

Semantics of Agent Communication Languages for Group Interaction

Sanjeev Kumar, Marcus J. Huber, David R. McGee, Philip R. Cohen, Hector J. Levesque*

Department of Computer Science and Engineering
Oregon Graduate Institute
20000 NW Walker Road
Beaverton, OR 97006, USA
{skumar, marcush, dmcgee, pcohen}@cse.ogi.edu

*Department of Computer Science
University of Toronto
Toronto, Ontario
M5S 3H5, Canada
hector@cs.toronto.edu

Abstract

Group communication is the core of societal interactions. Therefore, artificial agents should be able to communicate with groups as well as individuals. However, most contemporary agent communication languages, notably FIPA and KQML, have either no provision or no well-defined semantics for group communication. We give a semantics for group communication that we believe can profitably enrich the agent communication languages. In our semantics, individual communication is a special case of group communication wherein each communicating group consists of a single agent. One of the novel features of this semantics is that it allows senders to send messages even without knowing all the potential recipients of those messages – a typical scenario in broadcast communication.

Motivation

Artificial as well as human agents not only interact with individual agents, but they also need to communicate with groups of agents. We post messages to mailing lists and notice boards; participate in teleconferences and videoconferences; publish web pages and books; speak in meetings and classrooms; talk on radio and television; and advertise on pamphlets and banners. Agents will be assuming some of these responsibilities from humans and will therefore, need to be able to reason and communicate about group concepts. Moreover, in open multi-agent systems, where agents come and go dynamically, it will become ever more prevalent that agents will not know exactly to whom they are sending information or from whom they are requesting aid. These are compelling reasons to investigate developing support for group communication in multi-agent systems. It is no surprise, therefore, that a large number of distributed software systems inevitably use some incarnation of broadcasting and multicasting.

However, we observe that the major agent communication languages have either no provision or no well-defined semantics for group communication. For instance, in the FIPA ACL, the only way to inform a set of agents is to inform them individually, one at a time. Furthermore, semantics of the FIPA communicative acts imposes the precondition that the sender has certain beliefs about the mental

state of the (known) addressee. Consequently, there is no way to send messages to unknown agents – a typical scenario in broadcast communication.

KQML does offer several primitives, such as broadcast and recruit-all, that have group flavor but these primitives are merely shorthand for a request to do a series of other communicative acts. Proper semantics cannot be given to group requests such as “One of you, please, get me a slice of that pie.” We may safely conclude that support for group communication in the widely used agent communication languages does not exist.

Group communication is not just about sending a message to a large number of agents at the same time. As mentioned earlier, sometimes the sender does not know the specific recipients of a message. A person who posts the notice “Beware of dogs” may not know who will read that message. So the semantics of a communication language should allow for intentions with respect to “whoever gets this message,” while allowing for constraints on the intended recipients and identification of this constraint for correct illocutionary effect. Furthermore, the intended actor for a communication may be a subset of the recipients or a completely different set. By sending an email to the CSE101 mailing list requesting Becker to take the attendance in the next class, the instructor not only made a request to Becker to take attendance but also let the whole class know that she requested Becker to do it. Senders need not only be individuals but can also be groups. An invitation card from John and Betty is actually a request to attend from “them”. Individuals may be viewed as singleton groups. Therefore, the same communication primitives should work both for individual and group communication. We believe that any general-purpose agent communication language should be able to deal with these aspects of communication.

To summarize, we have argued that (1) Agent communication languages should support group communication where communication between individuals is a natural special case; (2) An agent communication language that supports group communication should account for the recipients being unknown, the sender being a group, and the intended actors being different from the recipients; (3) Semantics of an agent communication language should be in terms of group communication.

Constraints on Communication Languages

We believe that the properties of communication in human society should be an essential guiding principle in the design of agent communication languages. These properties are constraints that an agent communication language should address.

