M[[M_1]] a q)$
 $M[[[R]]] = \lambda a. \lambda q. (q \downarrow 1, \text{updatelist } (R[[[R]]]a) q \downarrow 2, q \downarrow 3, q \downarrow 4)$

R: Proposition $\rightarrow \text{Speech-Act} \rightarrow \text{String}$
 $R[[[X R]]] = \lambda a. (X[[[X]]]a) \text{ cons } (R[[[R]]]a)$
 $R[[[C]]] = \lambda a. a \downarrow 4$

X: Proposition $\rightarrow \text{Speech-Act} \rightarrow \text{String}$
 $X[[[Y Z]]] = \lambda a. Y[[[Y]]] \text{ cons } (Z[[[Z]]]a)$

Y: Modality $\rightarrow \text{String}$
 $Y[[[\text{believe}]]] = B \quad Y[[[\text{desire}]]] = D$
 $Y[[[\text{intend}]]] = I \quad Y[[[\text{know}]]] = K$

Z: Actor $\rightarrow \text{Speech-Act} \rightarrow \text{String}$
 $Z[[[R]]] = \lambda a. a \downarrow 2 \quad Z[[[S]]] = \lambda a. a \downarrow 1$

$F[[[F \mid]]]$ simply maps the speech act name to its denotable value

Figure 3 Simplified version of the valuation functions.

4.1.1 Evaluation of a simple ACL

We now evaluate l_i , a sample fragment of an ACL which has only one speech act (*query*) and no protocols, see Figure 4. This defines the semantics of a query as the sender's expression of a desire to know if the receiver believes the content. We want to be able to turn this ACL into a function from a speech act to a change in conversation state. If the performative of the incoming speech act matches the one in our ACL (*query*), then our function should add the sender's expressed desire to the state.

speech-act-semantics

[query]

desire S know S believe R C

Figure 4 Specification of a simple ACL.

Applying the **L** valuation function to l_i we have:

$\mathbf{L}[[l_i]] = S[[s_i]] \text{ nostatechange}$

Where $s_i = [\text{query}]$ desire S know S believe R C

$$\mathbf{S}[[s_i]] = \lambda c. \lambda a. \lambda q. a \downarrow 3 \text{ strequals } \mathbf{F}[[f_i]] \rightarrow \mathbf{M}[[m_i]] \text{ a } q \sqcap c \text{ a } q$$

Where $m_i = \text{desire S know S believe R C}$,
and $f_i = \text{query}$

$$\mathbf{F}[[\text{query}]] = \text{query}$$

$$\mathbf{M}[[m_i]] = \lambda a. \lambda q. (q \downarrow 1, \text{updatelist } (\mathbf{R}[[r_i]]a) q \downarrow 2, q \downarrow 3, q \downarrow 4)$$

Where $r_i = m_i = \text{desire S know S believe R C}$

All that remains is to simplify $\mathbf{R}[[r_i]]$:

$$\begin{aligned} \mathbf{R}[[r_i]] &= \mathbf{R}[[\text{desire S know S believe R C}]] \\ &= \lambda a. (\mathbf{X}[[\text{desire S}]]a) \\ &\quad \text{cons } (\mathbf{R}[[\text{know S believe R C}]]a) \\ &= \lambda a. (\mathbf{X}[[\text{desire S}]]a) \text{ cons } (\lambda a'. (\mathbf{X}[[\text{know S}]]a') \\ &\quad \text{cons } (\mathbf{R}[[\text{believe R C}]]a') a) \\ &= \lambda a. (\mathbf{X}[[\text{desire S}]]a) \text{ cons } (\lambda a'. (\mathbf{X}[[\text{know S}]]a') \\ &\quad \text{cons } (\lambda a''. (\mathbf{X}[[\text{believe R}]]a'') \text{ cons } (\mathbf{R}[[C]]a'') \\ &\quad a') a) \\ \mathbf{X}[[\text{desire S}]] &= \lambda a. \mathbf{Y}[[\text{desire}]] \text{ cons } (\mathbf{Z}[[S]]a) \\ &= \lambda a. D \text{ cons } ((\lambda a'. a' \downarrow 1) a) \\ &= \lambda a. D \text{ cons } a \downarrow 1 \end{aligned}$$

and similarly for the other X valuations. Note that superscripts are used merely to distinguish between separate identifiers represented by the same letter in nested abstractions.

