

Towards a standard for an Agent Communication Language

Yannis Labrou and Tim Finin

Computer Science and Electrical Engineering Department,
University of Maryland, Baltimore County,
Baltimore, MD 21250, USA
{jklabrou,finin}@cs.umbc.edu

Abstract

We discuss the problems of providing semantics for an Agent Communication Language standard and we briefly introduce an approach we have used to describe the semantics of the Agent Communication Language KQML. We argue that our approach adequately addresses or circumvents the aforementioned problems.

Introduction

The building of software agents suggests a new way of viewing existing technologies as tools to build software applications that dynamically interact and communicate with their immediate environment (user, local resources and computer system) and/or the world, in an autonomous (or semi-autonomous), task-oriented fashion. In this paradigm, the agent communication language (ACL) emerges as an indispensable component. The ACL is the medium through which the attitudes regarding the content of an exchange between software agents are communicated; the communication language suggests whether the content of the communication is an assertion, a request, some form of query *etc.* The primitives of such languages are usually thought and/or understood in the context of speech act theory.

We first discuss the problem of defining a standard for agent communication with particular attention to the issue of semantics for an ACL and its idiosyncrasies as opposed to the problem of providing semantics for speech acts. We then briefly introduce our approach and we discuss how it addresses the aforementioned problems. We go on to describe the basic considerations of our approach and we finally provide a simple example of how we have applied it for the semantics of the agent communication language KQML. We finally discuss our approach in the context of similarly intended research.

The Problem of a Standard for Agent Communication

The communication language has been an integral part of numerous multi-agent systems. But more often than not, the communication language is customized to the application environment and its assumptions. Whether it is the underlying agent theory, or the esoteric interaction protocols the agents follow, or the subtleties of the domain, such communication languages have primitives whose meaning is confined within the boundaries of the particular multi-agent system. Even when the communication primitives are supplemented with formal semantics, such formalisms are rarely useful for some other multi-agent system that introduces its own “agent universe.”

Over the past few years researchers, practitioners and standards’ bodies have grappled with the problem of a standard language for agent communication. Knowledge Query and Manipulation Language (KQML) and the Agent Communication Language proposal of the Foundation for Intelligent Physical Agents (FIPA) dominate the debate.¹ Both, provide a set of primitives (called *performatives* in the case of KQML) which allow agents to communicate attitudes regarding the content of the exchange to other agents and find other agents suitable to process their requests. Although they are superficially similar (at the syntactic level), they suggest substantially differing views on the issue of agent communication.

The task of defining a common ACL involves a slew of issues that have a pragmatic nature but one of the thornier ones is the issue of a clearly defined semantics for a language of communication acts. The issue of semantics for communication acts has received a fair share of attention. Cohen and Lesveque have introduced a model for rational agents (Cohen & Levesque 1990), which uses a *possible-worlds* formalism, that can

¹For the FIPA ACL proposal, as of August 1997, check <http://drogo.csel.t.stet.it/fipa/spec/f7612.html>

in turn be used as a substrate for the semantic description of illocutionary acts (Cohen & Levesque 1995; Smith & Cohen 1996). Sadek (Sadek 1992) has also taken on the similar task of defining rational agency and defining communicative acts on top of it. Finally, Singh proposes a model of agency (Singh 1993a), which differs from that of Cohen and uses it as a framework for the semantic treatment of speech acts (Singh 1993b).

Such approaches present two problems when used to describe the semantics of a common ACL: (1) they are tied to a specific agent theory that might not be applicable to all agents that want to use the ACL, and (2) they introduce complex formalisms that have no bearing to the implementation of agent systems. There is always of course the solution of treating such semantic descriptions as references that formally describe the semantics of the communication primitives and let the agent designers adhere to the semantics simply by declaring compliance even though there is no way to test the actual code's conformance to the semantic description.

Our view on the semantics of an ACL

We propose (Labrou 1996; Labrou & Finin 1997) an approach that attempts to circumvent the aforementioned problems and we have used it to provide the semantics of KQML. We treat KQML performatives as speech acts. We adopt the descriptive framework for speech acts and particularly illocutionary acts suggested by Searle (Searle 1969; Searle & Vanderveken 1985). The semantic approach we propose uses expressions that suggest the minimum set of preconditions and postconditions that govern the use of a performative, along with conditions that suggest the final state for the successful performance of the communication acts; these expressions describe the agents' states which are relevant to the exchange and use propositional attitudes like *belief*, *knowledge*, *desire*, etc. (this *intentional description* of an agent is only intended as a way of viewing the agent).

The preconditions/postconditions framework of our approach suggests a more operational view of the meaning of the performatives which we believe is more useful to implementors that have to provide the code that processes the communication primitives. Also, in order to achieve agent-theory independence, we provide no formal semantics (in a *possible-worlds* formalism or some similar framework) for the modal operators but we restrict the scope and use of these operators, so that they can be easily subsumed by similar modalities whose semantics could be provided by an intentional theory of agency. Apart from the complexity

of possible-worlds-like formalisms which can be prohibiting for the intended audience of a common ACL we want to avoid a tight coupling with a particular theory of agency. The latter concern seems to be shared by the community that works on the FIPA ACL proposal.