- *Addressee Constraint*: An ACL should support communication addressed to individuals as well as to groups. Moreover, a group may have a stable, known membership, as in a mailing list, or its membership may be unknown, as in a radio broadcast addressed to all listeners.
- *Sender Constraint*: An ACL should support communication sent by individuals as well as by groups. Typically, an individual acts on behalf of a group when the sender happens to be a group: for example, the invitation card from a couple, and an official letter from a company.
- *Recipient Constraint*: An ACL should support unintended recipients or over-hearers that are an inevitable part of group communication. For example, anybody may happen to read a notice addressed to CSE101 students on the school notice board. Similarly, an announcement on an airport public announcement system requesting Alfred Hopkins to meet someone at the bookstall may include everybody else who hears the announcement as an over-hearer.
- *Actor Constraint*: An ACL should support intended actors being wholly different from either the intended or the unintended recipients of a message. In most cases, however, the intended actors will be a subset of the intended recipients. For instance, Alfred Hopkins is the only intended actor in the above example.
- *Actor Awareness Constraint*: An ACL should support a requester's ignorance about the intended actors of the request. For example, a teacher should be able to request "all those who have done the homework" to raise their hands, without knowing in advance which students have done the homework.
- *Sender's Awareness Constraint*: An ACL should support a sender's ignorance about the individual members of a recipient group. This is typically the case with radio and television broadcasts, notices and banners, and authoring web pages and journal articles.
- *Recipient's Awareness Constraint*: An ACL should support the ignorance of a recipient about other recipients of the same message. The reader of a newspaper article may not know who else read that article, yet she may be able to make certain inferences about the mental state of others who have read or will be reading the same article.
- *Originator Constraint*: An ACL should support a recipient's potential ignorance of the originator or sender of a message. A sign "Authorized Personnel

Only" may not indicate the author, but it does communicate the appropriate intentions to anybody who reads the sign. Similarly, a note that I discover on the beach may let me make inferences about the intentions of "whoever wrote the note" even if I don't come to know or deduce its author from it.

In a later section, we will present the semantics of a request performative that satisfies these constraints. We note however, that the FIPA'97 specification (FIPA, 1997) supports the actor constraint to some extent.

Preliminaries

We use a modal language with the usual connectives of a first order language with equality, as well as operators for propositional attitudes and event sequences. (BEL $x p$) and (GOAL $x p$) say that p follows from x 's beliefs or choices respectively. (HAPPENS a) and (DONE a) say that a sequence of actions described by the action expression a will happen next or has just happened, respectively. (HAPPENS $x a$) and (DONE $x a$) also specify the agent for the action sequence that is going to happen or has just happened. BEFORE and AFTER are defined using HAPPENS. Knowledge (KNOW $x p$) is defined in the usual manner. Details of this modal language can be found in (Cohen and Levesque, 1990b). An action expression is built from variables ranging over sequences of events using constructs of dynamic logic: $a;b$ is action composition and $p?$ is a test action. Mutual belief is defined in terms of unilateral mutual belief or BMB (Cohen and Levesque, 1990b). However, unlike the previous work, we treat BMB between two agents as a semantic primitive in this paper.

BMB as a Semantic Primitive

We assume a modal structure \mathcal{M} that includes an accessibility relation \mathcal{B}_a for every agent a . We use the usual possible worlds representation where $\omega_1 \mathcal{B}_a \omega_2$ means that world ω_2 is belief accessible by agent a from world ω_1 . In this model, an agent x has unilateral mutual belief with agent y about proposition p if and only if $\forall \omega_1, \omega_2 \dots \omega_n$ such that $\omega_1 \mathcal{B}_x \omega_2, \omega_2 \mathcal{B}_y \omega_3, \omega_3 \mathcal{B}_x \omega_4, \omega_4 \mathcal{B}_y \omega_5, \dots, \omega_{n-1} \mathcal{B}_x \omega_n$ (where \mathcal{B}_a is \mathcal{B}_x if n is even or \mathcal{B}_y if n is odd), p is valid in the model \mathcal{M} in world ω_n . More formally,

$\mathcal{M}, \omega \models (\text{BMB } x y p) \equiv \forall n \forall (\omega, \omega') \in \mathcal{B}[x, y, n] \mathcal{M}, \omega' \models p$
where

$\mathcal{B}[x, y, n]$ is defined inductively by

(1) $\mathcal{B}[x, y, 1] = \mathcal{B}_x$

(2) $\mathcal{B}[x, y, n+1] = \mathcal{B}_x \circ \mathcal{B}[y, x, n]$

This semantics is similar to the semantics of common knowledge given in (Halpern and Moses, 1992). In order to generalize the above modal language concepts for groups of agents, we need a suitable representation for groups.