$$\mathbf{R}[[C]]a = \lambda a. a \downarrow 4$$

The full valuation of r_i is:

$$\begin{aligned} \mathbf{R}[[r_i]] &= \lambda a. (\lambda a'. D \text{ cons } a' \downarrow 1 a) \\ &\quad \text{cons } (\lambda a'. (\lambda a''. K \text{ cons } a'' \downarrow 1 a') \\ &\quad \quad \text{cons } (\lambda a'''. (\lambda a'''. B \text{ cons } a''' \downarrow 2 a'') \\ &\quad \quad \quad \text{cons } (\lambda a'''. a''' \downarrow 4 a'') \\ &\quad \quad \quad a') \\ &\quad a) \\ &= \lambda a. D \text{ cons } a \downarrow 1 \text{ cons } K \text{ cons } a \downarrow 1 \text{ cons } B \text{ cons } a \downarrow 2 \text{ cons } a \downarrow 4 \end{aligned}$$

\mathbf{R} is a function from speech act to string, effectively it takes the speech act a and puts the sender, receiver and content in the right places in the string. This completes the denotation of l_i .

4.1.2 Testing sample speech act inputs:

$\mathbf{L}[[l_i]]$ is a function, given a speech act and conversation state, it returns the new conversation state. Looking inside $\mathbf{L}[[l_i]]$ we find $\mathbf{S}[[s_i]]$, which is a function of *Speechactcase* onto *Speechactcase*. We see that if the performative of the speech act input to $\mathbf{L}[[l_i]]$ is anything other than *query*, the returned *Speechactcase* is *nostatechange* and hence the returned state is the same as the original:

$$\begin{aligned} \mathbf{L}[[l_i]] &= \lambda c. \lambda a. \lambda q. (a \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \\ &\quad \mathbf{M}[[m_i]] \text{ a } q \sqcap c \text{ a } q) \text{ nostatechange} \\ &= \lambda a. \lambda q. (a \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \\ &\quad \mathbf{M}[[m_i]] \text{ a } q \sqcap \text{nostatechange } a \text{ } q) \\ &= \lambda a. \lambda q. (a \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \\ &\quad \mathbf{M}[[m_i]] \text{ a } q \sqcap \lambda a'. \lambda s'. s' \text{ a } q) \end{aligned}$$

Let a_i be a speech act with a performative which is not *query*:

$$\begin{aligned} \mathbf{L}[[l_i]] a_i &= \lambda a. \lambda q. (a \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \\ &\quad \mathbf{M}[[m_i]] \text{ a } q \sqcap \lambda a'. \lambda s'. s' \text{ a } q) a_i \\ &= \lambda q. (a_i \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \mathbf{M}[[m_i]] \\ &\quad a_i q \sqcap \lambda a'. \lambda s'. s' a_i q) \\ &= \lambda q. (\lambda a'. \lambda s'. s' a_i q) \\ &= \lambda q. (q) \end{aligned}$$

This is a mapping from conversation state to conversation state which leaves the state unchanged. Now consider a speech act $a_2 = (\text{sender}_2, \text{receiver}_2, \text{query}, \text{content}_2, \text{id}_2)$:

$$\begin{aligned} \mathbf{L}[[l_i]] a_2 &= \lambda a. \lambda q. (a \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \\ &\quad \mathbf{M}[[m_i]] \text{ a } q \sqcap \lambda a'. \lambda s'. s' \text{ a } q) a_2 \\ &= \lambda q. (a_2 \downarrow 3 \text{ strequals } \mathbf{F}[[\text{query}]] \rightarrow \\ &\quad \mathbf{M}[[m_i]] a_2 q \sqcap \lambda a'. \lambda s'. s' a_2 q) \\ &= \lambda q. (\mathbf{M}[[m_i]] a_2 q) \\ &= \lambda q. (\lambda a'. \lambda q'. \\ &\quad (q' \downarrow 1, \text{updatelist } (\mathbf{R}[[r_i]]a') q' \downarrow 2 a_2 q, q \downarrow 3, q \downarrow 4) \\ &\quad = \lambda q. (q \downarrow 1, \text{updatelist } (\mathbf{R}[[r_i]] a_2) q \downarrow 2, q \downarrow 3, q \downarrow 4) \\ \mathbf{R}[[r_i]] a_2 &= D \text{ cons } (a_2 \downarrow 1) \text{ cons } K \text{ cons } (a_2 \downarrow 1) \\ &\quad \text{cons } B \text{ cons } (a_2 \downarrow 2) \text{ cons } (a_2 \downarrow 4) \\ &= D \text{ sender}_2 K \text{ sender}_2 B \text{ receiver}_2 \text{ content}_2 \end{aligned}$$

Let us call this string *string*₂

$$\mathbf{L}[[l_i]] a_2 = \lambda q. (q \downarrow 1, \text{updatelist } \text{string}_2 q \downarrow 2, q \downarrow 3, q \downarrow 4)$$

This is a function from conversation state to state which adds the proposition contained in *string*₂ to the state (as desired).