Semantics for an ACL standard

Providing semantics is the process of ascribing meaning and we claim that describing the state of an agent before sending a particular message and after receiving it, is a useful basis for ascribing meaning to the communication primitives. The two-step process is: (1) what agents' states we want to capture, and (2) how to describe these agents' states.

What constitutes the semantic description

We argue that the following constitutes a useful semantic description for each of the performatives:

1. A natural language description of the performative's intuitive meaning.
2. An expression that describes the content of the illocutionary act. For all practical purposes, this is a formalization of the natural language description.
3. Preconditions that indicate the necessary state for an agent in order to send a performative (**Pre(A)**) and for the receiver to accept it and successfully process it (**Pre(B)**). If the preconditions do not hold a *error* or *sorry* will be the most likely response.
4. Postconditions that describe the states of both interlocutors after the *successful* utterance of a performative (by the sender) and after the receipt and processing (but before a counter utterance) of a message (by the receiver). The postconditions (**Post(A)** and **Post(B)**, respectively) hold unless a *sorry* or an *error* is sent as a *response* in order to suggest the unsuccessful processing of the message.
5. A completion condition for the performative (**Completion**) that indicates the final state, after possibly a conversation has taken place and the intention suggested by the performative that started the conversation, has been fulfilled.
6. Any comments that we might find suitable to enhance the understanding of the performative.

We do not suggest that establishing the preconditions of a performative guarantees its successful execution and performance. We merely suggest that the preconditions indicate what can be assumed to

be the state of the interlocutors involved in an exchange. Similarly, the postconditions are taken to describe the states of the interlocutors assuming the successful performance of the communication primitive. The next task is that of introducing a language to describe agents' states.

How to describe agents' states

We use expressions in a meta-language to formally define (cognitive) states for agents and use them to describe the performative, the preconditions, postconditions and completion conditions associated with the use of a particular performative (speech act). In these expressions we use operators that stand for propositional attitudes and have a (informal) reserved meaning. The operators we use are:

1. BEL, as in BEL(A,P), which has the meaning that P is (or can be proven) true for A . P is an expression in the native language of agent A ².
2. KNOW, as in KNOW(A,S), expresses knowledge for S , where S is a state description (the same holds for the following two operators)
3. WANT, as in WANT(A,S), to mean that agent A desires the cognitive state (or action) described by S , to occur in the future.
4. INT, as in INT(A,S), to mean that A has every intention of doing S and thus is committed to a course of action towards achieving S in the future.

We also introduce two instances of actions:

1. PROC(A,M) refers to the action of A processing the KQML message M . Every message after being *received* is *processed*, in the sense that it is a valid KQML message and the piece of code designated with processing the performative for the application indeed processes it. PROC(A,M) does not guarantee proper processing of the message (or conformance of the code with the semantic description).
2. SENDMSG(A,B,M) refers to the action of A sending the KQML message M to B .

The argument of BEL is an expression P in the agent's implementation language. BEL(A,P) if and only if P is true (in the *model-theoretic* sense) for agent A ; we do not assume any axioms for BEL. Roughly, KNOW, WANT and INT stand for the psychological states of knowledge, desire and intention,

²The *native* language of the application may or may not have modal operators but in our analysis we do not assume any.

respectively. All three take an agent's state description (either a cognitive state or an action) as their arguments. An agent can KNOW an expression that refers to the agent's own state or some other agent's state description if it has been communicated to it. So, KNOW(A,BEL(A,"foo(a,b))) is valid, as is KNOW(A,BEL(B,"foo(a,b)")), if BEL(B,"foo(a,b)") has been communicated to A with some message, but KNOW(A,"foo(a,b)") is not valid because "foo(A,B)" stands for an expression in the agent's knowledge store and not for a state description.

Researchers have grappled for years with the problem of formally capturing the notions of *desire* and *intention*. Various formalizations exist but none is considered a definitive one. We do not adopt a particular one neither we offer a formalization of our own. It is our belief that any of the existing formalizations would accommodate the modest use of WANT and INT in our framework. For this reason we restrict the scope of the operators. We allow conjunctions (\wedge) and disjunctions (\vee) of expressions that stand for agents' states (the resulting expressions represent agents' states, also), but we do not allow \wedge and \vee in the scope of KNOW, WANT and INT. Propositions in the agent's native language can only appear in the scope of BEL and BEL can only take such a proposition as its argument. BEL, KNOW, WANT, INT and actions can be used as arguments for KNOW (actions should then be interpreted as actions that have already happened). WANT and INT can only use KNOW or an action as arguments. When actions are arguments of WANT or INT, they are actions to take place in the future.