Representing Groups

Similar to the assumptions of other researchers (Singh, 1993), we treat groups as simply a collection of entities. As such, we can regard a group as being defined by a member-

ship property. This can be captured by a predicate consisting of a free variable that ranges over individuals, and in general, ranges over subgroups as well.

Notation. We will underline the entities that represent groups when we need to emphasize their group status, and use the same symbol without the underline in a functional notation to denote the associated membership predicate. For example, $\underline{\tau}$ is a group having the membership predicate $\tau(z)$ where z is a free variable. An entity without underline can be either an individual or a group.

We introduce the notation $\langle \alpha \rangle$ to denote a formula defined by the following rule:

1. If α is a formula without any term of the form $\underline{\tau}$, then $\langle \alpha \rangle = \alpha$
2. If α is a formula with term $\underline{\tau}$, and z does not appear in α , and $\tau(z)$ is the property predicate that corresponds to $\underline{\tau}$, and $\alpha(z)$ is a formula formed by replacing $\underline{\tau}$ with z in α , then $\langle \alpha \rangle = \forall z. \tau(z) \supset \alpha(z)$

For example,

$\langle \text{BEL } x \ p \rangle = (\text{BEL } x \ p)$, if x is an individual agent.

$\langle \text{BEL } \underline{\tau} \ p \rangle = \forall z \tau(z) \supset (\text{BEL } z \ p)$

$\langle \text{BEL } \tau \ p \rangle$ cannot be further expanded until we know whether τ is an individual or a group.

In case of ambiguity, we will mark the starting angle bracket, and the group term that it applies to, with the free variable in the superscript.

$\langle \langle \text{BEL } \underline{\tau} \ (\text{BEL } x \ \langle \text{BEL } \underline{\tau} \ p \rangle) \rangle \rangle = \forall y \tau(y) \supset (\text{BEL } y \ (\text{BEL } x \ \forall z. (\tau(z) \supset (\text{BEL } z \ p))))$

If τ represents an individual agent, say x , then the superscript is dropped in the expansion.

$\langle \langle \text{BEL } \tau \ p \rangle \rangle = \langle \text{BEL } \tau \ p \rangle = (\text{BEL } \tau \ p) = (\text{BEL } x \ p)$

Sometimes, groups need to be treated as meta-agents with agent-like properties and not as a list of individuals. This distinction is discussed in the section on group action. In this case, the membership predicate will not be specified, the term representing this group will not be underlined, and the group will be treated as an individual agent.

Group Beliefs

Our semantics of group communication primitives based on speech acts deals with group beliefs. The simplest case is to consider the beliefs of all the members of a group when talking about group beliefs. The beliefs of more complex groups such as hierarchically composed organizations and institutions (Werner, 1989) can then be expressed in terms of the beliefs of an abstract group consisting of certain roles in that organization or institution.

Group Belief. Group belief may be defined in several ways, including *inclusive belief*: A group $\underline{\tau}$ believes p if all the individuals or the sub-groups that constitute the group believe p .

$$\begin{aligned} (\text{BEL } \underline{\tau} \ p) &\equiv \langle \text{BEL } \underline{\tau} \ p \rangle \\ &= \forall z \tau(z) \supset (\text{BEL } z \ p) \end{aligned}$$

For example, “the students of CSE101 believe p ” can be represented by

$$\begin{aligned} (\text{BEL } \underline{\text{StudentsOfCSE101}} \ p) &\equiv \\ \forall z (\text{student } z \ \text{CSE101}) &\supset (\text{BEL } z \ p) \end{aligned}$$

assuming that the domain membership predicate (student z CSE101) is defined.

Other possible definitions of group belief may include (1) *extensive belief*—mutual belief among all the constituents (individuals or sub-groups) of a group, (2) *existential belief*—belief by at least one constituent of a group, (3) *majority belief*—belief by a majority in a group, and (4) *extensive majority belief*—mutual belief among a majority in a group. For the purpose of this paper, we will only use inclusive belief (as also is done in Singh, 93).