4.1.3 The evaluation of two or more speech acts

Note how \mathbf{S} handles semantics for two speech acts:

$$\begin{aligned} \mathbf{L}[[\text{speech-act-semantics } S_1; S_2]] &= \mathbf{S}[[S_1; S_2]] \text{ nostatechange} \\ &= \lambda c. \mathbf{S}[[S_2]] (\mathbf{S}[[S_1]]c) \text{ nostatechange} \\ &= \mathbf{S}[[S_2]] (\mathbf{S}[[S_1]] \text{ nostatechange}) \\ &= \mathbf{S}[[S_2]] (\lambda c. \lambda a. \lambda q. a \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] \text{ a } q \\ &\quad \sqcap c \text{ a } q \text{ nostatechange}) \end{aligned}$$

Where f_i and m_i are the performative name and semantics of S_i .

$$\begin{aligned} &= \mathbf{S}[[S_2]] (\lambda a. \lambda q. a \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] \text{ a } q \\ &\quad \sqcap \text{nostatechange } a \text{ } q) \\ &= \lambda c. \lambda a. \lambda q. a \downarrow 3 \text{ strequals } \mathbf{F}[[f_2]] \rightarrow \mathbf{M}[[m_2]] \text{ a } q \sqcap c \text{ a } q \\ &\quad (\lambda a'. \lambda q'. a' \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] a' q' \\ &\quad \sqcap \text{nostatechange } a' q') \\ &= \lambda a. \lambda q. a \downarrow 3 \text{ strequals } \mathbf{F}[[f_2]] \rightarrow \mathbf{M}[[m_2]] \text{ a } q \sqcap \\ &\quad (\lambda a'. \lambda q'. a' \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] a' q' \\ &\quad \sqcap \text{nostatechange } a' q') a \text{ } q \end{aligned}$$

The functionality of this expression is:

$$\text{Speech-Act} \rightarrow \text{Conversation-State} \rightarrow \text{Conversation-State}$$

Consider an incoming speech act a_1 . If the performative name of this incoming act matches f_2 , we get:

$$\begin{aligned} &\lambda a. \lambda q. (a \downarrow 3 \text{ strequals } \mathbf{F}[[f_2]] \rightarrow \mathbf{M}[[m_2]] \text{ a } q \sqcap \dots) a \text{ } a_1 \\ &= \lambda q. \mathbf{M}[[m_2]] a_1 q \end{aligned}$$

i.e. the state change defined by semantics m_2 when a_1 is the speech act. If the performative name does not match f_2 , we get :

$$\begin{aligned} & \lambda q. (a_1 \downarrow 3 = \mathbf{F}[[f_2]] \rightarrow \mathbf{M}[[m_2]] a_1 q \square \\ & \quad (\lambda a'. \lambda q'. a' \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] a' q' \\ & \quad \square \text{nostatechange } a' q') a_1 q) \\ & = \lambda q. ((\lambda a'. \lambda q'. a' \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] a' q' \\ & \quad \square \text{nostatechange } a' q') a_1 q) \\ & = \lambda q. (a_1 \downarrow 3 \text{ strequals } \mathbf{F}[[f_1]] \rightarrow \mathbf{M}[[m_1]] a_1 q \\ & \quad \square \text{nostatechange } a_1 q) \end{aligned}$$

Which checks whether or not the performative name matches f_1 , if it does we get :

$$\lambda q. \mathbf{M}[[m_1]] a_1 q$$

i.e. the state change defined by semantics m_1 when a_1 is the speech act. Otherwise we get :

$$\lambda q. (\text{nostatechange } a_1 q) = \lambda q. (\lambda a. \lambda s. s a_1 q) = \lambda q. (q)$$

This function leaves the state unchanged. The denotation of the semantics for three or more acts can proceed in a similar fashion. The semantic function \mathbf{S} effectively passes on the parameters to the next nesting of the function if the speech act doesn't match.