A negation of a mental state is taken to mean that the mental state does not hold in the sense that it should not be inferred (we will use the symbol not). When \neg qualifies BEL, *e.g.*, \neg (BEL(A,X)), it is taken to mean that the :content expression X is not true for agent A , *i.e.*, it is not provable in A 's knowledge base. Obviously, what "not provable" means will depend on the details of the particular agent system, for which we want to make no assumptions.

A simple example that showcases our approach

In (Labrou 1996) we provide the semantics for all the KQML performatives. We only present the case of *tell* to showcase our approach.

tell(A,B,X)

1. A states to B that A believes the content to be true.
2. BEL(A,X)

3. **Pre(A)**: $BEL(A,X) \wedge KNOW(A,WANT(B,KNOW(B,S)))$
Pre(B): $INT(B,KNOW(B,S))$
 where S may be any of $BEL(B,X)$, or $\neg(BEL(B,X))$.
4. **Post(A)**: $KNOW(A,KNOW(B,BEL(A,X)))$
Post(B): $KNOW(B,BEL(A,X))$
5. **Completion**: $KNOW(B,BEL(A,X))$
6. The completion condition holds, unless a *sorry* or *error* suggests B's inability to acknowledge the *tell* properly, as is the case with any other performative.

Objections may be raised regarding the meaning we chose to attribute to *tell*. Our semantics suggest that an agent can not offer unsolicited information to some other agent. This can be easily amended by introducing another performative, let us call it *proactive-tell* which has the same semantic description as *tell* with the following difference: **Pre(A)** is $BEL(A,X)$, and **Pre(B)** is empty. Following KQML's tradition of an open standard, the KQML users' community should decide the performative names to be associated with whatever semantic description.

Brief discussion

By attempting a semantics for communication acts without a theory of agency, *i.e.*, formal semantics for the propositional attitudes (operators), we certainly give up interesting inferencing. For example, if some agent A wants to inform another agent B about its belief regarding a proposition A, *e.g.*, by sending the KQML performative **tell(A,B,X)**, and later commits the act **tell(A,B,X \rightarrow Y)**, B will not be able to infer that A believes Y since we do not even want to assume a *KD45* model for belief. We see this rather as an advantage. Why should it be the case that B infers that $BEL(A,Y)$? Agent A might be a simple agent with no inferencing or might not want to be held liable to other agents about its own private inferencing. Even the simplest assumptions of current agent theories present constraints that might not be to everybody's liking. When defining a standard, accommodating the widest possible audience becomes a primary concern and we have tried to take this concern into consideration in our approach. So, we suggest that the additional information of the agent theory that holds for the agent can be supplied as an additional parameter in the exchange and subsequently taken into consideration for further inferencing. In the end, we trade a formal semantics for the propositional attitudes (which inevitably defines a *model of agency* that is unlikely to be universal for all agents) for a simpler formalism and agent theory independence.

In Conclusion

We argued that providing semantics for a standard agent communication language requires agent theory independence. We suggest an approach for ascribing meaning to the communication primitives of an ACL which although it makes reference to an intentional agent description limits the scope of the used modalities so that they can be subsumed by any intentional agent theory. Moreover the preconditions/postconditions framework we suggest has an operational flavor which, as we claimed facilitates the process of producing the code for processing the communication primitives.

References

- Cohen, P. R., and Levesque, H. J. 1990. Intention is choice with commitment. *Artificial Intelligence* 42:213–261.
- Cohen, P. R., and Levesque, H. 1995. Communicative actions for artificial agents. In *Proceedings of the 1st International Conference on Multi-Agent Systems (ICMAS'95)*. AAAI Press.
- Labrou, Y., and Finin, T. 1997. Semantics and conversations for an agent communication language. In *Proceedings of the Fifteenth International Joint Conference on Artificial Intelligence (IJCAI-97)*.
- Labrou, Y. 1996. *Semantics for an Agent Communication Language*. Ph.D. Dissertation, University of Maryland, Baltimore County.
- Sadek, M. 1992. A study in the logic of intention. In *Proceedings of the 3rd Conference on Principles of Knowledge Representation and Reasoning (KR'92)*, 462–473.
- Searle, J., and Vanderveken, D. 1985. *Foundations of illocutionary logic*. Cambridge, UK: Cambridge University Press.
- Searle, J. R. 1969. *Speech Acts*. Cambridge, UK: Cambridge University Press.
- Singh, M. 1993a. A logic of intentions and beliefs. *Journal of Philosophical Logic* 22:513–544.
- Singh, M. 1993b. A semantics for speech acts. *Annals of Mathematics and Artificial Intelligence* 8(I-II):47–71.
- Smith, I. A., and Cohen, P. R. 1996. Toward a semantics for an agent communications language based on speech-acts. In *Proceedings of the 13th National Conference on Artificial Intelligence*. AAAI/MIT Press.