Group BMB. An entity τ_1 has unilateral mutual belief about a proposition p with another entity τ_2 , when τ_1 believes that there is mutual belief between itself and τ_2 about p . It is possible to define different variations of group BMB corresponding to the various types of group beliefs mentioned above. For inclusive beliefs that we assume in this paper, we define four different categories of BMB.

- 1) *Unilateral Mutual belief between two individuals*: This is the degenerate case in which the two groups happen to be singleton groups. The semantics of $(\text{BMB } x \ y \ p)$ has been given in a previous section. The semantics of all other cases will be expressed in terms of the semantics of this base case.
- 2) *Unilateral Mutual belief between an individual and a group*: Agent x has unilateral mutual belief about proposition p with every member of group $\underline{\tau}$ separately.
 $(\text{BMB } x \ \underline{\tau} \ p) \equiv \langle \text{BMB } x \ \underline{\tau} \ p \rangle$
 $\equiv \forall z \tau(z) \supset (\text{BMB } x \ z \ p)$
- 3) *Unilateral Mutual belief between a group and an individual*: Every individual in the group $\underline{\tau}$ has unilateral mutual belief about proposition p with agent x .
 $(\text{BMB } \underline{\tau} \ x \ p) \equiv \langle \text{BMB } \underline{\tau} \ x \ p \rangle$
 $\equiv \forall z \tau(z) \supset (\text{BMB } z \ x \ p)$
- 4) *Unilateral Mutual belief between two groups*: A group $\underline{\tau}_1$ has unilateral mutual belief about proposition p with another group $\underline{\tau}_2$ when everybody in group $\underline{\tau}_1$ has unilateral mutual belief with every member of group $\underline{\tau}_2$ separately.
 $(\text{BMB } \underline{\tau}_1 \ \underline{\tau}_2 \ p) \equiv \langle \langle \text{BMB } \underline{\tau}_1 \ \underline{\tau}_2 \ p \rangle \rangle$
 $\equiv \forall z \tau_1(z) \supset (\text{BMB } z \ \underline{\tau}_2 \ p)$

Group Mutual Belief. Given the above definitions of unilateral mutual belief, the entities τ_1 and τ_2 have mutual belief about proposition p when both τ_1 and τ_2 have unilateral mutual beliefs about proposition p with respect to the other entity.

$$(\text{MB } \tau_1 \ \tau_2 \ p) \equiv (\text{BMB } \tau_1 \ \tau_2 \ p) \wedge (\text{BMB } \tau_2 \ \tau_1 \ p)$$

This is a straightforward generalization of the mutual belief defined for two agents in (Cohen and Levesque, 1990b).

Group Action

Researchers in multi-agent systems have attempted to answer questions such as what it means for a group to do an action (Grosz, B. J. and Kraus, S., 1996). However, we are mainly interested in the meaning of terms such as (HAP-

PENS τa) and (DONE τa) where a is an action expression and τ is a group.

For the purpose of this paper, all we need is to be able to distinguish between (1) a group doing an action as an entity (or meta-agent), and (2) everybody in a list of individuals performing the action. For instance, a request to CSE101 students to move the teacher's desk is a request to the students as a whole. It may entail the CSE101 students deciding which students would do the action of moving the heavy desk and how the individual actions of those students would be coordinated. On the other hand, a request to everybody in CSE101 to submit the homework is a request to every student in the class to submit their homework individually. An agent communication language should be able to properly convey these nuances of a requester's intentions about the performers of an action. We distinguish between these two cases in our semantics by requiring that the group be treated as a meta-agent in the first case – the membership predicate should not be specified. Terms such as (HAPPENS τa) do not decompose further and it is a part of the problem solving process of the group to decide how the group does the action a . The second case requires specification of a membership predicate and terms such as (HAPPENS τa) will be defined as $\langle \text{HAPPENS } \tau a \rangle$. This term expands to $\forall z (\tau(z) \supset (\text{HAPPENS } z a))$ requiring every member of the group τ to do the action a .