Persistent Social Relations	
Domain	$v \in \text{Persistent-Soc-Rel} = \text{String-List}$
Operations	
$\text{newpersistentsocrel}$	$\text{Persistent-Soc-Rel}$
$\text{newpersistentsocrel}$	newlist
Conversation Array	
Domain	$a \in \text{Conv-Array} = \text{cid} \rightarrow \text{Conversation-State}$
Operations	
newArray	Conv-Array
newArray	$\lambda i. \text{newconversation}$
accessArray	$\text{cid} \rightarrow \text{Conv-Array} \rightarrow \text{Conversation-State}$
accessArray	$\lambda i. \lambda a. a(i)$
updateArray	$\text{Id} \rightarrow \text{Conversation-State} \rightarrow$
	$\text{Conv-Array} \rightarrow \text{Conv-Array}$
updateArray	$\lambda i. \lambda c. \lambda a. [i \mapsto c]a$
Social State	
Domain	$s \in \text{Social-State} = \text{Persistent-Soc-Rel} \times \text{Conv-Array}$
Operations	
newsocial	Social-State
newsocial	$(\text{newpersistentsocrel}, \text{newArray})$
Current Social State	
Domain	$s \in \text{Cur-Social-State} =$
	$\text{Persistent-Soc-Rel} \times \text{Conversation-State}$
Cases of Speech Acts	
Domain	$c \in \text{Speechactcase} =$
	$\text{Speech-Act} \rightarrow \text{Cur-Social-State} \rightarrow \text{Cur-Social-State}$
Operations	
nostatechange	Speechactcase
nostatechange	$\lambda a. \lambda s. s$
errorstate	Speechactcase
errorstate	$\lambda a. \lambda s. (\text{error}, \text{error})$

Figure 5 Additional semantic algebras.

4.2 Complete Definition

The additional semantic algebras are shown in Figure 5. We now have a new structure: the *Social-State* as described in section 3.1, it contains a dynamic array of conversation states where the speech act's *cid* selects the current conversation. In the evaluation of the language (Figure 6 – only the more important valuation functions are included) we have replaced the *Conversation-State* with the *Social-State*. We have also added the converse function, it is a *Conversation-State* to *Conversation-State* mapping and it looks at the input state to determine the updated contingent social relations for the output state. The converse function is incorporated in the \mathbf{L} evaluation function so that contingent social relations are updated after the semantic functions are processed. The converse function can also stand alone, and this is used if an agent is initiating a new protocol, in which case a state containing only the protocol variable and the string “initiator” is passed into the function and the appropriate initial speech acts are returned.