Group Extension of Basic Concepts

We adopt an attempt-based semantics (Cohen and Levesque, 1990b) to illustrate the semantics of our group communication performatives. Here we extend the basic semantic concepts using the group formulation developed in the previous sections. The reader may assume any of the definitions for group and organizational beliefs suggested in the previous sections. It is important to note that the definitions to follow allow for both groups and individuals, as τ may either be an individual or a group.

Definition 1. Persistent Goal

$$\begin{aligned} (\text{PGOAL } \tau p q) \equiv & (\text{BEL } \tau \neg p) \wedge (\text{GOAL } \tau \diamond p) \wedge \\ & (\text{KNOW } \tau [\text{UNTIL } [(\text{BEL } \tau p) \vee \\ & (\text{BEL } \tau \square \neg p) \vee \\ & (\text{BEL } \tau \neg q)] \\ & (\text{GOAL } \tau \diamond p)]). \end{aligned}$$

Persistent goal formalizes the notion of commitment (Cohen and Levesque, 1990a). An entity (agent or group) τ having a persistent goal p is committed to that goal. The entity τ cannot give up the goal that p is true in the future, at least until it believes that one of the following is true: p is accomplished, or is impossible, or the relativizing condition q is untrue.

Definition 2. Intention

$$(\text{INTEND } \tau a q) \equiv (\text{PGOAL } \tau [\text{HAPPENS } \tau (\text{BEL } \tau (\text{HAPPENS } a)); a] q)$$

Intention to do an action a is a commitment to do the action knowingly. The entity τ is committed to being in a mental state in which it has done the action a and just prior

to which it believed that it was about to do the intended action next (Cohen and Levesque, 1990a).

Definition 3. Attempt

$$\begin{aligned} (\text{ATTEMPT } \tau e p q t) \equiv & t?; [(\text{BEL } \tau \neg p) \wedge \\ & (\text{GOAL } \tau (\text{HAPPENS } e; \diamond p?)) \wedge \\ & (\text{INTEND } \tau t?; e; q? (\text{GOAL } \tau (\text{HAPPENS } e; \diamond p?)))]?; e \end{aligned}$$

An attempt to achieve p via q is a complex action expression in which the entity τ is the actor of event e and just prior to e , the actor chooses that p should eventually become true, and intends that e should produce q relative to that choice. So, p represents some ultimate goal that may or may not be achieved by the attempt, while q represents what it takes to make an honest effort (Cohen and Levesque, 1990b; Smith et. al., 1998).

Definition 4. Persistent Weak Achievement Goal

$$\begin{aligned} (\text{PWAG } \tau_1 \tau_2 p q) \equiv & [\neg(\text{BEL } \tau_1 p) \wedge (\text{PGOAL } \tau_1 p)] \vee \\ & [(\text{BEL } \tau_1 p) \wedge (\text{PGOAL } \tau_1 (\text{MB } \tau_1 \tau_2 p))] \vee \\ & [(\text{BEL } \tau_1 \square \neg p) \wedge (\text{PGOAL } \tau_1 (\text{MB } \tau_1 \tau_2 \square \neg p))] \vee \\ & [(\text{BEL } \tau_1 \neg q) \wedge (\text{PGOAL } \tau_1 (\text{MB } \tau_1 \tau_2 \neg q))] \end{aligned}$$

This definition adapted from (Smith et. al., 1996) states that an entity τ_1 has a PWAG with respect to another entity τ_2 when the following holds: (1) if entity τ_1 believes that p is not currently true, it will have a persistent goal to achieve p , (2) if it believes p to be either true, or to be impossible, or if it believes the relativizing condition q to be false, then it will adopt a persistent goal to bring about the corresponding mutual belief with entity τ_2 . PWAG is a basic concept in joint intentions and is used in the definition of request.

A Generalized Communication Primitive

We now present a definition of the request performative with group semantics. This definition is a generalized versions of the individual communication performative defined in (Smith et. al., 1998). The terms α , β , and γ in the following definition can represent either groups or individuals. Here, α is the entity performing the request, β is the recipient (including the "over-hearers") of the request message, and γ is the intended actor.

Definition 5. Request

$$\begin{aligned} (\text{REQUEST } \alpha \beta \gamma e a q t) \equiv & (\text{ATTEMPT } \alpha e \phi \psi t) \\ \text{where } \phi = & \langle^z (\text{DONE } \gamma^z a) \wedge \\ & [\text{PWAG } \gamma^z \alpha (\text{DONE } \gamma^z a) \\ & (\text{PWAG } \alpha \gamma \langle^w \text{DONE } \gamma^w a \rangle q)] \rangle \\ \text{and } \psi = & [\text{BMB } \beta \alpha (\text{BEFORE } e [\text{GOAL } \alpha \\ & (\text{AFTER } e [\text{PWAG } \alpha \gamma \phi q]]))] \end{aligned}$$

Substituting for ϕ and ψ in the definition of attempt (definition 3), we get the goal and the intention of the request respectively. The goal of the request is that the intended actor γ eventually does the action a and also has a PWAG with respect to the requester α to do a . The intended actor's PWAG is with respect to the requester's PWAG (towards her) that she does the action a . The requester's PWAG is itself relative to some higher-level goal

q . The intention of the request is that the recipient β believe there is a mutual belief between the recipient and the requester that before sending the request, the requester α had a goal that after sending the request he (the requester) will have a PWAG with respect to the intended actor γ about the goal ϕ of the request.

The recipient β and the intended actor γ never quantify into the beliefs of the requester α - meaning thereby that the requester α does not need to know who β and γ are. Let us consider the general case in which β and γ are groups with specified membership predicate and α could be either a group or an individual i.e. consider (REQUEST $\alpha \beta \gamma e a q$ t). The term $\langle^z(\text{DONE } \gamma^z a)\dots\rangle$ in ϕ expands to $(\forall z \gamma(z) \supset \dots)$ with γ^z replaced by z everywhere. After plugging ϕ into the definition of attempt (definition 3) and simplifying, we get (GOAL $\alpha \dots(\forall z \gamma(z) \supset \dots)$) which means that the requester does not have to know about the members of the group γ . The PWAG conjunct of ϕ has the requester's PWAG as its relativizing condition (PWAG $\alpha \gamma \dots$). However, the γ in (PWAG $\alpha \gamma \dots$) is not specified as γ^z so it does not get replaced by the z that appears in $\langle^z(\text{DONE } \gamma^z a)\dots\rangle$ and hence γ does not quantify into the requester's PWAG as a result of expanding the angle brackets in ϕ . From definition 4, the (PWAG $\alpha \gamma \dots$) expands to terms of the form [(BEL αp) \wedge (PGOAL α (MB $\alpha \gamma p$))]. Expanding the MB in terms of BMB and between two groups, the only relevant term that we get is of the form (PGOAL α (BMB $\gamma \alpha p$) ...). Using the definition of inclusive BMB given earlier, this expression further simplifies to

$$(\text{PGOAL } \alpha [\forall z. \gamma(z) \supset (\text{BMB } z \alpha p)] \dots)$$

where z is a variable that has not been used anywhere else in the expansion of request. Here also, γ does not quantify into the beliefs of the requester α . It is important to note, however, that any other definition of group BMB (such as exclusive BMB) will also not quantify γ into the beliefs of the requester α in the term (PGOAL $\alpha \dots$).

By plugging ψ into attempt, and with similar reasoning we find that the term $(\dots[\text{PWAG } \alpha \gamma \phi q]\dots)$ does not quantify the intended recipient γ into the beliefs and goals of the requester α . Moreover, the term (INTEND...[BMB $\beta \alpha$]) in the expansion of attempt after plugging ψ , never quantifies the recipient β into the beliefs and goals of the requester α , as can be seen by similar expansion and reasoning. Hence, we see that *our definition of request never requires a requester to know who the recipients (both intended and unintended) or the intended actors are.*

We now illustrate examples of usage of this request.

Example 1. A request from one agent x to another agent y . This is the degenerate case in which each of the communicating groups consists of a single agent. The recipient of the message and the intended actor will be the same agent. Using the rules for expanding our macro notation, the above definition reduces to the following:

$$\begin{aligned} (\text{REQUEST } x y e a q t) &\equiv (\text{ATTEMPT } x e \phi \psi t) \\ \text{where } \phi &= [(\text{DONE } y a) \wedge \\ & \quad [\text{PWAG } y \alpha (\text{DONE } y a) \\ & \quad (\text{PWAG } x y (\text{DONE } y a) q)]] \end{aligned}$$

$$\text{and } \psi = [\text{BMB } y x (\text{BEFORE } e [\text{GOAL } x (\text{AFTER } e [\text{PWAG } x y \phi q])])]$$

As expected, this expression is same as the definition of request between two agents in (Smith et. al., 1998) with the exception of BEFORE and AFTER predicates that more precisely describe when the mental states should hold.