L: Language $\rightarrow \text{Speech-Act} \rightarrow \text{Social-State} \rightarrow \text{Social-State}$	
$\mathbf{L}[\text{converse-function } \mathbf{C} \text{ protocol-semantics } \mathbf{P} \text{ speech-act-semantics } \mathbf{S}] = \lambda a. \lambda q.$	
let $z =$ $(\mathbf{P}[[\mathbf{P}]] \text{ proterrstate } a ((\mathbf{S}[[\mathbf{S}]] \text{ nostatechange}) a$	
$(q \downarrow 1, \text{accessArray } a \downarrow 5 \ q \downarrow 2)))$	
in let $y = \mathbf{C}[[\mathbf{C}]] (z \downarrow 1, (z \downarrow 2 \downarrow 1, z \downarrow 2 \downarrow 2, \text{newlist}, \text{updateHistory}$	
$a \downarrow 6 \ z \downarrow 2 \downarrow 4)))$	
in let $x = (y \downarrow 1, \text{updateArray } a \downarrow 5 \ y)$ in	
$q \downarrow 2 \downarrow 3$ (“obliged “ $\text{cons } \mathbf{F}[[a \downarrow 1]] \text{ cons } \mathbf{F}[[a \downarrow 3]]$	
$\text{cons ” (“ } \text{cons } \mathbf{F}[[a \downarrow 2]] \text{ cons ”)”) } \rightarrow x$	
$\square (q \downarrow 2 \downarrow 3$ (“permitted “ $\text{cons } \mathbf{F}[[a \downarrow 1]] \text{ cons$	
$\mathbf{F}[[a \downarrow 3]] \text{ cons ” (“ } \text{cons } \mathbf{F}[[a \downarrow 2]] \text{ cons ”)”) } \rightarrow x$	
$\square (q \downarrow 1, \text{updateArray } a \downarrow 5$	
$(\text{update violator } a \downarrow 1 \ q \downarrow 2 \downarrow 1, \text{updatelist } \text{“violation”}$	
$q \downarrow 2 \downarrow 2, \text{newlist}, \text{updateHistory } a \downarrow 6 \ z \downarrow 2 \downarrow 4)))$	
P: Protocol-Semantics $\rightarrow \text{Speechactcase} \rightarrow \text{Speechactcase}$	
$\mathbf{P}[[P_1; P_2]] = \lambda p. \mathbf{P}[[P_2]] (\mathbf{P}[[P_1]]p)$	
$\mathbf{P}[[[T] S]] = \lambda p. \lambda a. \lambda q. \text{cases } (\text{access “protocol” } q \downarrow 2 \downarrow 1) \text{ of}$	
$(\text{isString}(x) \rightarrow x \square \text{isNat}(y) \rightarrow \text{error}) \text{ end}$	
$\text{strequals } \mathbf{T}[[[T]]] \rightarrow \mathbf{S}[[\mathbf{S}]] \text{ errorstate } a \ q \square p \ a \ q$	
C: Converse-Function $\rightarrow \text{Cur-Social-State}_\perp \rightarrow \text{Cur-Social-State}_\perp$	
$\mathbf{C}[[C_1; C_2]] = \lambda q. \mathbf{C}[[C_2]] (\mathbf{C}[[C_1]]q)$	
$\mathbf{C}[[[T] O]] = \lambda q. \text{cases } (\text{access “protocol” } q \downarrow 2 \downarrow 1) \text{ of}$	
$(\text{isString}(x) \rightarrow x \square \text{isNat}(y) \rightarrow \text{error}) \text{ end}$	
$\text{strequals } \mathbf{T}[[[T]]] \rightarrow \mathbf{O}[[O]]q \square q$	

Figure 6 Updated and additional evaluation functions.

4.3 Specifying an ACL with auction protocol

This is an English auction protocol, the price is increased at each iteration until no more bidders are prepared to bid, the last successful bidder being the winner. We use only three speech acts. The main iteration consists of an ⁽¹⁾*announce* of the new price from auctioneer to bidders, and an ⁽²⁾*accept* of a price from bidder to auctioneer. The auction terminates with a ⁽³⁾*declare* from auctioneer to all participants. There are three roles, a member of the role ⁽¹⁾*Bidder* becomes ⁽²⁾*Buyer* when its accept causes the ⁽³⁾*Auctioneer*'s next announce. We assume there is a commonly agreed global timer variable that can be referred to in any conversation.

```

converse-function
[auction]
Initiator:
announce(#Bidder,{protocol=auction,
  item=#item,price=#price,timeout=#timeout});
if #time>#timer then obliged #Auctioneer,
  declare (#Bidder U #Buyer,null);
if intend Auctioneer Auctioneer,sell
  (#item,member of #Bidder,>=#startprice)
  then permitted #Bidder,
  accept(#Auctioneer U #Bidder U #Buyer);
if intend #Buyer
  #Buyer,Buy(#item,#Auctioneer,#price) then
  obliged #Auctioneer, announce(#Bidder U
  #Buyer, price=#price+#increment));
protocol-semantics
[auction]
[announce]
timer:=#time+#timeout;
if #Auctioneer=null then
  (Auctioneer:=S;Bidder:=R);
intend #Auctioneer #Auctioneer,sell
  (#item,member of #Bidder,>=#price);
intend #Auctioneer #Auctioneer,
  (wait(#timeout) then declare (#Bidder U
  #Buyer,{protocol=auction,item=#item,
  price=#price,timeout=#timeout}))
proposal=Buy(#item,#Auctioneer,#price);
desire #Auctioneer know #Auctioneer
  intend #Bidder #proposal
[accept]
if #Buyer!=null then not intend #Buyer
  #Buyer,Buy(#item,#Auctioneer,#price);
Bidder:= #Bidder U #Buyer;Buyer:= S;
Bidder:= #Bidder \ S;oldprice:=#price;
[declare]
p.obliged #Buyer,buy(#item,S,#oldprice);
p.obliged S,sell(#item,#Buyer,#oldprice);
terminate;
speech-act-semantics
[announce] C
[declare] C
[accept]
if #proposal!=null then intend S,#proposal;