Example 2. “All those who have done the homework raise their hands”.

Here, the requester α is a single agent – the teacher. The recipient β is a group—all students in the class. The intended actor γ is also a group—all the students in the class who have done their homework. The action a is “raise hand”. Formally, this request may be expressed as

$$\begin{aligned} (\text{REQUEST } &\text{teacher} \\ &\text{students in class} \\ &\text{students done homework} \\ &e \\ &\text{raise_hand} \\ &\text{homework_due(now)} \\ &t) \end{aligned}$$

Let us assume that the membership predicate for γ i.e. `students_done_homework` is `(doneHomework z)`.

The goal term ϕ in the definition of request expands to the following:

$$\begin{aligned} \phi = \forall z. (\text{doneHomework } z) \supset [& (\text{DONE } z \text{ raise_hand}) \wedge \\ & [\text{PWAG } z \text{ teacher } (\text{DONE } z \text{ raise_hand}) \\ & (\text{PWAG } \text{teacher } \text{students_done_homework} \\ & \langle \text{DONE } \text{students_done_homework } \text{raise_hand} \rangle \\ & \text{homework_due(now)} \\ &]] \end{aligned}$$

The goal part of the request is that every student z that has done the homework eventually does the action of raising her hand. Moreover, the student z should also have a PWAG with respect to the teacher that she (the student z) does the action of raising her hand. Furthermore, this PWAG should be with respect to the teacher's PWAG with “the students who have done their homework” that all students who have done their homework do the action of raising their hands. The intention of the request is to have mutual belief with all students (irrespective of whether or not they have done the homework) in the class about this goal.

Meeting the Constraints. What makes this definition of request uniquely powerful is that it satisfies all the constraints on agent communication languages identified earlier. The addressee and the sender constraints are satisfied because α and β can be groups as well as individuals. The recipient constraint is satisfied because β includes all the recipients—intended as well as unintended. The actor constraint is satisfied because we have a separate term γ for the intended actor. The only place where the recipient β is used in the definition of request is in [BMB $\beta \alpha \dots$]. From the definition of inclusive BMB used in this paper, we see that the members of β do not need to know who the other members of β are. Therefore, the recipient's awareness constraint is supported where it is needed. The originator constraint is satisfied because the requester α does not quantify into the beliefs of the recipient β in the term [BMB $\beta \alpha \dots$]. The most intriguing part of the request definition, however,

is that it even satisfies the actor awareness constraint and the sender's awareness constraints as seen by the following theorems.

Theorem 1: A request can be performed even when the requester does not know about the intended actor. Formally,
 $(\text{Done } \alpha (\text{REQUEST } \alpha \beta \gamma \text{ e a q t})) \wedge \neg \exists z.(\text{BEL } \alpha \gamma(z))$
 is satisfiable.

Proof sketch: Construct a possible worlds model that satisfies both the conjuncts. We use the situation in example 2 to construct such a model. Let the real world w_0 be the world just after the request event has taken place. Let w_1 and w_2 be the worlds that are both belief and goal accessible by the teacher. Let the proposition $(\text{Done } \alpha (\text{REQUEST } \alpha \beta \gamma \text{ e a q t}))$ be true in w_1 and w_2 . Let y_1 and y_2 be two students who have done their homework and hence are the intended actors of the request. Suppose that w_1 is belief accessible from w_0 by y_1 , and w_2 be belief accessible from w_0 by y_2 . The proposition $(\text{doneHomework } y_1)$ is true in w_1 and $(\text{doneHomework } y_2)$ is true in w_2 . However, it is not the case that $\exists z.(\text{BEL teacher } (\text{doneHomework } z))$ because the z in w_1 and w_2 differ. Since w_1 is the only accessible world for y_1 and w_2 is the only accessible world for y_2 , y_1 believes that it has done the homework in w_1 , and y_2 believes it has done the homework in w_2 , because both y_1 and y_2 know that they individually satisfy $(\text{doneHomework } z)$. Therefore, it is possible for the teacher to have the goal that whoever has done the homework be able to evaluate the implication $(\forall z.(\text{doneHomework } z) \supset (\text{DONE } z \text{ raise_hand}) \wedge (\text{PWAG } z \text{ teacher } \dots))$. This is the goal part of the request that we get after plugging ϕ in the attempt. Similarly, using a membership predicate for the class and constructing worlds in which these propositions hold, the intention part of the request can be satisfied. Therefore, $(\text{Done } \alpha (\text{REQUEST } \alpha \beta \gamma \text{ e a q t})) \wedge \neg \exists z.(\text{BEL } \alpha \gamma(z))$ is satisfiable in this model.