```

Figure 7 ACL specification containing an auction protocol

Figure 7 shows the ACL specification for an ACL containing this protocol. When rules are applied, variables prefixed with a # are replaced by their actual value. Note in particular the protocol specific semantics specified for the opening announce: the auctioneer expresses the intention to sell the item *item* to a member of the group of *bidders* at a price not less than *price* and to wait until the time *timeout* has elapsed for a response. These italicised variables do not yet exist in the conversation state, the auctioneer must choose values for them. These rich semantics are vital to enable the auctioneer to reason about how the values will be chosen and to enable the receiver to reason about how to update its internal state. For example when choosing a value for *startprice*, the auctioneer must know that *startprice* means the minimum selling price for the item. Note how the context independent semantics of announce and declare are identical, it is only in the context of the protocol that they are given different meanings. This is similar to the way in which many words in natural language are synonymous in some contexts and not in others. The additional protocol specific meaning of declare is the

creation of two persistent social relations: ⁽¹⁾the bidder is committed to buying the item and ⁽²⁾the auctioneer is committed to selling it to him (at the previously announced price). A new accept in the next iteration of the auction revokes the commitments of the previous iteration. The *accept* has a general context independent meaning which is the expression of an intention to perform an action that the intended recipient of the *accept* has previously proposed. This is included to show the kinds of complex intuitive meanings which can be specified in this framework, and which can be given an exact semantics by means of a function which is easily implemented in any procedural language. An animation of the protocol is shown in Table 1. Only the more interesting propositions and variables are shown (to save space).

5. EVALUATION AND CONCLUSIONS

The practical engineering aspects of developing and implementing agent communication languages were considered in the design of this framework. Emphasis was therefore placed on minimising the effort of the developer by separating concerns and facilitating a modular development. In particular we cater separately for the following four aspects of context, i.e. the meaning of an act can depend on:

1. The protocol being followed in the conversation.
2. The relationship between this speech act and the remainder of the discourse, for example, in section 4.3 the semantics of *accept* and the previous *announce* are related.
3. The domain in which the conversation takes place.
4. The status or authority of the participants (via roles).

Separate specifications constitute flexible building blocks which can be re-used for specific solutions. Designers can extend the set of primitive building blocks, adding speech acts and protocols which can be published and used by others for different protocols. Although these elements of meaning have been separated, they are brought together in the semantic evaluation functions and are thereby interlinked. In particular, there is a well defined relationship between the semantics of protocols and the acts they are composed of since protocols are defined using the state created by the individual acts.

The formal semantics and public perspective mean that it is possible for an outside observer to determine when agents are not complying with the specifications. Enforcing compliance can be achieved by: ⁽¹⁾Sentinel agents may be placed in the society to observe interactions and punish offending agents (by evicting or terminating them for example). ⁽²⁾We may introduce notions like politeness and trustworthiness. Agents that consistently violate contingent commitments and speak out of turn may be branded as impolite and possibly ostracised by the society. Similarly, agents that are known to violate persistent commitments may be deemed untrustworthy and may not be offered contracts.

6. FURTHER RESEARCH

With a formal semantics based on externals, verification is possible, but not trivial. This is the subject of ongoing work. When both the ACL specification and the internals are well defined mathematical functions, further functions can be derived which can prove a certain outcome will result, given a certain initial state. We are currently working on developing a design tool and a compiler for ACLs which uses a diagrammatical representation of protocols based on UML diagrams.