Theorem 2: A request can be performed even when the requester does not know everyone who will get the message. Formally,
 $(\text{Done } \alpha (\text{REQUEST } \alpha \beta \gamma \text{ e a q t})) \wedge \neg \exists z.(\text{BEL } \alpha \beta(z))$
 is satisfiable.

Proof sketch: This follows from the proof of the above theorem when a model is constructed to satisfy the intention part of the request.

Discussion

Although there has been considerable work in agent communication languages (FIPA 1997; Labrou, 1997), and researchers, including us, have investigated group intentions and group action (Grosz and Kraus, 1996; Singh 1993), group communication has not been addressed in a comprehensive manner. We believe the present work provides a first step in this direction. We identified a set of constraints for agent communication languages, presented a generalized request performative that can handle both group and individual communication, and showed that this performative is novel in that it satisfied all the identified constraints.

We note that the implementation of an agent communication language and the design of its semantics are two distinct issues. Future work includes the specification and implementation of a complete agent communication language with group semantics. A treatment of roles and responsibilities in teams, organizations, and institutions is also needed for a better understanding of what happens in group-communication in these complex groups. Furthermore, the impact of group communication semantics on communication protocols needs to be investigated.

Acknowledgement

We gratefully acknowledge the support of the DARPA CoABS Program (contract F30602-98-2-0098, A0 G352) for the research presented in this paper and the anonymous reviewers for their constructive suggestions.

References

- Cohen, P. R., and Levesque, H. J. 1990a. Intention is Choice with Commitment. *Artificial Intelligence*, 42(3).
- Cohen, P. R., and Levesque, H. J. 1990b. Performatives in a rationally based speech act theory. In *Proceedings of the 28th Annual Meeting of the Association for Computational Linguistics*, Pittsburgh, Pennsylvania.
- Cohen, P. R., and Levesque, H. J. 1991. Confirmations and Joint Action. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, San Mateo, Calif.
- FIPA Specification. 1997. <http://www.fipa.org/>
- Grosz, B. J. and Kraus, S. 1996. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2).
- Halpern, Y. J., and Moses, Y. 1992. A Guide to Completeness and Complexity for Modal Logics of Knowledge and Belief. *Artificial Intelligence*, 54(3).
- Labrou, Y., and Finin, T. 1997. Semantics and Conversations for an Agent Communication Language, In *Proceedings of the 15th International Joint Conference on Artificial Intelligence*, Nagoya, Japan.
- Singh, M. P. 1993. Intentions for Multi-agent Systems. *MCC Technical Report*, KBNL-086-93.
- Smith, I. A., and Cohen, P. R. 1996. Toward a Semantics for an Agent Communications Language Based on Speech-Acts. *Proceedings of the Annual Meeting of the American Association for Artificial Intelligence*, Portland, Oregon.
- Smith, I.A.; Cohen, P.R.; Bradshaw, J. M.; Greaves, M.; and Holmback, H. 1998. Designing conversation policies using joint intention theory. *International Joint Conference on Multi-Agent Systems*, Paris, France.
- Werner, E. 1989. Cooperating agents: A unified theory of communication and social structure, in *Distributed Artificial Intelligence*, Gasser, L., and Huhns, M.N. (eds.), Morgan Kaufmann, Publishers, Inc., San Mateo, California.