Table 1 Animation of a sample Auction

Speech act (<i>sender, receiver, performative, content, conversation-identifier, sequence</i>)	Social State			
	Conversation State (<i>history omitted</i>)			Persistent social relations
	Variables	Propositions	Contingent social relations	
Joe, {Bob,Pat}, announce, {protocol=auction, item=book, price=40, timeout=3}, 6, 1	Auctioneer=Joe Bidder={Bob,Pat} protocol=auction item=book; price=40 timeout=3; timer=37 time=34; proposal=Buy(book,Joe,40)	<i>I</i> Joe Joe, sell(book, member of {Bob,Pat}, >=40) <i>I</i> Joe Joe, (wait(3) then declare({Bob,Pat})) <i>D</i> Joe <i>K</i> Joe <i>I</i> {Bob,Pat} {Bob,Pat}, Buy(book,Joe,40)	permitted Bob, accept({Joe,Pat}) permitted Pat, accept({Joe,Bob})	
Bob, {Joe,Pat}, accept, null, 6, 2	Bidder=Pat Buyer=Bob	<i>I</i> Bob Bob, Buy(book,Joe,40)	obliged Joe, announce({Bob,Pat}, price=40+null)	
Joe, {Bob,Pat}, announce, price=43, 6, 3	time=36; price=43 timer=39	<i>I</i> Joe Joe, sell(book, member of {Pat}, >=43)	permitted Pat, accept(Joe,Bob)	
Bob, {Joe,Pat}, accept, null, 6, 4	Bidder=Bob Buyer=Pat	<i>I</i> Pat Pat, Buy(book,Joe,40)	obliged Joe, announce({Bob,Pat}, price=43+null)	
Joe, {Bob,Pat}, announce, price=46, 6, 5	time=37; price=46 timer=40	<i>I</i> Joe Joe, sell(book, member of {Bob}, >=46)	permitted Bob, accept(Joe,Pat)	
	time=41; timer=40		obliged Joe, declare {Bob,Pat}, null)	
Joe, {Bob,Pat}, declare, null, 6, 6		terminate		obliged Pat, buy(book,Joe,43) obliged Joe, sell(book,Bob,43)

It is apparent that some of the variables that we have placed in the conversation state are really external to the conversation, for example the timer variable. Also there are exogenous inputs that are not speech acts which do have an effect on the social state. Communication is but one component of social interaction and it is dependent on other components. We have already made this explicit by extending the conversation state to a social state which subsumes it. The next step is to do the same for the inputs that effect changes in this state, (i.e. speech acts are only one type of input) and also for the specification which defines the effect of actions on the state of the society (i.e. the ACL specification is but one component of this).

In human societies characteristics such as being helpful are typically determined uniformly by different observers (because there are certain conventions of society by which behaviour is judged) i.e. it is not a subjective opinion, it becomes public knowledge. We would like the social state to include such public information. For this purpose we intend to make a specification of the rules governing the conventions of social behaviour, and to extend the language evaluation function (maps a speech act onto a change in the social state) so that it can infer the changes in the states of these high level parameters. An explicit representation of these high level concepts would allow us to specify a domain constraint which requires that agents are helpful, and to enforce sanctions on those who violate the constraint.

We have proposed a development method where the rules agents use to update their internal states are specified at design time by a human. In practice this may be difficult or impossible. Agents may move to new domains and learn new protocols via published specifications. Such specifications will not include *add* and *select* functions, nor should they, since different agents should have the freedom to choose their own strategies. This is analogous to a human who has not been to an auction before, after reading the rules, the individual can decide upon a strategy for participating.

If we consider that an agent has some rationality, and can access formal specifications of protocols, there should be a way for the agent to automatically produce a strategy of its own. For this reason we are working on specifications of meta-rules that an agent can use to decide how to update its internal state in a new protocol. This research has once again highlighted the importance of rich semantics capturing the intuitive meaning of acts so that an agent can reason intelligently to interpret and respond to them.

7. REFERENCES

- [1] Artikis, A., Guerin, F., Pitt, J. Integrating Interaction Protocols and Internet Protocols for Agent-Mediated E-Commerce, In Dignum, F., Cortes, U., *Agent-Mediated Electronic Commerce III*. Springer-Verlag. (2001).
- [2] FIPA, [OC00003] FIPA 97 Part 2 Version 2.0: Agent Communication Language Specification. Foundation for Intelligent Physical Agents. (1997).
- [3] Gazdar, G. Speech Act Assignment. *Elements of Discourse Understanding*. Cambridge University Press. (1981).
- [4] Genesereth, M., Ketchpel, S., Software Agents. *Communications of the ACM*, (July 1994).
- [5] Cohen, P. and Levesque, H. Communicative Actions for Artificial Agents. ICMAS'95. MIT Press, Cambridge, Massachusetts, USA; 1995; p. 65-72. (1995).
- [6] Schmidt, D. A. Denotational Semantics: A Methodology for Language development. Allyn and Bacon Inc. (1986)
- [7] Singh, M. A Social Semantics for Agent Communication Languages. *Proceedings of the IJCAI Workshop on Agent Communication Languages*, Springer-Verlag. (2000).
- [8] Singh, M. Agent Communication Languages: Rethinking the Principles. *IEEE Computer*. vol.31, no.12; p.40-7. (1998).